AD-A219 146

PROBLEM SOLVING AND REASONING

Technical Report AIP - 85

James G. Greeno &
Herbert A. Simon

Department of Psychology
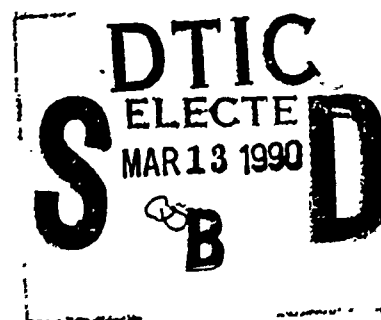Carnegie Mellon University
Pittsburgh, PA 15213

1988

The Artificial Intelligence and Psychology Project

Departments of
Computer Science and Psychology
Carnegie Mellon University

Learning Research and Development Center
University of Pittsburgh

90 03 12 036

PROBLEM SOLVING AND REASONING

Technical Report AIP - 85

James G. Greeno &
Herbert A. Simon

Department of Psychology
Carnegie Mellon University
Pittsburgh, PA 15213

1988

REPORT DOCUMENTATION PAGE

1a. REPORT SECURITY CLASSIFICATION Unclassified		1b RESTRICTIVE MARKINGS		
2a. SECURITY CLASSIFICATION AUTHORITY		3 DISTRIBUTION / AVAILABILITY OF REPORT Approved for public release; Distribution unlimited		
2b DECLASSIFICATION / DOWNGRADING SCHEDULE				
4 PERFORMING ORGANIZATION REPORT NUMBER(S) AIP - 85		5 MONITORING ORGANIZATION REPORT NUMBER(S)		
6a. NAME OF PERFORMING ORGANIZATION Carnegie-Mellon University	6b OFFICE SYMBOL (If applicable)	7a NAME OF MONITORING ORGANIZATION Computer Sciences Division Office of Naval Research		
6c ADDRESS (City, State, and ZIP Code) Department of Psychology Pittsburgh, Pennsylvania 15213		7b ADDRESS (City, State, and ZIP Code) 800 N. Quincy Street Arlington, Virginia 22217-5000		
8a. NAME OF FUNDING / SPONSORING ORGANIZATION Same as Monitoring Organization	8b OFFICE SYMBOL (If applicable)	9 PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER N00014-86-K-0678		
8c ADDRESS (City, State, and ZIP Code)		10 SOURCE OF FUNDING NUMBERS p4000ub201/7-4-86		

PROGRAM ELEMENT NO	PROJECT NO	TASK NO.	WORK UNIT ACCESSION NO
N/A	N/A	N/A	N/A

11 TITLE (Include Security Classification)
Problem Solving and Reasoning

12 PERSONAL
James Yao & Herbert A. Simon

13a. TYPE OF REPORT Technical	13b TIME COVERED FROM 86Sept15 TO 91Sept14	14 DATE OF REPORT (Year, Month, Day) 1988	15. PAGE COUNT 82

16 SUPPLEMENTARY NOTA.,
Appeared in Chapter , In Richard C. Atkinson, Richard J. Herrnstein, Gardner Lindzey &
R. Duncan Luce Steven's Handbook of Experimental Psychology, 2nd ed., Vol. II, New York,
Wiley & Sons, 1988

17	COSATI CODES		18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)
FIELD	GROUP	SUB-GROUP	problem solving, reasoning, cognitive processes, induction, problem representation, computer simulation

19 ABSTRACT (Continue on reverse if necessary and identify by block number)

ABSTRACT

Important advances were made in the 1960s and 1970s in the scientific study
of thinking. They have resulted from new methods for formulating models of
the cognitive processes and structures underlying performance in complex
tasks, and the development of experimental methods to test such models. A
major accomplishment was the discovery of general forms of cognitive
activity and knowledge that underlie human problem solving and reasoning.
This chapter surveys the major theoretical concepts and principles that have
been developed, presents some of the evidence that supports these
principles, and discusses the empirical and theoretical methods.

20. DISTRIBUTION / AVAILABILITY OF ABSTRACT ☐ UNCLASSIFIED/UNLIMITED ☒ SAME AS RPT ☐ OTIC USERS	21 ABSTRACT SECURITY CLASSIFICATION	
22a NAME OF RESPONSIBLE INDIVIDUAL Dr. Alan L. Meyrowitz	22b TELEPHONE (Include Area Code) (202) 696-4302	22c. OFFICE SYMBOL N00014

DD FORM 1473, 84 MAR   83 APR edition may be used until exhausted.   SECURITY CLASSIFICATION OF THIS PAGE
All other editions are obsolete.

PROBLEM SOLVING AND REASONING

James G. Greeno, *Stanford University*

Herbert A. Simon, *Carnegie-Mellon University*

INTRODUCTION

Important advances were made in the 1960s and 1970s in the scientific study of thinking. They have resulted from new methods for formulating models of the cognitive processes and structures underlying performance in complex tasks, and the development of experimental methods to test such models. A major accomplishment was the discovery of general forms of cognitive activity and knowledge that underlie human problem solving and reasoning. This chapter surveys the major theoretical concepts and principles that have been developed, presents some of the evidence that supports these principles, and discusses the empirical and theoretical methods

that are used in this domain of scientific study. This introductory section gives an overview of the major concepts that will be described in the chapter. We discuss relations between these concepts and issues that have been investigated in experimental psychology as well as some general methodological issues.

Overview of Concepts

The concepts that have been developed can be placed in two groups: hypotheses about the form of cognitive action and hypotheses about the form of cognitive representation. The hypotheses about cognitive action extend analyses of behavior that were developed in general behavior theory by investigators such as Thorndike (1923), Tolman (1928); Skinner (1938), and Hull (1943). The hypotheses about representation extend analyses that were developed by Gestalt psychologists such as Köhler (1929), Duncker (1935/1945), Katona (1940), and Wertheimer (1945/1959). One of the important insights reached in the analysis

of problem solving is that hypotheses about these issues of action and representation are complementary; both are necessary components of a theory of human thought.

Form of Cognitive Action

Hypotheses about cognitive action can be considered at two levels: basic action knowledge and strategic knowledge.

A consensus has developed that human knowledge underlying cognitive action can be represented in the form of *production rules*, a formalism introduced by Post (1943) to represent reasoning in mathematics, and adapted for application to psychology by Newell and Simon (1972). Models in which knowledge for action is represented as a set of production rules are referred to as *production systems*.

Any theory of performance must include hypotheses about the process of choice whereby individuals select the actions that they perform. A production system provides a framework for expressing hypotheses about this process in specific detail. A production rule (or, more simply, a *production*) consists of a condition and an action. The *condition* specifies a pattern of information that may or may not be present in the situation. The *action* specifies something that can be performed. The general form of action based on productions is simply: If the condition is true, perform the action.

In a production system, the basic problem of choice among actions is solved by specifying conditions that lead to the selection of each action that can be performed. The condition of each production rule is a *pattern of information* that the system can recognize. These patterns include features of the external problem situation (the stimulus). They also include information that is generated internally by the problem solver and held in short-term memory. The internal information includes *goals* that are set during problem solving. It also can include information in memory, such as past attempts to achieve specific goals. Thus, production rules, which represent basic action knowledge, consist of associations between patterns of information and actions. An action is chosen when the individual has a goal with which the action is associated, and the external stimulus situation as well as information in memory include features associated with the action.

An important component of a model of cog-

nitive activity is its representation of *strategic knowledge*. This includes processes for setting goals and adopting general plans or methods in working on a problem. Models of general problem-solving strategies have been developed to simulate performance in novel problem situations where the individual has little or no experience. One important model is based on a process of means–ends analysis (Newell & Simon, 1972) in which goals are compared with current states, and actions are selected to reduce differences that are identified. General strategies also include processes for setting subgoals when the current goal cannot be achieved directly. Analyses of strategic knowledge in specific domains also have been developed to simulate performance by problem solvers who have received special training (e.g. Greeno, 1978). Strategic knowledge of experienced problem solvers includes global plans for solving classes of problems and knowledge of subgoals that are useful in classes of problem situations.

The general ideas used in formulating hypotheses about cognitive activity in production systems build upon the concepts developed and used in general behavior theory, particularly the formulations of Tolman (1928) and the later forms of Hull's (1952) theory. Early expositions of behavior theory emphasized the direct relations between stimuli and responses, with rather deliberate inattention to intervening events in the brain. Thorndike (1923) emphasized that actions are chosen because of their associations with stimulus conditions. In Skinner's (1938) formulation, actions are performed under the 'control' of external stimulus features. Tolman (1928), on the other hand, emphasized internal goals and information stored in memory in the determination of response selection. Tolman used such terms as 'means–end expectation' and 'means–end readiness' in referring to these factors. In Hull's theory, concepts of covert anticipatory responses (1930) and incentive motivation (1952) were used. In discussions of problem solving, Maltzman (1955) and Staats (1966) postulated stimulus–response units at different levels of generality. The idea of knowledge about action at different levels is used in more recent formulations of strategic knowledge, especially in hypotheses about planning, some of which we discuss in the sections on well specified problems and on problems of design and arrangement.

The concept of a production rule is consistent with these formulations, and behavior theory, even in the terms used by Watson and Skinner, can be expressed as a system of productions (Millenson, 1967). However, as production rules are used in contemporary information processing theory, they more explicitly emphasize the motivational states and memories of prior experiences that combine with external stimulus conditions to determine response choice. Modern production system models of problem solving and similar cognitive processes may be viewed as an (lengthy) extrapolation of Tolman's research program whereby the roles of external environment (stimulus) and inner environment (motivational states and memory contents) as determinants of respose are symmetrical. It also makes exactly how those two sources of information control responses much more explicit. We characterize the extrapolation as lengthy because it postulates not only that many of the essential components of the stimulus lie in the brain, but also that a large part of the response to a production (or all of it) may be internal—consisting, for example, of a change in content of short-term (STM) or long-term memory (LTM). We do not want to underestimate the magnitude of the shift in viewpoint, but we do emphasize that it is a continuous development from the experimental psychology that preceded it. That is presumably what Miller, Galanter, and Pribram (1960) meant when they described the new approach (half jokingly) as 'subjective behaviorism.' Subjective, of course, referred to the minds of the subjects, not to the scientific methods of the investigators.

One major difference between recent hypotheses about cognitive activity and those developed in general behavior theory (in addition to the shift to internal events in behavior) is that recent formulations are much more definite and specific. Models have been formulated as production systems with sufficient specificity to be expressed as computer programs that simulate actual performance of solving specific problems. It is not sufficient to postulate the existence of stimulus–response associations and goals, even at differing levels of generality, to do this. It is necessary also to formulate hypotheses about just what the stimuli, responses, and goals are. Hypotheses about specific structures of knowledge concerning actions and goals in the problem domain must be constructed, and

processes must be designed to recognize specific, relevant patterns of information in the task situation. Hypotheses about strategic knowledge have to specify the conditions in which goals will be set and plans adopted.

Again, we prefer to emphasize continuity in this development. Nothing in the new fine-grained mechanisms is antithetical to the grosser level of description of the earlier theories. In fact, important progress has been made in explaining in detail (sometimes quantitatively) the rich body of experimental data provided within the behavioral framework (Simon & Feigenbaum, 1964; Gregg & Simon, 1967). The impact from achieving this higher level of resolution in our theoretical models and their predictions has led to significantly greater understanding of the psychological processes involved in problem solving and reasoning.

Hypotheses about Representation

Hypotheses about cognitive representations of problems are formulated using the idea of a *problem space*. The problem space includes an individual's representation of the *objects* in the problem situation, the *goal* of the problem, and the *actions* (operators) that can be performed as well as strategies that can be used in working on the problem. It also includes a knowledge of *constraints* in the problem situation—restrictions on what can be done, as well as limits on the ways in which objects or features of objects can be combined.

In developing hypotheses about the representation of problems, much use has been made of concepts developed in analyses of language understanding, including networks of propositions (Anderson, 1976, Kintsch, 1974, Quillian, 1968), procedural representation of concepts (Feigenbaum, 1963; Hunt, Marin, & Stone, 1966, Winograd, 1972), and schemata (Hayes & Simon, 1974; Norman & Rumelhart, 1975, Schank, 1972, Schank & Abelson, 1978). Representations of problems differ from those usually postulated for the understanding of language in that they are constrained to provide information needed for solving the problem. Hypotheses about knowledge used in representing problems include processes for recognizing features that are relevant to actions, strategies, and constraints of the problem domain, and for constructing representations with information that can be used in the cognitive processes of problem solving.

Hypotheses about problem representations address some of the issues of understanding principles and structure in problem solving that were emphasized by some educational, developmental, and Gestalt psychologists (Brownell, 1935; Duncker, 1945; Judd, 1908; Katona, 1940; Köhler, 1929; Piaget, 1952; Wertheimer, 1959). As with hypotheses about cognitive activity, current hypotheses about representation are more definite and specific than those of previous discussions. The hypotheses specify cognitive processes and structures that actually construct representations from texts or other presentations of problem information (Hayes & Simon, 1974; Larkin, McDermott, Simon, & Simon, 1980; Riley, Greeno, & Heller, 1983). Hypotheses about understanding of problem structure and general principles include cognitive structures that specify just what is understood about the problem and how the understanding is achieved (Greeno, 1983; Greeno, Riley, & Gelman, 1984). Another characteristic of recent discussions is that hypotheses about understanding are coordinated with hypotheses about cognitive activity in problem solving, so the significance of understanding, as well as the specific information that it provides for the problem solver, is made clear.

Methodology

The use of computer programming languages as formal systems for psychological theory has been a major factor in the development of the concepts and empirical results discussed in this chapter. The standards that are now common for adequacy of a hypothesis include its expression in a computer program that simulates actual solution of problems—that is, a description of the problem can be given as input for the program, and the program carries out steps that result in the problem's solution. To meet this standard, the theorist must develop specific hypotheses about many aspects of the psychological process that had been unspecified. Representations of specific stimulus situations must be postulated, including relations among cues that are assumed to provide important information for the subject Knowledge structures and processes required for comprehension of stimulus situations must also be specified, leading to specific forms of information that are assumed to constitute the subject's cognitive representations of the stimuli. Assumptions about knowledge in the subject's memory are specified in detail, including associative structures of information and production rules in which specific actions are associated with specific stimulus conditions. The actions include overt responses and internal actions such as setting goals and choosing plans.

To provide evidence to evaluate these more detailed hypotheses, more detailed data are required. A major source of these data has been the increased use of thinking-aloud protocols. These protocols provide a more detailed description of behavior, enabling inferences about intermediate steps such as subgoals and attention to specific aspects of the problem. Protocol statements are treated not as introspective descriptions of psychological processes, but rather as overt reports of mental activity that the subject would be aware of in any case, but usually would not announce. Indeed, subjects are instructed to avoid trying to explain their behavior, but only to give reports of things they notice or think about as they are working (Ericsson & Simon, 1980). Statements in protocols provide data to be explained by models that constitute hypotheses about the process. Thus, protocol statements have the same status as other detailed observations, such as specific patterns of error by individuals on sets of problems, latencies of response when information for problems is presented sequentially, or eye fixations during processing of problem information.

The remainder of this chapter is organized in five sections. "Well-Specified Problems" deals with problems in which a definite goal or solution procedure is specified. "Problems of Design and Arrangement" considers problems in which goals are specified in terms of general criteria, rather than as definite states or procedures. In ' Induction" and "Evaluation of Deductive Arguments" we consider tasks that are often called reasoning, rather than problem solving. Finally, we present conclusions and unifying concepts.

WELL-SPECIFIED PROBLEMS

This section concerns problem solving in relatively well structured situations in which a definite goal is specified. The problem solver is given an initial situation or problem state, a set of operators that can be used to change the situation, and a goal state. The task is to find a

sequence of actions, restricted to use of the permitted operators, that results in the goal state. Problems discussed here include (1) goal-directed problems for which the problem solver has little or no specific knowledge or experience and must resort to what are sometimes called 'weak methods,' (2) solution of problems of the same structure for which individuals have received special training or experience, (3) problems that specify a procedure rather than a goal, and, (4) the representation of problems for which the individual has received special training.

General Knowledge for Novel Problems with Specific Goals

A substantial body of research has been conducted on the solution of well-structured puzzle-like problems that require relatively little domain-specific knowledge. The research strategy of focusing on such problems has some advantages beyond those of making the experiments simpler and the data easier to interpret. In difficult problem domains requiring special knowledge, we are likely to learn from our subjects principally what they know and how they have organized and represented their knowledge in memory, because much of an individual's success depends on whether he or she knows the specific principles and procedures of the domain.

In experiments in domains that are relatively free of specialized content and where subjects are relatively naive, we may still find significant differences in behavior from subject to subject and from domain to domain, but we are also likely to discover some of the commonalities of behavior that characterize problem solving, at least by novices, over a wide range of domains. We are also likely to detect the flexible, general-purpose techniques that people fall back on when they do not have special knowledge or methods adapted specifically to the task at hand. These fallback techniques, often called weak methods, are the only weapons that are available for attacking truly novel problems. Hence, understanding them should contribute to an understanding of discovery processes and creative problem solving.

The *problem space* consists of the problem solver's representation of the materials of the problem, along with knowledge that is relevant

to the task. This includes a representation of the problem goal and operators that can be used. These may be specified in the problem description or supplied by the problem solver's knowledge. The operators include actions that can be performed and conditions that are required for performance of the actions. The problem space also includes the problem solver's strategic knowledge, which may include methods previously acquired through experience in the domain, as well as general problem-solving methods.

The tasks discussed in this section have definite goals specified in the problem instructions. Subjects solving these problems are usually not experienced in the tasks. The problem-solving operators also are specified in the problem instructions, rather than being known in advance by the problem solvers, and the problem solvers must rely on general problem-solving strategies—that is, on weak methods. The principal methods of this kind employ a general problem-solving heuristic called *means–ends analysis*, in which the current state is compared with the goal of the problem or a subgoal that the problem solver is trying to achieve, and an operator is selected that can reduce differences between the current state and the goal.

Research has been conducted on several tasks of this general kind, two of which we discuss here: proof discovery exercises in logic (Newell & Simon, 1972), and water-jar problems (Atwood & Polson, 1976). These studies illustrate two empirical methods. Newell and Simon's study of logic-proof discovery used detailed analyses of thinking-aloud protocols obtained from a few subjects, with data from a larger group of subjects to check the representativeness of some general features of performance. Atwood and Polson's study of water-jar problems used frequencies of responses that occurred during problem solving to evaluate a model of problem solving expressed in quantitative form.

Discovering Proofs in Logic

Discovering proofs for mathematical theorems of one kind or another is a task all of us have faced. One domain in which theorem proving has been studied extensively is elementary symbolic logic (Moore & Anderson, 1954; Newell, Shaw, & Simon, 1957; Newell & Simon, 1972). The propositional calculus is defined by

only two rules of inference and a dozen axioms. In the studies discussed here the task was presented as a syntactic game of transforming strings of uninterpreted symbols according to rules given as symbolic formulas. This ensured that subjects could not draw readily on common-sense knowledge they may have had of the laws of reasoning (The studies of syllogistic reasoning discussed later directly address the question of subjects' knowledge of formal logical rules.)

DEDUCTION AND INDUCTION IN PROBLEM SOLVING

At the outset we must deal with one common misconception about proof-finding tasks. Logic is the science of deductive reasoning from premises to conclusions. A proof is a sequence of expressions starting with axioms (or previously proved expressions) and terminating with the desired theorem; each step of the proof must satisfy the laws of deduction. Its validity can be checked, step by step, by applying those laws systematically.

Finding the proof of a theorem is another matter. We have a known starting point, the axioms, and a known goal, the theorem. However, in most mathematical domains there is no systematic rule for constructing a path from axioms to theorem. That path must be discovered, and the method usually used is to search for it; the amount of trial and error required depending on how selectively the search is carried out. Hence, while a *proof* is an example of a logical deduction, the problem-solving activity involved in searching for a proof is an inductive search.

THE MOORE–ANDERSON LOGIC PROBLEMS

In the logic task designed by Moore and Anderson (1954), subjects were not told that they were discovering proofs in symbolic logic, but were simply instructed to 'recode' certain strings of symbols into other specified strings, using a given set of transformation rules. The rules were displayed on a sheet of paper that was available to the subjects at all times. A typical rule (there were twelve, some with subparts) was.

$$A \vee B \rightarrow B \vee A,$$

which was to be interpreted: The expression $A \vee B$ may be transformed into the expression $B \vee A$, where A and B are variables for which any parts of an expression can be substituted. The connectives in such expressions were referred

to by the experimenter as *wedge* ($\vee$), *dot* ($\cdot$), *horseshoe* ($\supset$), and *tilde* ($\sim$), instead of being given their usual interpretations in logic of *or*, *and*, *implies*, and *not*. Subjects were run on this task by Carpenter, Moore, Snyder, and Lysansky (1961) at Yale, and by Newell and Simon (1972) at Carnegie Institute of Technology.

Several kinds of data can be obtained in problem-solving tasks of this kind. The times to solution can be recorded, as well as the times for making each successive transformation of an expression. Numbers of correct solutions can be counted, and errors can be classified and analyzed.

THINKING-ALOUD PROTOCOLS

The richest data, however, are obtained by instructing subjects to think aloud while solving the problem. The verbal protocols provide a higher temporal density of data than is usually obtained by other methods (except perhaps from records of eye movements). Typically, subjects speak at an average rate of about two words per second, although there are substantial differences among subjects and from one part of a task to another.

In order for thinking-aloud data to be used correctly and effectively to help understand subjects' cognitive processes, answers are needed to several questions, especially: (1) which processes, or what parts of the processes, are verbalized, and (2) to what extent does verbalization alter or in any way affect the problem-solving process itself? A recent extensive review of relevant literature (Ericsson & Simon, 1980) supports three general conclusions. First, subjects mainly verbalize a subset of the symbols that pass through the STM as the task is being performed. The verbalizations are more complete (i.e., give a fuller record of successive STM contents) when the problem is solved in terms of verbal symbols than when the STM contents have to be translated from some other modality (i.e., visual image). Second, the *process* of recognizing some familiar visual or auditory stimulus does not produce any intermediate symbols in STM that can be reported; only the *result* of the recognition process can be reported. Third, in most problem-solving tasks, the cognitive processes are the same in the thinking-aloud as in the silent condition. Moreover, the speed of task performance is generally neither increased nor decreased by the instruction to think aloud.

The protocols under discussion here are those produced by subjects while they are performing the cognitive task. In using retrospective protocols as data, additional factors must be taken into consideration. First, only such information can be reported retrospectively as has been transferred to LTM and retained there. Second, unless the instructions call for recall of specific events, subjects may engage, in a variety of ways, in active reconstruction of the event or process that is being probed. Hence, retrospective protocols must be interpreted in the light of what we know about the laws of memory and forgetting (Bartlett, 1932; Nisbett & Wilson, 1977).

The most detailed analysis of problem-solving protocols calls for reconstructing from them the successive cognitive states of subjects as they work toward solution of the problem. 'Cognitive state' means what the subject knows or has found out about the problem up to the time of the protocol fragment being examined, along with information, such as subgoals and evaluations, that has been generated by the subject from decisions and judgments. Typically, in tasks like the logic-theorem proving task, subjects verbalize the symbolic expressions they produce and those they are actively considering, the operators they are applying to transform expressions, and often the goals they are trying to attain—such as the final theorem or expressions they think would bring them closer to it (Newell & Simon, 1972). As they proceed, subjects often evaluate their progress and the suitability of steps they have just taken.

From such protocol statements we can usually reconstruct the problem space in which a subject is operating. Formally, a *problem space* is defined by a set of symbol structures, corresponding to the cognitive states that can be generated as the subject works on the task, and a set of cognitive operators, or information processes that produce new cognitive states from existing ones. The problem-solving efforts of a subject may be described as searches through a problem space from one cognitive state to another until the solution (a particular cognitive state) is found or the search is abandoned.

Given a description of the problem space inferred from a protocol, a search tree called a *Problem Behavior Graph* (PBG) can be constructed to represent the course of the subject's search. The size and shape of the PBG discloses the extent of the subject's skill and knowledge and the consequent selectivity he is able to achieve. With the PBG, the experimenter can construct a simulation program for a computer which, if given the same problem, would generate the same PBG as that generated by the subject.

The accuracy of fit of the simulation program to the strategy that guides a subject's behavior can be judged by comparing the program's trace step-by-step with the problem-solving protocol. Formal methods for judging goodness of fit in a statistical sense are not available, but departures of trace from protocol are easy to detect. These discrepancies form the basis for modifying the simulation program to fit the protocol more closely. Except for the fact that the data in this case are not numerical, the process of fitting a computer program to protocol data is identical in principle to the process of fitting a system of differential equations to time series data.

A *basic problem space* for the logic task is one in which the subject's cognitive state is defined by the logic expressions thus far derived from the initial given expression, and by the legal operators for generating new expressions from these. Since the protocol normally discloses both what operators are being applied and what expressions are obtained from the application, a great deal of redundancy is contained in the available information, with which the consistency of the interpretation is tested. Many protocols allow a richer problem space to be inferred—one in which the subject notes similarities and differences among logic expressions, and chooses the next step in those terms. When the subject's choice of actions is also guided by goals and subgoals, these are added to the description of the problem space.

SOLUTION PROCESSES

No single strategy, or simulation program based on such a strategy, can be expected to describe the problem-solving behavior of all subjects. However, the behavior of many subjects in tasks like the proving of logic-theorems reveals that a small number of common mechanisms are central features of the problem-solving process. One of the most important of these is means–ends analysis, first introduced into the problem-solving literature by Duncker (1935/1945). Means–ends analysis requires a problem space rich enough to contain not only logic expressions

and operators, but also symbol structures describing differences between pairs of logic expressions and other symbol structures that describe goals. Thus, a subject operating in such a problem space might say, "I have an expression whose main connective is a horseshoe, and my goal expression has a wedge. Let me look for an operator that will change horseshoe to wedge."

In broadest outline, means–ends analysis can be described by the following set of productions, where S is the present state or expression, G is the goal expression, D is a difference between two expressions, and O is an operator:

If the goal is to remove difference D between S and $G \rightarrow$ find a relevant operator O and set the goal of applying it.

If the goal is to apply O to S, and condition C for applying O is unsatisfied $\rightarrow$ set the goal of satisfying C by modifying S.

If the goal is to apply O to S $\rightarrow$ make application

If there is a difference D between S and G $\rightarrow$ set the goal of removing it.

If there is no difference between S and G $\rightarrow$ halt and report problem solved.

While the production system displayed here does not describe all the details of the control of search, it provides the main outlines of means–ends analysis. The system seeks to detect a difference between the present position in the problem space and the goal position. Given such a difference, it searches memory for an operator that is relevant for removing the difference and attempts to apply it. If all the conditions for operator application are not satisfied, it expresses the discrepancy as a new difference and establishes the goal of reducing it. The scheme operates recursively, and when one difference has been removed it looks for another. An important component of the strategy not represented in the productions is the use of memory to store goals that have been tried, so the problem solver can avoid looping through the same cycle of repeated unsuccessful attempts of a goal that cannot be achieved.

A clear distinction can be made between the general strategy of means–ends analysis and domain-specific knowledge that is required for the strategy to be used in solving any particular problem. The general strategy is represented in the productions shown above. To use these productions, a problem solver must be able to represent the state, S, and the goal, G and identify differences between them. In the domain of logic, states correspond to expressions, and differences involve different letters, different connectives, and different arrangements of letters and connectives. The problem solver also must know what operators can be used, what conditions permit each operator to be applied, and what kinds of difference are removed by use of each operator. In logic, the operators are the rules for transforming expressions. The conditions are patterns that are specified in the rules, and the relevant differences for a rule can be inferred by comparing the two sides of the rule. For example, $A \cdot B \rightarrow A$ requires a pattern in which two subexpressions are connected by a dot, and has the effect of removing a letter or a subexpression, as well as removing the dot. $A \supset B \leftrightarrow \sim A \vee B$ does not remove or add any letters, it can be applied to a pattern with a horseshoe to change the horseshoe to a wedge or vice versa, and it changes the sign of one of the letter or subexpressions.

The general strategy of means–ends analysis has been implemented in a program called the *General Problem Solver* (GPS) and shown to be sufficient for providing solutions in over a dozen problem domains, including puzzles such as the Tower of Hanoi and tasks such as integral calculus, given appropriate representations of the states, operators, and connections between operators and differences in the specific domains (Ernst & Newell, 1969).

In the experiments conducted with the logic task, subjects were not experienced in the domain. The operators were presented as part of the task instructions, and it is reasonable to presume that subjects relied primarily on general problem-solving strategies, rather than on knowledge that was specific to that task. If that is correct, and if the subjects' general problem-solving strategies have the properties of GPS, then their performance in the logic task should be similar to that of the program when it is run on the task. The results supported this hypothesis.

KINDS OF EVIDENCE

The hypothesis was evaluated at three levels. First, specific protocols were examined that compared the statements made by subjects with the steps in solutions by specific versions of GPS. For these simulations, GPS was varied by supplying it with differing priorities of differences. Second, a set of protocols [all those obtained by Newell and Simon (1972) on one moderately difficult problem] were coded, and each was translated into a *problem behavior graph* (PBG) showing a succession of cognitive states that was inferred from the statements and problem-solving operators to account for the transitions between states. The state-to-state transitions were classified and the categories were compared with categories of activity performed by GPS. Third, some summary statistics were compiled for Newell and Simon's subjects and for the subjects run at Yale, involving the frequencies of occurrence of several intermediate steps in solutions of the problems. These statistics were compared in order to detect any gross abnormalities in Newell and Simon's data, with the results from a larger group of subjects at Yale who solved the problem with pencil and paper and without the requirement of thinking aloud.

As Table 9.1 illustrates, individual protocols can often be simulated in great detail, but there will undoubtedly be differences among individuals in their problem solving methods, and hence in the production systems that would describe them. For purposes of psychological theory, we are often less interested in the details of a particular simulation (except as a strong test of the theory) than we are in the structure of a program that simulates the main mechanisms revealed in a set of protocols. The problem of averaging over groups of subjects can also be handled formally by comparing the statistics of the behavior of a program with the statistics of the human subjects as a group. This section examines comparisons of programs in detail with individual protocols, and the statistical approach is described in the next section.

INDIVIDUAL PROTOCOLS

Newell and Simon have presented several protocols in which activities of subjects reflect processes similar to those in GPS. The illustration in Table 9.1 shows a segment of one subject's protocol along with a trace of a version of GPS working on the same problem. In the protocol and the GPS trace, L0 refers to the goal expression and L1 refers to the initial expressions of the problem. The expressions L2, L3, and so on refer to additional expressions that are generated by the problem solver by applying operators to L1 and other previously generated expressions. The operators that are referred to in this segment are

$$R6: A \supset B \leftrightarrow {\sim} A \vee B$$

$$R7: A \vee (B \cdot C) \leftrightarrow (A \vee B) \cdot (A \vee C)$$

$$A \cdot (B \vee C) \leftrightarrow (A \cdot B) \vee (A \cdot C)$$

The protocal segment in Table 9.1 began near the end of the first minute of work on the problem and lasted slightly more than three minutes.

In this segment, the goal of both the subject and GPS was to delete the letter R from the initial expression. Both problem solvers considered rule R7 as a possible way to do so. The rule R7 cannot be applied to L1 because its connectives are wrong, so a subgoal was set to change the connective of L1. This led to use of R6, but the two occurrences of R in the transformed expression have opposite signs. When attempts were made to change one of the signs, the horseshoe was returned to the subexpression. At this point the subject and the specific version of GPS that produced this run were unable to continue on this line of work.

This protocol and GPS trace are alike in an impressive degree of detail. However, the important finding is not that the subject and GPS tried to use the same rules in the same sequence. The precise sequence of rules used by GPS can be tailored fairly arbitrarily, and other versions of GPS would not try to use R6 and R7 in this situation. The important finding involves the general character of the subject's performance, involving goals related to differences between the current expression and the problem goal and subgoals to make operators applicable. The protocol provides several clear illustrations of activities that are consistent with the hypothesis of a GPS-like problem-solving process.

PROBLEM BEHAVIOR GRAPHS

It is important to consider whether activities like those in Table 9.1 are typical of problem solvers or are relatively rare. Newell and Simon addressed this question by examining Problem

Table 9.1. Comparison of GPS with protocol data
Source. (Newell & Simon. 1972)

GPS trace	Subject protocol
L0: ~ (~Q · P)	
L1: (R ⊃ ~P) · (~R ⊃ Q)	
Goal 1: Transform L1 into L0	
Goal 2: Delete R from L1	
.	
.	
.	
Goal 2: (reinstated)	Now I'm looking for a way to get rid
Goal 9: Apply R7 to L1	of the horseshoe inside the two
Goal 10: Change connective to ∨	brackets that appear on the left
in left (L1)	and right sides of the equation.
Goal 11: Apply R6 to left (L1)	And I don't see it.
Produce L4:	Yeh, if you apply R6 to both sides
(~R ∨ ~P) · (~R ⊃ Q)	of the equation,
	From there I'm going to see if I can
	apply R7.
	[E writes L2: (~R ∨ ~P) · (R ∨ Q)]
Goal 12: Apply R7 to L4	I can almost apply R7, but one R needs
Goal 13: Change connective to ∨	a tilde. So I'll have to look for
in right (L4)	another rule.
Goal 14: Apply R6 to	I'm going to see if I can change that R to
right (L4)	a tilde R. As a matter of fact, I should
Produce L5:	have used R6 on only the left hand side
(~R ∨ ~P) · (R ∨ Q)	of the equation. So use R6, but only on
	the left hand side.
	[E writes L3: (~R ∨ ~P) · (~R ⊃ Q)]
Goal 15: Apply R7 to L5	Now I'll apply R7 as it is expressed.
Goal 16: Change sign of	Both... excuse me, excuse me, it can't be
left (right (L4))	done because of the horseshoe. So...
Goal 17: Apply R6 to	now I'm looking... scanning the rules
right (L5)	here for a second, and seeing if I
. Produce L6:	can change the R to a ~R in the second
(~R ∨ ~P) · (~R ⊃ Q)	equation, but I don't see any way of
	doing it.
	(Sigh) I'm just sort of lost for a second.
Goal 18: Apply R7 to L6	
Goal 19: Change connective to	
∨ in right (L6).	
Reject	
Goal 16: (reinstated)	
Nothing more	
Goal 13: (reinstated)	
Nothing more	
Goal 10: (reinstated)	
Nothing more	

Note. From Allen Newell, Herbert A. Simon, HUMAN PROBLEM SOLVING, (c) 1972, p. 482. Adapted by permission of Prentice-Hall, Inc., Englewood Cliffs, NJ

Behavior Graphs (PBGs) obtained from the protocols of several subjects working on a moderately difficult problem.

An example of a PBG is shown in Figure 9.1. The numbers prefixed by B on the left, correspond to lines of the transcribed protocol. This PBG was obtained from the protocol that includes the segment given in Table 9.1, which corresponds to the section of the PBG starting at B10 and ending just before B? Information included in the cognitive sta' 'n the rectangles; operators are shown   nes that connect
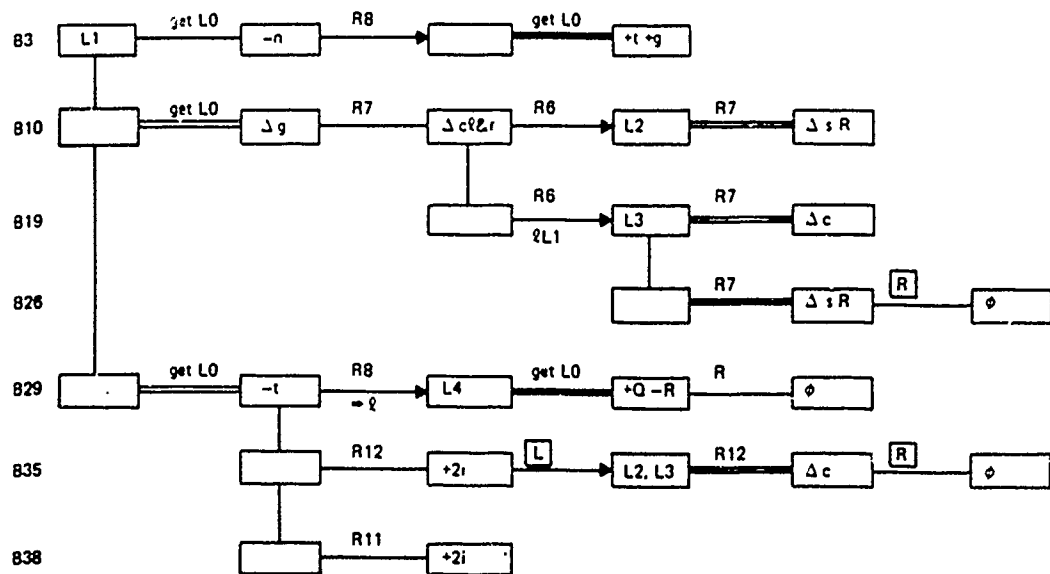
Figure 9.1. Problem behavior graph for a protocol, including the segment in Table 9.1. From Allen Newell, Herbert A. Simon, HUMAN PROBLEM SOLVING, (c) 1972, p. 468. Adapted by permission of Prentice-Hall, Inc., Englewood Cliffs, NJ.

the rectangles. Information in the rectangles refers to new expressions that were written (e.g., L2 or L3 in the protocol), or differences between a current expression and the goal that the subject was considering. For example 'Δg' refers to a difference in grouping of terms and 'Δc ℓ & r' refers to the differences between connectives in the given expression and the goal of applying R7 (horseshoes in both the left and right sides of L1 and wedges or dots needed to apply R7).

Most of the operators refer to the rules; we mentioned R6 and R7 earlier. When a rule is applied successfully, there is an arrowhead on the line between rectangles. A rule shown with a line but no arrowhead indicates that there was a goal of applying the rule but it was not achieved. Double lines indicate repeated attempts to apply rules.

The relation between the protocol and the PBG can be illustrated by examining the first few lines of Table 9.1 and the PBG starting at B10. The instruction "get L0" refers to consideration of the goal, this led to recognition of the difference in grouping between L0 and L1 (Δg). The subject then attempted to apply R7, this led to identifying the differences in connectives noted in the third rectangle (Δc ℓ & r). An attempt to apply R6 was then successful, resulting in line L2. The subject attempted to apply R7

a second time and noticed that there was a difference in the signs of the R terms in the two subexpressions (ΔsR). From time to time, the subject returned to an earlier state, as when he decided that R6 should be applied only to the left side of L1. This is indicated by a vertical line drawn down from the cognitive state that the subject returned to. The rule R6 was applied to the left subexpression of L1, giving line L3; then R7 was attempted again, but the subject noticed the horseshoe, an incorrect connective for R7. The subject returned to the goal of changing the sign of R in expression L2, but the search for an appropriate rule (indicated by R in a box) failed to produce anything helpful.

PBGs were compiled from protocols of seven subjects working on the problem in Table 9.1. The transitions between states were classified, and the categories were compared with activities that occur when GPS works on a problem. The categories, and their frequencies in the seven PBGs, are shown in Table 9.2. Frequencies in the second and third columns of Table 9.2 are for subcategories of the categories in columns to the left. For example, the 258 occurrences of means-ends analysis consisted of 89 steps toward goal objects, 151 steps involving operator applicability, and 18 steps to avoid consequences.

Most of the categories shown in the table

Table 9.2. Total frequencies of occurrences of GPS-like mechanisms in seven protocols

Source: (Newell & Simon, 1972).

Category	Frequency	
Means-ends analysis	258	
towards goal object		89
operator applicability		151
overcome difficulty		143
further specify		5
resolve uncertainty		3
*avoid consequences		18
avoid difficulty		17
prepare desired result		1
Working forward	41	
systematic scan and evaluate		37
input form similarity		3
do something different		1
Working backward	2	
output form similarity		2
Repeated application	230	
·after subgoal		93
to overcome difficulty		58
to further specify		11
to resolve uncertainty		2
to avoid consequences		12
to correct error		8
to process interruption		2
implementation		97
for plan		84
*to command experimenter		13
*review		40
Other	27	
*noticing		6
*repeated application		11
*new application		10
Total	558·	

Note. From Allen Newell, Herbert A. Simon. HUMAN PROBLEM SOLVING, (c) 1972, p. 493. Adapted by permission of Prentice-Hall, Inc., Englewood Cliffs, NJ.

correspond to GPS-like activities. Those that do not are marked with asterisks, accounting for about 18 percent of the transitions in the PBGs. The most interesting discrepancies involved choice of operators to avoid undesirable consequences ('*avoid consequences*'), and the noticing of features of the problem not related to the present goal ('*noticing*'). Simulation of these would require significant additions to the problem-solving processes of GPS. The remaining discrepancies involve activities that relate to the requirement of giving protocols ('*command*

experimenter' and '*review*') or points in the protocol where there was insufficient information to determine whether the transition was related to one of the GPS-like categories ('*other,*' except for those in the subcategory *noticing*).

AGGREGATE FREQUENCIES

The data in Table 9.2 were obtained from a small group of subjects who were required to think aloud as they worked. It is possible that the subjects were atypical, or that the instruction to think aloud caused major distortions in the problem-solving method.

Newell and Simon compared some summary statistics from their subjects with data obtained by Carpenter et al. (1961) at Yale University. The larger number of subjects run at Yale (64) solved the problems with pencil and paper, without thinking aloud. If the data for the Carnegie subjects do not differ from the Yale data in significant ways, then there is evidence that the general characteristics of their problem solving were not caused by individual idiosyncracies, or by the requirement of verbalizing protocols while working on the problems.

The summary statistics involved a division of expressions into categories. Each category consists of an expression from the problem, such as the left subexpression of expression L1, and other expressions that can be formed from it by making minor transformations. Minor transformations for this purpose are those involving rules that change the order of terms, the connectives, or the signs, but do not change the terms in an expression. The data for each group of subjects are the proportions of all the expressions written that fall into the categories. The categories of expressions are listed in the left column. For example, expressions in Class L1 are those that can be formed by applying one of the minor transformations to expression L1 as shown in Table 9.1. The categories used are not arbitrary; they are motivated by the observation that differences that depend on changing the terms in expressions are more difficult to remove, and thus, require higher priority in solving the problems.

Data for the problem in Table 9.1 are shown in Table 9.3. The comparison of the two groups of subjects does not show exact agreement but indicates no major differences in their problem-solving processes. A statistical test shows that the difference between the category frequencies

Table 9.3. Proportions of expressions

Class of expressions	Carnegie (78 expressions)	Yale (519 expressions)
L1	37	.29
Left of L1	.14	.16
Right of L1	12	.22
L0	.24	.17
Other	13	.16

in the two groups was not significant $[\chi^2(4) = 8.86; p > .05]$. (The independence assumption of the chi-square test was not met in these data, since several expressions were written by each subject. However, this would generally make it more likely that a significant difference would be obtained, so the conclusion seems warranted.)

Data are shown in Table 9.4 for a somewhat harder problem, in which the given expression was L1 $= (P \vee Q) \cdot (Q \supset R)$, and the goal was L0 $= P \vee (Q \cdot R)$. Again, the agreement is not exact, but the difference is not large enough to reject the hypothesis that the two sets of responses were produced by a single underlying process $[\chi^2(8) = 15.27, p > .05]$.

PLANNING STRATEGY

A second strategy of broad applicability and wide use that was identified in the logic protocols is *planning*. Its underlying idea is that some gaps between the initial situation and the goal are more important and potentially harder to remove than others. If the problem space is simplified by abstracting the problem expressions—removing from them the less important features—the

Table 9.4. Proportions of expressions

Class of expressions	Carnegie (97 expressions)	Yale (487 expressions)
L1	.33	.28
Extended L1	.02	.04
Left of L1	.14	.19
Right of L1	.14	.15
$(R \vee P)$	.13	.07
$(P \vee Q) \cdot (P \vee R)$	.03	.01
L0	.03	.01
Rule 9	.16	.18
Other	.01	.07

simplified expressions will define a much smaller space through which the search can be conducted expeditiously. If a solution can be found to the simplified problem, the omitted details can be restored and this solution is used as a guide for searching in the original problem space.

To use the planning strategy subjects not only must be able to apply means-ends analysis, but must have enough knowledge of the problem space to be able to distinguish important from unimportant differences between expressions. For example, in the domain of logic, subjects gradually learn that it is easier to change the connectives in logic expressions than to change the letters. The planning space is then a space in which expressions like $(R \supset \sim P) \cdot (\sim R \supset Q)$ are replaced by $(RP)(RQ)$. The sequences of proof steps in the original space, $R \supset \sim P$, $\sim R \supset Q$, $\sim Q \supset R$, $\sim Q \supset \sim P$, $Q \vee \sim P$, $\sim (\sim Q \cdot P)$, becomes the simpler sequence in the planning space, RP, RQ, PQ. The second step of the search in the planning space corresponds to two separate steps in the original space, and the third step in the planning space corresponds to three steps in the original space—a reduction of one-half in the length of the derivation, and of a much larger factor in the amount of search required to find it.

Evidence for planning was obtained in protocols like the following, obtained in a problem with four given expressions: L1 $= P \vee Q$; L2 $= \sim R \supset \sim Q$; L3 $= S$; L4 $= R \supset \sim S$; and the goal: L0 $= P \vee T$. Rule R9, mentioned in the protocol, is $A \rightarrow A \vee X$, a rule for adding a term to an expression.

Well, one possibility right off the bat is when you have just a $P \vee T$ like that the last thing you might use is that R9. I can get everything down to a P and just add a $\vee T$. So that's the one thing to keep in mind.

Well, maybe right off the bat, I'm kinda jumping into it, I maybe can work everything down to just a P; I dunno if that's possible. But I think it is, because I see that steps 2 and 4 are somewhat similar; if I can cancel out the Rs, that would leave me with just an S and Q; and if I have just an S and Q, I can eventually get step 3, get the Ss to cancel out and end up with just a Q; and if I end up with just a Q, maybe the Qs will cancel out; so you see, all the way down the line. I dunno, it looks too good to be true, but I think I see it already.

Water-Jar Problems

Water-jar problems, studied extensively by Luchins (1942), are transformation problems with definite goals, involving a set of three jars of different capacities. In the form studied by Atwood and Polson (1976), the largest jar is full in the initial state, and the goal is for the water to be divided equally between two jars. For example, the capacities may be: jar A, 8 oz; jar B, 5 oz; jar C, 3 oz. Then in the initial state, jar A contains 8 oz of water, and jars B and C are empty. The goal is to have 4 oz of water each, in jars A and B. The problem-solving operators involve pouring water from a source jar into a target jar. Water can be poured into the target jar until it is full if there is enough water in the source jar; water can be poured out of the source jar until it is empty if there is enough room in the target jar. Intermediate actions are not possible.

In the water-jar task, differences between any state and the problem goal consist of discrepancies between the contents of the three jars in that state and the contents that are specified in the goal. Atwood and Polson hypothesized that subjects would judge their progress by combining the discrepancies, forming an overall *evaluation function* for the current state, and would try to select moves that would improve the value of this function. They assumed that the evaluation of a specific state i was

$$e_i = |C_i(A) - G(A)| + |C_i(B) - G(B)|,$$

where $C_i(A)$ and $C_i(B)$ are the actual contents of jar A and jar B in state i and $G(A)$ and $G(B)$ are the contents of jar A and jar B in the goal state. (The contents of jar C are redundant with those of A and B.)

Atwood and Polson formulated a process model, based on the means–ends strategy of attempting to reduce the evaluation to zero. They assumed that at each move, subjects would consider various pouring operations that could be made legally and would try to choose one that would make the evaluation function smaller, or at least would not increase its current value by more than a threshold amount. This strategy differs from the means–ends strategy of GPS in one significant respect: GPS considers all the ways in which the current state and the goal differ and selects a move to reduce the most important of these qualitative differences. Atwood and Polson's model combines the dif-

ferences into a single numerical index—the value of the evaluation function—and tries to reduce that difference by at least a threshold amount. This difference probably does not have a significant effect on predictions of performance in the water-jar task, but there are situations in which strategies based on global evaluations and on individual qualitative differences would lead to significantly different performance.

Atwood and Polson also made specific assumptions about memory capacity; they assumed that a limited short-term memory would hold information about states that would be produced by alternative moves, and that each state reached in solving the problem would be stored in long-term memory with a fixed probability.

The model also specifies a sequence of processes for selecting a move. The sequence includes calculating the evaluation function for alternative moves, storing information about alternatives in STM, recognizing states that have occurred before on the basis of information in LTM, and deciding whether to make a given move under consideration. The assumptions of the model allow for several possibilities: (1) A move may be selected if it leads to an acceptable state (this was assumed to be less likely if the state was recognized as having occurred before); (2) the moves stored in STM may be examined, with a selection of those stored in LTM from previous occurrences; (3) a move may be chosen at random from the set of possible moves; or (4) the subject may decide to return to the initial state of the problem.

Atwood and Polson tested their model with data from human subjects who solved different versions of the problem. Problems were presented at computer terminals, and the moves made by each subject were recorded. The model was implemented as a computer program which was run with various values of the parameters, but because it contains probabilistic processes, it does not produce a single sequence of moves in solving a problem. The model was run many times with each set of parameter values, and its performance was summarized by the average frequency of each of the possible problem states. The parameter values were chosen for which the set of frequencies for two problems (jar sizes of 8, 5, 3 oz. and of 24, 21, 3 oz.) that approximated the frequencies obtained from the human subjects. These parameter values seem quite reasonable. The size of STM was set at three alternative

moves, states reached in the problem were stored in LTM with a probability of .90; and the threshold of acceptability for a new state was set at 1.0 above the value of the current state.

Results of the simulation are shown in Figure 9.2. Each set of predictions was based on running the model 250 times. The data for each problem came from a group of about 40 subjects and were different from the data used to estimate the parameters, for which only one problem (8, 5, 3) was used. The model correctly predicted the order of difficulty of the four problems. For two problems, Figure 9.2a and 9.2b, the detailed predictions of response frequency were not significantly different the data by a statistical test. In the two harder problems, Figure 2c and 2d, the general shapes of the frequency distributions agreed with the data, but the model erred by predicting too many returns to states at the beginning of a path that led to the goal. As Atwood and Polson noted, this defect could be corrected by assuming that the probability of recognizing a previous state depends on the number of times it has been encountered.

Conclusions

Problem solving in situations that are novel to the problem solver, in which a definite goal and a set of legal problem-solving operators are described by the instructions, requires some general problem-solving strategy. In situations of this kind the strategy of means–ends analysis represents the major feature of human problem-solving performance. The evidence discussed in this section consists of individual thinking-aloud protocols and aggregate response frequencies in two tasks. Findings that fit this general pattern have been obtained in a wide range of problem-solving tasks, including puzzles such as the Tower of Hanoi (Anzai & Simon, 1979) and physics textbook problems (Simon & Simon, 1978), which are discussed below in "Problems with Specified Procedures."

Means–ends analysis is perhaps the single most important strategy that people employ to search selectively through large problem spaces. The selectivity is powerful because it points search in the direction of the goal, selecting operators on the basis of their relevance to reducing the distance from that goal. Use of means–ends analysis requires some domain-specific knowledge; for example, it can be employed efficiently only if the subject has

learned enough about the problem domain to associate particular differences with particular operators that remove them. However, it is basically a weak method, applicable in situations where the problem solver has little specific knowledge based on experience in the problem domain.

Domain-Specific Knowledge for Familiar Problems with Specified Goals

We now turn to problems solved by individuals who have specialized knowledge, acquired either through instruction or practice. The first subsection concerns problem solving in a domain of school mathematics—high school geometry. We will then discuss problem-solving set or Einstellung, which we interpret as resulting from domain-specific knowledge structures.

Geometry Exercises

In school subjects such as geometry, the knowledge for solving problems is imparted intentionally, through instruction. Research conducted by Greeno (1978) had the goal of investigating and characterizing the knowledge that is acquired by students who learn successfully in the course.

The main data were obtained in a series of interviews conducted weekly with six students who were taking a standard high school geometry course. In each interview, an individual student worked for about 20 minutes to solve three or four problems. Most of the problems were typical of homework or test problems that the class was working on at the time. Students were asked to think aloud as they worked, and their protocols were recorded and transcribed.

One of the problems solved in an early session (during the second month of the course) is shown in Figure 9.3. The problem as it was presented is shown in Figure 3a. The upper right diagram (in Figure 3b) provides notation for referring to the various angles in diagram (3a). The seven steps shown below the diagrams are a formal solution with inferences and justifying reasons. The students were not required to write the solution steps of this problem formally but they were required to state aloud the intermediate inferences they made. Most of the students solved the problem in Figure 9.3 correctly. Specific aspects of their solutions are discussed
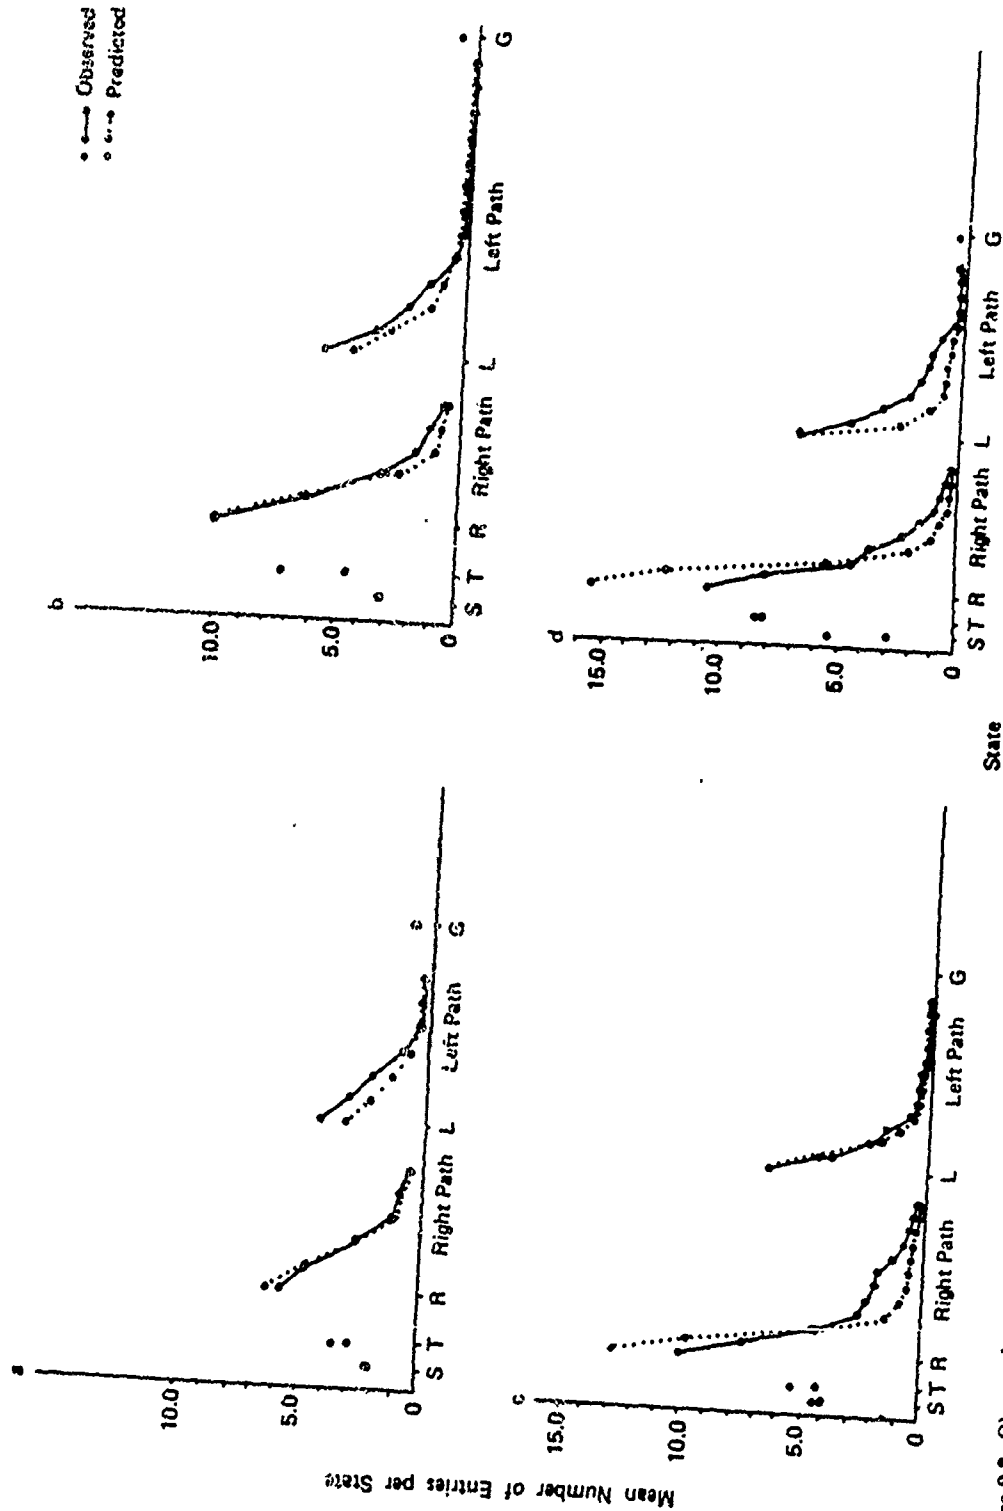
Figure 9.2. Observed and predicted values of mean visits per state for four water-jar problems. (a) problems 8, 5, 3; (b) 12, 7, 4; (c) 14, 9, 5; (d) 16, 10, 3. From "A Process Model for Water Jug Problems" by M.E. Atwood and P.G. Polson, 1976, *Cognitive Psychology, 8*, pp. 206–207. Copyright 1976 by Academic Press. Adapted by permission.
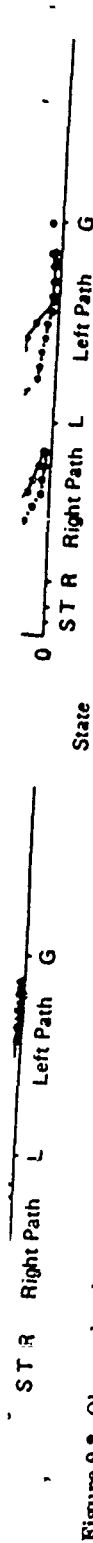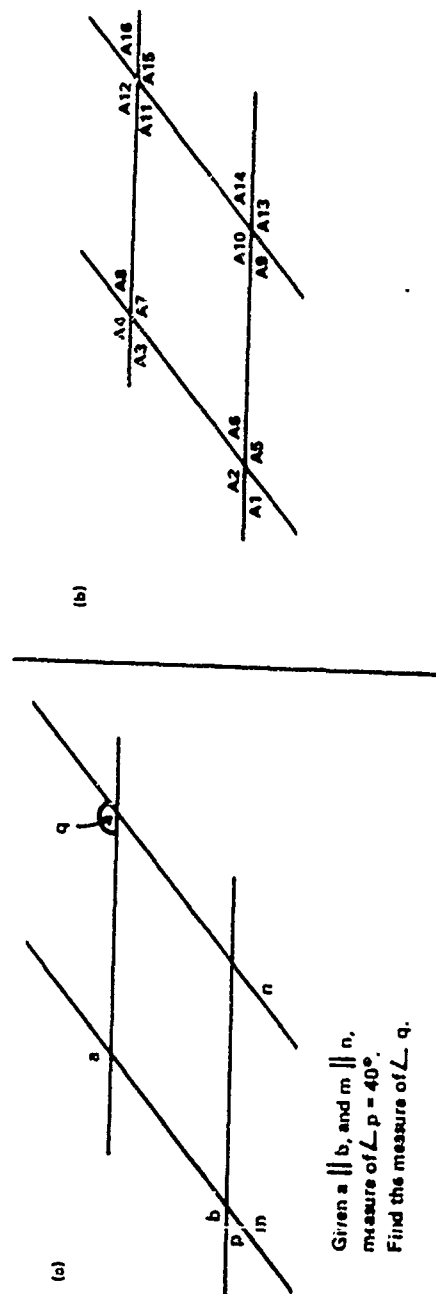
Figure 9.2. Observed and predicted values of mean visits per state for four water-jar problems. (a) problems 8, 5, 3; (b) 12, 7, 4; (c) 14, 9, 5, (d) 16, 10, 3. From "A Process Model for Water Jug Problems" by M.E. Atwood and P.G. Polson, 1976, *Cognitive Psychology, 8*, pp. 206–207. Copyright 1976 by Academic Press. Adapted by permission.

(a)

Given a ∥ b, and m ∥ n,
measure of ∠ p = 40°.
Find the measure of ∠ q.

Statement	Reason
1. Measure A1 (P) = 40°	1. Given
2. A1 (P) ≅ A6	2. Vertical angles
3. A6 ≅ A8	3. Corresponding angles
4. A8 supplementary A12 (Q)	4. Interior angles on same side
5. A6 supplementary A12 (Q)	5. Substitution
6. A1 (P) supplementary A12 (Q)	6. Substitution
7. Measure A12 (Q) = 140°	7. Definition of supplement

(b)

Figure 9.3. A solved problem in geometry. A1, A6, etc., in the solution refer to the positions of angles in the upper right diagram

below. They are generally similar to the solution shown in Figure 9.3.

The solution shown in Figure 9.3 was given by a computational model called Perdix that was formulated to simulate the students' performance. The structures and processes represented in Perdix are hypotheses about the knowledge that students acquire in a geometry course.

PROBLEM-SOLVING KNOWLEDGE

Perdix contains three kinds of knowledge, all represented as production rules: (1) problem-solving operators that make inferences, (2) perceptual concepts that recognize patterns in diagrams, and (3) strategic processes that set goals and select plans for problem-solving activities.

Problem-solving operators in geometry correspond to the theorems, postulates, and definitions that are used as reasons to justify steps in a problem solution. Examples include "Vertical angles are congruent" (a theorem), "Corresponding angles are congruent" (a postulate), and "If two angles are supplementary, the sum of their measures is 180°" (a definition). When the antecedent of one of these propositions is satisfied in a problem, then the consequent can be inferred. For example, because $A1$ and $A6$ are vertical angles in Figure 9.3, the inference that $A1$ and $A6$ are congruent is permitted. The propositions that correspond to the problem-solving operators are prominent in geometry instruction. They are represented in Perdix as production rules, with the antecedents as conditions and the relations that can be inferred as actions.

Patterns of information in the problem have to be recognized to determine that a problem-solving operator can be applied. For example to apply the inference rule, "Vertical angles are congruent," in Figure 9.3 and thus infer that $A1$ and $A6$ are congruent, the problem solver must first recognize that $A1$ and $A6$ are vertical angles In the geometry course, perceptual concepts are taught with examples using diagrams. In Perdix, knowledge for recognizing patterns is represented by discrimination networks, similar to the structures in the *Elementary Perceiver and Memorizer*, EPAM (Feigenbaum, 1963) and the *Concept Learning System*, CLS (Hunt et al., 1966). Perdix's recognition system is based on features of a diagram, such as sides of two angles that are collinear, along with other

information that may be given or inferred, such as statements that lines are parallel or perpendicular. An example is shown in Figure 9.4 which represents the process that can recognize a pair of vertical angles, a pair of angles formed by bisecting an angle, and other patterns that involve pairs of angles that have a single vertex.

Strategic knowledge is needed for setting goals that organize problem-solving activity. In the example problem of Figure 9.3, the main goal is to find the measure of angle Q. This cannot be achieved directly, and the problem solver must know that a way of finding the measure of an angle is to find a quantitative relationship (e.g., congruent or supplementary) of the unknown angle with one that has a known measure. This can be represented as a production: when the current goal is to find the measure of an angle, and the measure of another angle is known, set a subgoal of finding a quantitative relation between the unknown angle and the known angle.

The importance of strategic knowledge is illustrated in the protocol in Table 9.5. The student was working on the problem shown in Figure 9.3. The student marked several angles in a copy of the diagram; these are indicated in parentheses in the protocol of Table 9.5 in relation to the diagram in Figure 9.3(b). For example, "P would equal one ($\rightarrow A1$)" indicates that a label '1' was written on the angle in the student's diagram at position $A1$.

The student seems to have known the problem-solving operators and the geometric patterns needed to apply them (this was confirmed in another part of the interview) but was unable to solve the problem. A likely hypothesis is that the student lacked knowledge of the problem-solving strategy needed in this problem. The strategy involves forming a chain of angles that are related by congruence. Knowledge of this strategy involves setting a series of goals, when the problem requires a relation between two angles and none can be recognized, one must first find an angle related to one of them by congruence and then try to relate that angle to the other angle. This strategic procedure can be applied recursively until an angle is found that is related to the goal angle by one of the geometric relations from which a quantitative relation can be inferred.

Four of the six students who were interviewed in Greeno's study solved the problem in

Start: Bind V1 and V11 to angles; V2 to their shared vertex.

T1: Are both sides of V1 the same as sides of V11?
- yes → Add relation (CIRCPR V1, V11)
- no →

T2: Is one side of V1 the same as one side of V11?
- yes → Bind V3 to shared side
- no →

T3: Is one side of V1 collinear with one side of V11?
- yes → Bind V3 and V13 to the collinear sides
- no → FAIL

T4: Are the remaining sides of V1 and V11 collinear?
- yes → Add relation (VERT V1, V11)
- no → FAIL

T5: Are the remaining sides of V1 and V11 collinear?
- yes → Bind V6 to line of sides
- no →

T6: Is V6 perpendicular to the line of V3?
- yes → Add relation (BOTHRIGHT V1,V11)
- no → Add relation (LINPR V1,V11)

T7: Do the remaining sides of V1 and V11 form an angle that V3 bisects?
- yes → Add relation (BISECTPR V1, V11)
- no →

T8: Are the remaining sides of V1 and V11 perpendicular with V3 in the smaller angle?
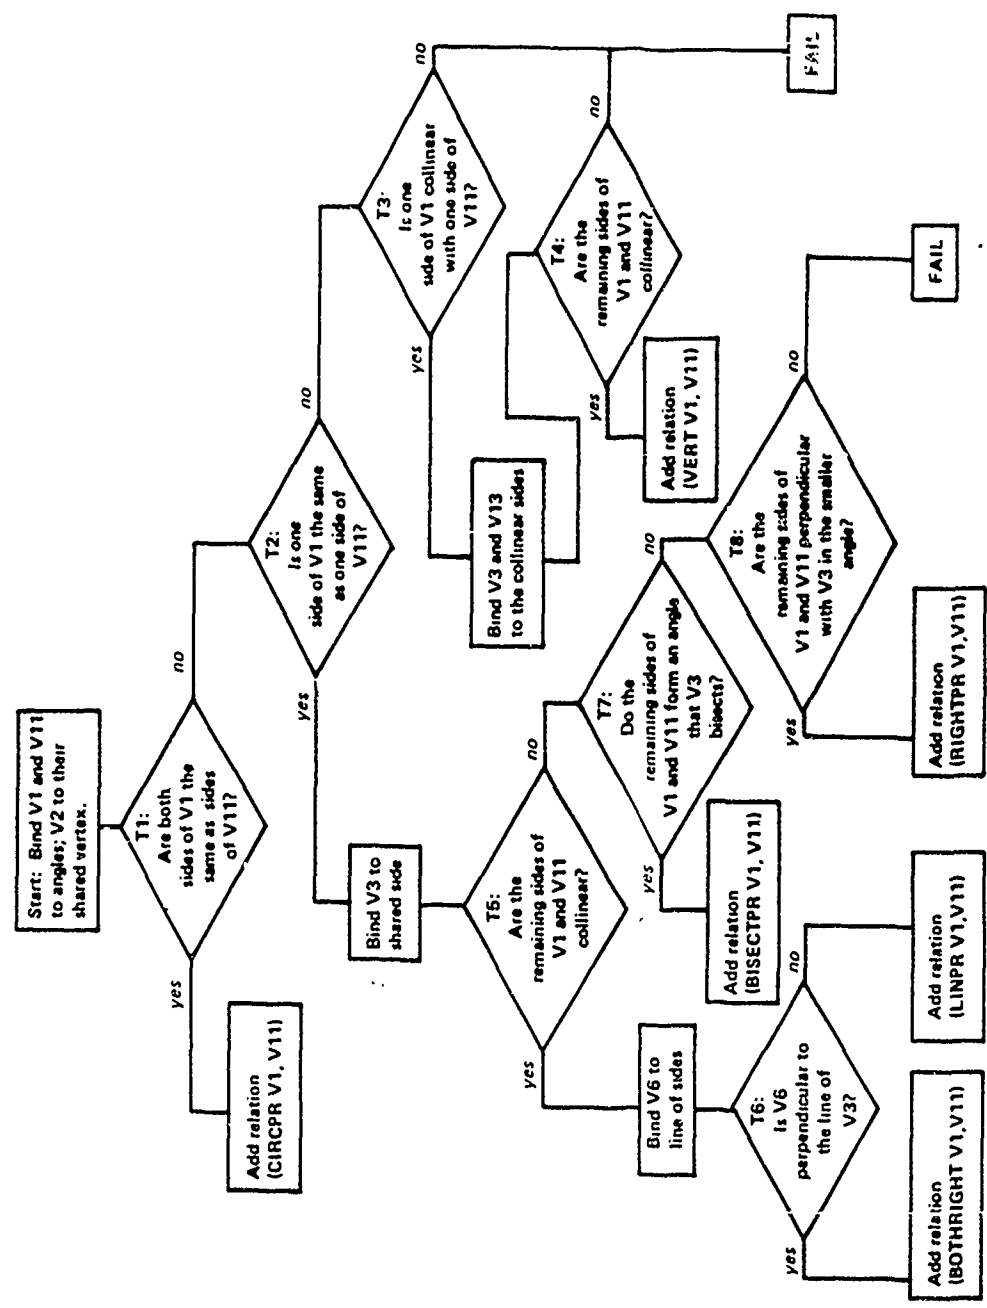- yes → Add relation (RIGHTPR V1,V11)

Figure 9.4. Pattern recognizing system in Perdix for analyzing angles with a shared vertex. From "A Study of Problem Solving" by J. G. Greeno in *Advances in Instructional Psychology: Vol. 1* (p. 21), R. Glaser (Ed.), 1978. Hillsdale, NJ: Erlbaum. Copyright 1978 by Lawrence Erlbaum Associates. Reprinted by permission.

Table 9.5.   Protocol of an attempt to solve Figure 9.3

S:	All right, I would put, like, P would equal one ($\rightarrow$A1).
E:	Okay.
S:	And then, two ($\rightarrow$A6).
E:	Put in two there, right.
S:	And then three ($\rightarrow$A15); no, wait—three ($\rightarrow$A15) and four ($\rightarrow$A12), I guess.
E:	Okay. Now, why did you put two there?
S:	Well, I don't know. It could have something to do with vertical angles.
E:	Okay.

...

S:	All right, the first thing I guess I should try to do, I would try to find if there were any alternate interior or corresponding angles?
E:	Okay.
S:	Or any of those.
E:	Mm-hm.
S:	I guess I would say that . . . well, wait a minute. I guess maybe I would put five there ($\rightarrow$A16).
E:	Okay.
S:	I don't know if I would need this.
E:	Okay.
S:	These two are supplementary.
E:	Right.
S:	That doesn't help much. And then, the measure of angle five . . . would it equal the measure of angle one?
E:	Well, you might have to work that out.
S:	How . . . if this equals . . . this equals forty.
E:	That's right.
S:	Oh, all right. Wait, the measure . . . I can't, I don't know. I don't know how to do these.
E:	Okay.

Figure 9.3 successfully, apparently applying the strategy of forming a chain of congruent angles. The students used different specific sequences of angles, which could result from differences in the way they scanned the diagram looking for angles to add to the chain, or differences in the ease with which they recognized various geometric patterns. About a week after one unsuccessful student gave the protocol in Table 9.5, that student successfully solved a different problem that also required the chaining strategy.

In geometry instruction, very little strategic knowledge is taught explicitly; it has to be inferred by the students from example problems. Inference appears to be a common feature of instruction in domains requiring acquisition of knowledge for problem solving, and, in the light of results of basic research on cognitive processes in problem solving, we consider the explicit teaching of problem-solving strategies to be a potentially productive development for instruction.

Strategic knowledge is represented in Perdix by productions that select plans for work on problems. A *plan* is a general approach to the problem, based on information in the problem situation. GPS forms such plans using its general planning strategy, described on page 601. Perdix has specific cognitive structures for plans that are used frequently for geometry problems. Forming a chain of congruent angles is one such plan. Another is using congruent triangles to prove that two angles or two line segments are congruent.

The organization of planning knowledge in Perdix is similar to that developed by Sacerdoti (1977), called a *procedural network*. In a procedural network, there are units of knowledge corresponding to actions at different levels. Each of these knowledge units includes information about the prerequisites and consequences of an action that can be performed. In Perdix, knowledge of each plan includes information about goals that can be achieved using the plan (its consequences), conditions in problems that make the plan promising (its prerequisites), and subgoals that should be set if the plan is adopted.

Perdix's strategic knowledge constitutes the main way in which it differs from GPS. Strategic knowledge in GPS is the general means–ends strategy that can be used in any domain for which the problem solver is taught the operators together with the productions that connect operators with differences, and is given the goal of a problem. The hypothesis represented in Perdix is that instruction in a domain such as geometry leads to acquisition of strategic knowledge specific to that domain, such as the schematic knowledge that represents plans to use chains of congruent angles or congruent triangles. Both GPS and Perdix construct plans that are more general than the actions that must be performed in solving the problem. The difference is that GPS forms plans using its general means–ends strategy, whereas Perdix's plans are based on knowledge of specific geometry strategies.

When GPS plans, it uses the strategic process

If 2 SIDES OF $\triangle \cong$, then the angles opposite those SIDES ARE $\cong$

Given $\triangle ABC$; $AC \cong BC$

Prove : $\angle A \cong \angle B$

1. If $\triangle ABC$; $AC \cong BC$
2.
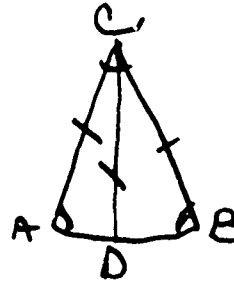
1. Given

Figure 9.5. Written work and drawing by a student on the problem, "Prove that if two sides of a triangle are congruent then the angles opposite those sides are congruent. From "Theory of Constructions and Set in Problem Solving" by J.G. Greeno, M.E. Magone, and S. Chaiklin, 1979, *Memory and Cognition, 7*, p. 447. Copyright 1979 by the Psychonomic Society. Reprinted by permission.

of means–ends analysis in a problem space that contains features taken directly from the basic representation of the problem. The planning space of GPS can be acquired by learning which features of objects should be given first priority. In Perdix, planning uses schematic knowledge of specific methods applicable to problems in the domain of geometry. These schemata include general subgoals, such as proving that triangles are congruent or finding an angle with a relation based on parallel sides, that can be used as intermediate steps. The associations of these subgoals with the goals they help to achieve have to be acquired by students; they are not explicitly given as goals of problems in which they are used.

SOLUTION OF ILL-STRUCTURED PROBLEMS

A hypothesis that is consistent with the analysis of geometry problem solving is that domain-specific strategic knowledge may provide the main basis for solving ill-structured problems. Problems may lack definite structure for many reasons. One important source of indefinite structure is that a problem may require knowledge from several different sources, with the result that its solution requires coordinated work in several disparate problem spaces (Simon, 1973).

A modest form of this kind of problem arises in geometry, involving problems that require

construction of auxiliary lines. The problem space that is presented, including a diagram, given information, and a goal to be proved, must be augmented in order for the problem to be solved. Greeno, Magone, and Chaiklin (1979) proposed that the solutions of such problems can be based on an individual's knowledge of plan schemata. In the Perdix model the need for an auxiliary line is recognized when a plan's prerequisites are partly satisfied in the problem situation. This leads to the definition of a subproblem; the goal is to complete the pattern of features that constitute the prerequisites, which is achieved in a problem space with operators appropriate to that goal.

An example is shown in Figure 9.5, the drawing and written work of a student on the following problem: Prove that if two sides of a triangle are congruent, then the angles opposite those sides are congruent. The protocol given by this student is in Table 9.6. After drawing the triangle ABC, the student added the line CD, which is not specified in the initial problem space. student's comments at *1 and *2, along the retrospective comment at *3, provi ence that construction of the auxiliary li. vas related to a plan of proof involving ngruent triangles, and the construction com- ed a pattern that is required for that plan to be applied—that is, the presence of two triangles in the diagram. Perdix simulates solutions like

Table 9.6.   Protocol for the problem of Figure 9.4

	S:	Okay, if two sides of a triangle are congruent, so . . . draw a triangle.
	E:	Okay.
	S:	Then the angles opposite those sides are congruent. Okay, so, like, if I have . . . given: triangle ABC—I'l' letter it ABC.
	E:	Right.
	S·	And then I have . . . prove: . . . do I already have these two sides given? Okay. Two sides of a triangle are given.
	E:	Mmm-hmm.
	S:	Let me go back to my given and say that segment AC is congruent to segment BC.
	E:	Okay.
	S:	And I want to prove that angle A is congruent to angle B.
	E:	Good.
	S:	All right. Let me write down my given. Okay. And mark my congruent sides. Okay, so I want to prove that angle A is congruent to angle B. Now, let's see. Do you want . . . ?
	E:	Yeah. Why are you drawing a line there?
*1	S:	I don't know yet.
	E:	Oh, that's okay. Don't erase it.
	S:	I'm going to do it, no, I just . . .
	E:	Oh, okay, fine.
	S:	Okay . . . okay, then I could . . . if I drew a line . . .
	E:	Mmm-hmm.
*2	S:	That would be the bisector of angle ACB, and that would give me . . . those congruent angles . . . no. (Pause.) Yeah, well, that would give me those congruent angles, but I could have the reflexive property, so this would be equal to that. Okay, I've got it.
	E:	Okay.
	S:	Okay.
	E:	Now, before you go ahead and write it all down, when you said you were going to draw the line . . .
	S:	Yeah.
	E:	And I said why are you doing that, and you said you didn't know yet, what do you think happened to give you the idea of making it the bisector?
*3	S:	Okay, well, I have to try to get this . . . I have to try to get triangle ACD congruent to BCD. Because, if I do that, then angle A is congruent to angle B because corresponding parts of congruent triangles are congruent.
	E:	So you were drawing the line to give yourself triangles, is that the idea?
*4	S:	No, to . . . to get a side that was in both triangles.
	E:	Okay.
	S:	And to get congruent angles.
	E:	So that's why you drew it as the bisector.
	S:	Yeah.

this with a process of pattern recognition that identifies partial patterns of two triangles missing a line, and uses special problem-solving operators to complete the patterns.

Another way in which problems can be ill structured involves the way in which goals are formulated. Goals in well-structured problems are presented as specific objects (e.g., a specific logic expression to be derived or a specific distribution of water among some jars). In ill-structured problems, goals are often underdetermined, with several alternative ways in which they might be satisfied. Examples are frequently cited from art or science, such as the goal of composing a fugue, or of designing an interesting experiment. In school geometry, the goals of problems are usually well specified, but a subgoal that arises in many problems functions as an indefinite goal for experienced problem solvers. This is the goal of proving that two triangles are congruent. There are several ways in which congruence of triangles can be proved, involving different patterns of congruent components such as side–side–side, side–angle–side, and so on. Beginning learners treat these as definite subgoals, trying one after another

until one works (Anderson, Greeno, Kline, & Neves, 1981). More experienced students do not mention specific patterns in their protocols, and appear to engage in a relatively diffuse search for congruent components of triangles with a kind of monitor that identifies whatever pattern of congruent components happens to emerge Greeno (1976) hypothesized that experienced students acquire an integrated structure of knowledge in the form of a pattern-recogn.zing system that represents the goal of proving that triangles are congruent. A version of this that was implemented in Perdix is shown in Figure 9.6.

ACQUISITION OF PROBLEM-SOLVING SKILL

An important question is how the knowledge that is needed for solving problems in a domain such as geometry is acquired. Studies of learning involving the three kinds of knowledge needed for problem solving have been undertaken: these are problem-solving operators, perceptual concepts for pattern recognition, and strategic knowledge.

Anderson (1982) based an analysis of problem-solving operators on observations of three students as they studied and worked problems in the early sections of a geometry text. He simulated processes of acquiring problem-solving skill in a version of his ACT model (cf. Anderson, 1983).

A major aspect of Anderson's model is a process that acquires cognitive procedures from declarative information. This model learns new procedures by working on problems. When ACT encounters a problem for which it has not learned a procedure, it uses general problem-solving methods along with information that is available. For example, a geometry problem may require finding a theorem that can justify a step in a proof. The ACT model has a general procedure for searching in a list of theorems and for matching features of theorems to the information in a problem. When an applicable theorem is found, ACT asserts that theorem to solve that part of the problem.

ACT has a learning process called *procedur alization*, which forms new production rules that are added to ACT's procedural knowledge. A new production can be formed when a theorem has been found and applied successfully in problem solving. The new production has conditions corresponding to selected features in the problem situation, and an action that asserts

the theorem. The production is a new problem-solving operator. ACT has acquired a new ability to assert a theorem in appropriate conditions without having to search through the list of theorems in the text. It has learned the theorem. not in the sense of having memorized it, but in the sense of being able to recognize when it is applicable, and to apply it.

Acquisition of perceptual concepts for pattern recognition in problem solving was studied by Simon and Gilmartin (1973) in the domain of chess. The learning mechanism used was adapted from the EPAM model (Feigenbaum, 1963), which simulates acquisition of discrimination networks like that in Figure 9.4. Simon and Gilmartin developed an EPAM-type model that acquired knowledge of patterns of chess pieces from presentations of board positions. This knowledge was used to simulate performance in a task of reconstructing positions after brief presentations, a task known to differentiate among players according to their level of skill (Chase & Simon, 1973; deGroot, 1965; also see "Chess and Go").

Acquisition of strategic knowledge for solving problems has been studied empirically by Schoenfeld (1979). Four students in upper-division college mathematics courses were given special instruction in the use of five heuristic strategies for working on problems: drawing a diagram, arguing by induction, arguing by contradiction or contraposition, considering a simpler problem with fewer variables, and establishing subgoals. Each strategy was presented in a training session, lasting about one hour, including an explanation of conditions in which the strategy is useful as well as practice in using the strategy. Students took a pretest and a posttest with problems not included in the training. These students had a list of the strategies available during the posttest and were reminded from time to time to try one of the strategies if they were not progressing well on a problem. Performance of these students was superior to that of another group of students who had worked on the same training problems as the instructed group, but without explanation of the strategies. Thinking-aloud protocols confirmed that students considered and used strategies that they had been trained to use. The training was especially effective with strategies that have clear cues for their application. the fewer-variables strategy, cued by the presence of many variables,
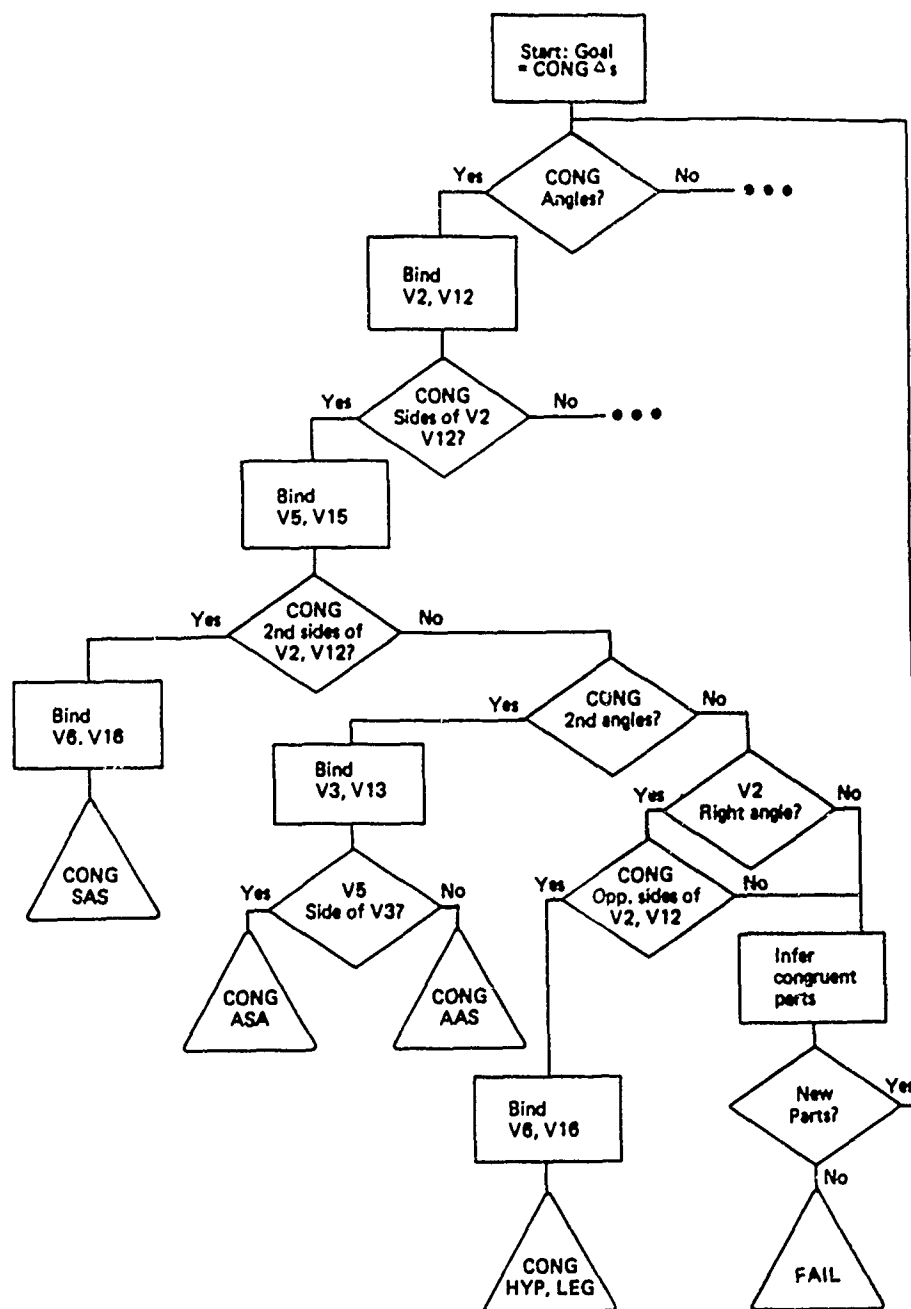
Figure 9.6. Part of Perdix's goal structure for proving congruence of triangles, represented as a pattern-recognizing system Cong = congruent, SAS = side angle side, ASA = angle-side-angle, AAS = angle-angle side; HYP, LEG = hypotenuse-leg. From "Indefinite Goals in Well-Structured Problems" by J.G. Greeno, 1976, *Psychological Review*, *83*, p 486 Copyright 1976 by the American Psychological Association. Adapted by permission of the author.

and arguing by induction, cued by an integer argument.

Processes of acquiring strategic knowledge have been addressed in theoretical analyses by Anzai and Simon (1979) and by Anderson, Farrell, and Sauers (1984). Anzai and Simon observed and simulated acquisition of a strategic concept in the Tower of Hanoi puzzle. The

concept involves movement of a set of disks requiring a sequence of individual moves, with the sequence considered as a global action. Anderson et al. simulated the acquisition of knowledge for applying techniques in learning the programming language LISP. In both theoretical analyses, important factors in acquiring strategic knowledge are the activation of a problem goal that can be achieved by a sequence of actions and the acquisitions of a production in which the action of setting the goal is associated with appropriate conditions in the problem situation.

Einstellung (Set)

The context in which problem solving occurs may have an important influence on the process. As a consequence of previous tasks that a subject has engaged in or previous stimuli that have been presented, certain responses may become more readily and speedily available and others less readily available. The subject has acquired a 'set' for the familiar stimuli and responses.

One experimental design that has often been used to demonstrate the effects of set is to present subjects with a sequence of tasks that induce set, then a new sequence of tasks in which this set either facilitates or impedes performance relative to that of control subjects who were not exposed to the first sequence. Luchins (1942) conducted a well-known set of experiments using this design, with water-jar tasks.

In Luchins's version of the water-jar task, subjects must measure a specified amount of water, using a given set of ungraduated measuring jars. A source of water is assumed to be available, so that any of the jars can be filled to its capacity if the subject so chooses. Water can be poured from one jar to another, until the target jar is filled or the source jar is empty. Also, the contents of the jar can be discarded.

The series of problems that Luchins used is shown in Table 9.7. All the problems except the first and the ninth can be solved by filling jar B, then pouring from it to fill A, and then filling C twice ($X = B - A - 2C$). But problems 5, and 7 through 11, can also be solved using only jars A and C—by either adding the contents of C to the contents of A, or subtracting the contents of C from A, and for problem 9, the $B - A - 2C$ procedure does not work.

Table 9.7. Problems used by Luchins (1942)

Problem Number	Measuring Jugs			Required Amount
	A	B	C	
1	29	3		20
2	21	127	3	100
3	14	163	25	99
4	18	43	10	5
5	9	42	6	21
6	20	59	4	31
7	23	49	3	20
8	15	39	3	18
9	28	76	3	25
10	18	48	4	22
11	14	36	8	6

Subjects given problems 7 through 11 immediately after solving problem 1 generally use the two-jar procedure just described. Subjects who are first given problems 1 through 6 generally use the $B - A - 2C$ procedure, which is more complex than necessary for problems 7 through 11, and they have considerable difficulty with problem 9.

Set effects can be the result of several cognitive processes of which three that have been put forward will be discussed.

First, set may be the result of a bias in retrieving knowledge structures from memory. A standard assumption is that the alternative concepts or cognitive procedures that might be retrieved have varying strengths or levels of activation which determine the probabilities of their retrieval. If a cognitive unit has been used successfully several times in the immediate past, a relatively high level of activation for that unit results.

Schemata used in planning provide one kind of structure that can account for set. An example is in the domain of geometry, where Greeno et al. (1979) developed a simulation model with planning schemata, described above in "Geometry Exercises." Luchins (1942) included a study of geometry problem solving in his investigations of Einstellung. Figure 9.7 shows the kind of problem used as a test. The proof can be obtained in one step; $\angle AMC$ and $\angle BMD$ are vertical angles. However, if subjects were first given a series of problems where they used congruent triangles in proofs, they were likely to construct the more complex proof for Figure 9.7 in which triangles AMC and BMD are proved congruent by side–side–side. An explanation is provided if
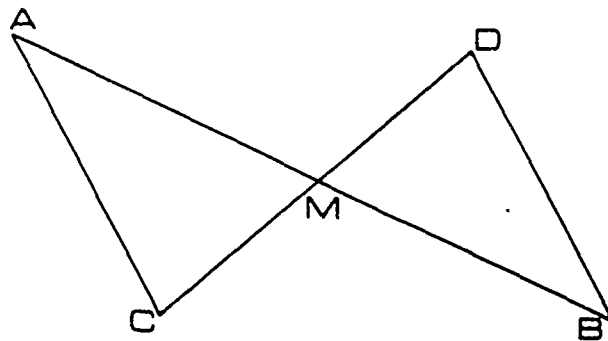
Given· M is the midpoint of $\overline{AB}$ and $\overline{CD}$;
       $\overline{AC} \cong \overline{BD}$.

Prove $\angle AMC \cong \angle BMD$

Figure 9.7. An einstellung problem in geometry.

we assume that students have a schema corresponding to the plan of using congruent triangles for a proof, and that this schema has a high level of activation because of its use in the initial series of problems. Greeno et al. (1979) reported an experiment with a test problem that could be solved by using either congruent triangles or angles formed by a transversal with parallel lines, but either method required construction of an auxiliary line. Subjects were given series of problems to solve before the test problem, involving either congruent triangles or parallel lines. They were strongly biased toward solving the test problem in the same way that they had solved the trial problems.

Set based on activation may either facilitate task performance or impede it, depending on whether the memory elements that are activated contain the information that is needed for performance. Sweller and Gee (1978) showed that the tendency to use a previously successful rule can greatly facilitate solution of a relatively complex problem, presumably by eliminating the need to search in a large space of possibilities, even when in the same situation it prevents subjects from noticing a simpler solution method. Such situations are common, since set is bound to arise wherever memory organization is not neutral with respect to the problem-solving process—that is, wherever there are alternative ways of storing information in memory, one of which may be more conducive

to retrieval in a given problem context than another.

A second possible explanation of Einstellung is provided by composition of productions, investigated first by Lewis (1978). *Composition* is a process in which a newly acquired production performs actions that required two or more productions in the previous knowledge structure. Composition generally makes performance more efficient by providing a way to act directly rather than requiring several steps to achieve a goal. The new productions created by composition usually have conditions that are relatively specific, and in some production systems (including ACT) this leads to their being preferred to productions with less specific conditions. Anderson (1982) noted that this would simulate the performance observed by Luchins (1942) on problems like Figure 9 7

Third, some setlike phenomena could also be produced by the basic problem-solving procedure that a subject uses. We have already noted that subjects frequently use the heuristic of means–ends analysis—that is, comparing situation with goal and taking an action that seems to reduce the difference between them. In their analysis of behavior of subjects solving water-jar problems, Atwood and Polson (1976) showed that where alternative actions could be taken, most subjects selected the one that led to a situation that most resembled the goal situation. As with the more specific sets induced by

Luchins's manipulation, this general set to pick paths that lead toward the desired goal can sometimes interfere with problem solution. Where memory limitations prevent subjects from looking far ahead, this goal-oriented strategy may sometimes produce a myopic preoccupation with immediate progress and the avoidance of paths that lead to the goal only indirectly Jeffries, Polson, Razran, and Atwood (1977) showed that, without looking ahead, subjects solving the Missionaries and Cannibals puzzle would have difficulty (as, in fact, they do) on the step where they were required to bring two persons back from the farther bank of the river to which they were trying ultimately to transport them all.

Problems with Specified Procedures

The present section examines tasks in which the problem presents material for a procedure, and the task is to apply the procedure to find the result. While the tasks discussed in "General Knowledge for Novel Problems with Specific Goals" and "Domain-Specific Knowledge for Familiar Problems with Specified Goals" specify a goal and require finding a method to get there, the tasks in this section specify a method and ask where the method leads.

The tasks chosen for discussion come from arithmetic. Many tasks in mathematics involve applying procedures, for example, finding a derivative in calculus or the product of two expressions in algebra. Such tasks may not be thought to involve problem solving, since they require knowledge of a procedure rather than search in a space of possible solutions. However, students who receive these tasks as homework assignments and presumably the teachers who assign them consider them to be problems.

More significantly, the knowledge required for these procedure-based tasks is similar to the knowledge that students acquire when they learn to solve problems that do not specify solution methods, such as geometry proof exercises or water-jar problems. Knowledge for planning in geometry consists of a set of procedures that the student has acquired for solving various kinds of problems. In geometry use of these procedures requires recognition of their applicability, which is not required if the problem calls for the operators *subtract* or *differentiate*. Nevertheless, characteristics of

the procedural knowledge that have been identified by theoretical analyses of the various tasks are more notable for their similarities than for their differences.

This section focuses on empirical methods that have been used to infer the nature of procedural knowledge, on inferences based on patterns of errors that occur in elementary arithmetic and on inferences from latency data.

Diagnosis of Cognitive Procedures from Patterns of Errors

Brown and Burton (1980) analyzed children's knowledge for solving subtraction problems with multidigit numbers. Their data were obtained in an arithmetic achievement test taken by 1325 school children. Although performance on tests is ordinarily used to assign a simple score for each student, thus allowing judgments of which students have learned a satisfactory amount, Brown and Burton's analysis showed that test data are potentially much richer and can be used to make stronger inferences about the nature of children's knowledge.

The more powerful theoretical use of test data depends on two conditions. First, performance on the test is not characterized simply by the number of problems correct, but by the specific answers given to all the problems, with particular attention to the incorrect answers. Second, the analysis of each student's test performance consists of a model of a procedure for solving the problems.

The idea of using patterns of errors to infer underlying psychological processes is not new, either in the psychological or the educational literature. Earlier psychological models were simpler, and the inferences about processes were correspondingly less powerful, an example is Polson, Restle, and Polson's (1965) use of errors to identify a stage of learning in which similar stimuli have not yet been discriminated. In the educational literature more complex psychological distinctions have been made, for example by Brownell in 1941. However, analyses of underlying psychological processes was informal in that work, consisting of verbal descriptions of procedures hypothesized to produce observed error patterns, and, as Brown and Burton documented, verbal descriptions of procedures turn out to be ambiguous in important ways.

An example of an individual student's

Table 9.8. One student's performance on subtraction problems
Source (Brown & Burton, 1978)

8	99	353	633	81
3	79	342	221	17
5	20	11	412	64
4769	257	6523	103	7315
0	161	1280	64	6536
4769	96	5243	139	779
1039	705	10038	10060	7001
44	9	4319	98	94
1995	76	15719	10962	7007

Note. From Diagnostic Models for Procedural Bugs in Basic Mathematical Skills by J.S Brown and R.R Burton, 1978, *Cognitive Science. 4.* p 178 Copyright 1978 by the ABLEX Publishing Co Reprinted by permission.

performance is shown in Table 9.8. This table contains six errors (the fourth problem in the second row, and all the problems in the third row), not a very good score. However, all but one of the errors were apparently caused by a single flaw in the student's procedure. When the subtraction required borrowing and the numeral to be decreased was zero, the student replaced the zero by a nine, but did not take the further step of subtracting *one* from the preceding digit.

Brown and Burton developed a general model of subtraction for which various flawed versions can be represented as variants. The desired outcome was that the performance of each individual child, like the one shown in Table 9.8, should correspond as closely as possible to one of the variants of the general model. The general model has the form of a procedural network, the formalism developed by Sacerdoti (1977) and used by Greeno et al. (1979) to explain constructions and set in geometry problem solving. The main features of a procedural network are that units of knowledge correspond to actions at differing levels of generality, and each action unit includes information about conditions for performing the action, and the action's consequences.

Figure 9.8 shows the action components in Brown and Burton's procedural network for subtraction. The diagram shows component procedures and their subprocedures, but does not show the control information that is also required. For example, the diagram includes a procedure *Subtract Column*, and three subprocedures, *Borrow Needed, Do-Borrow,* and *Complete-Column.* Control knowledge involving these subprocedures includes the information that *Borrow-Needed* is a test that determines whether it is necessary to borrow before finding the difference in the column, and the outcome of that test determines whether *Do-Borrow* will be called.

Brown and Burton formulated models of faulty performance by varying components of the procedural network for correct subtraction. For example, the flaw of borrowing from zero is modeled by removing some of the control processing from the procedure *Borrow-Ten* in the *Do-Borrow* subprocedure The change involves removing the decision *Find-Next-Column* if a zero is found, resulting in a procedure that just changes zero to nine and adds ten to the original column.

The family of models that Brown and Burton arrived at included 60 procedural flaws of the kind described above. They provide explanations for many of the patterns of performance found in the test data, and more students' performance is explained if combinations of elementary flaws are included in the analysis. About 40 percent of the students' error patterns were explained reasonably well by single flaws or combinations of two elementary flaws. In examining additional sets of data, more elementary flaws have been identified (115 were in the data base in 1982), and adequate explanations are typically provided for about 40 percent of students who make errors (VanLehn, 1982).

An alternative analysis of subtraction errors was provided by Young and O'Shea (1981), who developed a relatively simple production system that simulates correct subtraction performance and, by deleting individual productions, simulates faulty performance. Young and O'Shea's analysis provides explanations for about the same proportion of students as Brown and Burton's model. On the other hand, it provides explanations for only a small proportion of the patterns of performance that have been observed. While many patterns occur rarely, their existence provides evidence for a relatively complex generative system.

Another significant development was an effort by Brown and VanLehn (1980) and VanLehn (1983) to formulate a system that explains the production of flawed procedures. These formulations distinguish between a cognitive structure of partial knowledge of subtraction, and a fall-
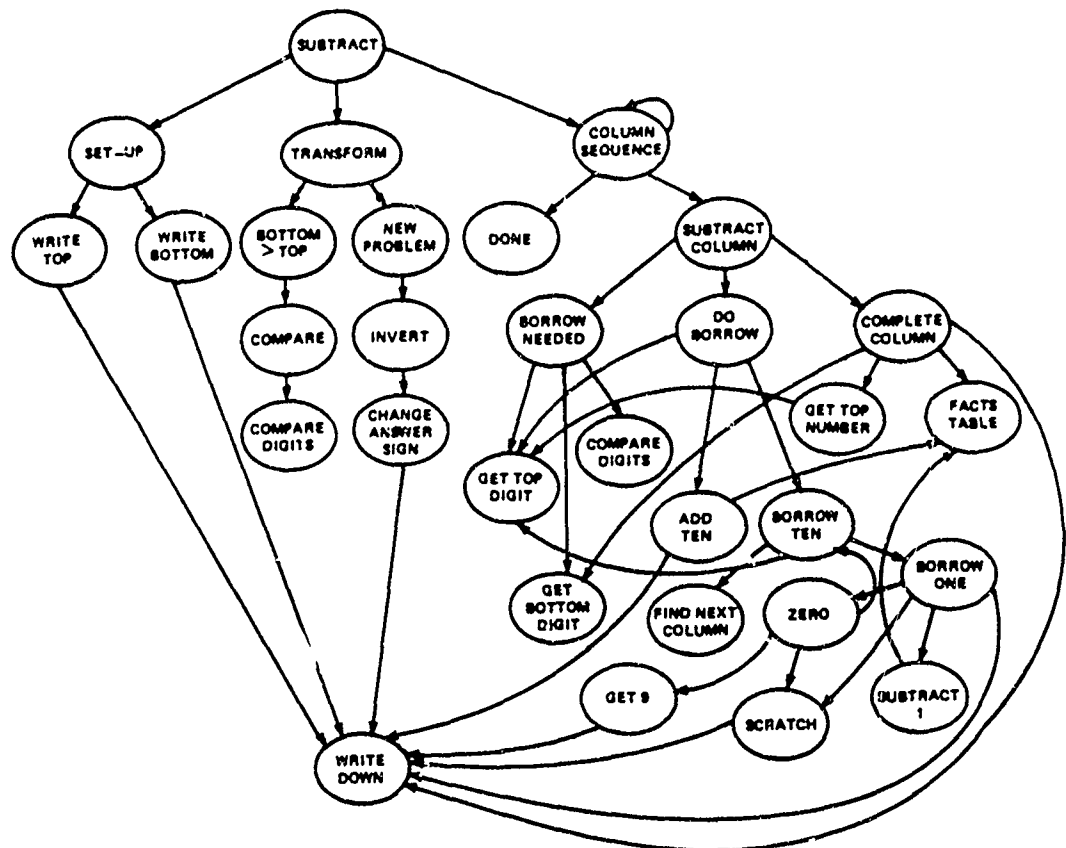
Figure 9.8. A procedural network for subtraction. From "Diagnostic Models for Procedural Bugs in Basic Mathematical Skills" by J.S. Brown and R.R. Burton,, 1978, *Cognitive Science*, *4*, p. 162. Copyright 1978 by the ABLEX Publishing Co. Reprinted by permission.

back process of problem solving that is used when a situation is encountered for which the partial knowledge is not adequate. In Van-Lehn's (1983) version, the underlying cognitive structures (core procedures) result from a combination of partial learning and deletion of components of procedural knowledge. A core procedure might, for example, lack a component for dealing with a zero during borrowing. When such as impasse occurs it is assumed that the problem solver applies a general problem-solving method in order to continue. Methods available include skipping an operation, applying the operation to a different problem element, and using an alternative operation that is applicable in a similar problem situation. One form of evidence that supports the theory comes from data obtained by giving students repeated tests Many students perform differently on two tests separated by two or three days, but the performance can be explained by assuming

a single core procedure for which different problem solving methods have been used.

VanLehn (1983) conducted theoretical investigations in which a small set of problem-solving methods is combined with a plausible set of core procedures to generate flawed subtraction procedures. The generative system that has been developed can account for about half of the flawed procedures that have been observed, amendments that would increase the theory's empirical adequacy could easily be devised but would not have strong theoretical motivation. Part of the progress that has been made involves identifying some general features of the system. It can be argued, on the basis of general properties of flaws, that the system has a push-down memory for recalling past goals, that goals are organized hierarchically, and that the representation of a goal includes the problem components to which the goal applies.

Another line of analysis that has developed

from the study of subtraction flaws involves analysis of cognitive structures for understanding general arithmetic principles that underlie correct subtraction procedures. See "Problem Representation in Mathematics and Physics."

Inferences Based on Latencies

An arithmetic task that is even simpler than multidigit calculation is the solution of basic addition problems such as 3 + 5. The main data used in the analyses of this task are latencies. Patterns of latencies of individual subjects are used to diagnose their solution processes.

In an empirical study by Groen and Resnick (1977), five preschool children who knew how to count and could recognize the numerals 1 to 9, but who did not know about addition were used as subjects. These children were taught a method for addition using blocks. The procedure was to count out two piles, each having one of the numbers in it, and then count how many were in the two piles together For example, for 3 + 5, the child could count out a pile of three, then a pile of five, and then count the complete set to find eight as the answer In showing the child the method, the experimenter sometimes started with the number on the left of the problem, and sometimes with the number on the right.

The problems used were basic addition facts involving the digits 1 to 5, omitting 5 + 5. After a child could solve all 24 of the problems correctly using blocks, a new apparatus was introduced The blocks were no longer provided, and the child answered problems by pressing buttons labeled 1 to 9. Children were shown how to count out answers on their fingers if it was necessary Children received from four to seven sets of problems with this apparatus, with about 25 problems per set.

The latency data were analyzed with regression techniques; models of cognitive processes were employed to determine the values of independent variables Two models were used. According to one, the process of finding the answer to each problem was much like the procedure that the children were taught. In that procedure, a number of sets must be counted; in fact, the total number of counts equals double the number of the answer. If we assume that a fairly uniform amount of time is used each time something is counted, the total amount of time needed is

$$T = A + B(2S),$$

where S is the sum of the two numbers (i.e. the answer), and A and B are constants. In the second model, the process is considerably simpler. The sum can also be found by starting with the larger of the two addends and counting up the number of the smaller addend. According to this model, the time it takes to find the answer is

$$T = A - B(M),$$

where M is the minimum addend, and A and B are constants. These two models are called the *sum model* and the *min model*, respectively.

Comparison of these two models with the data of children's performance is interesting primarily because of the possibility that children spontaneously change their procedure for solving addition problems. If they use the procedure they were taught, their performance should agree with the sum model. However, performance consistent with the min model would reflect a more efficient procedure, and would indicate that children had spontaneously modified their problem-solving procedures. It would thus indicate a significant capability for discovery or invention.

To apply either the sum or the min model to the data, problems are grouped according to the number of counting operations they require. Because the models specify different counting operations, they imply different groupings of items. For example, according to the sum model, the problems 6 + 1, 5 + 2, and 4 + 3 all require the same number of operations, but these problems require different numbers of counts according to the min model. On the other hand, the problems 4 + 3 and 3 + 5 require the same number of counts by the min model, but are different according to the sum model.

If a model is approximately correct, the regression based on it should give accurate predictions of problem latency. The criterion of fit used by Groen and Resnick was the proportion of variance R^2 accounted for by the regression. Higher values of R^2 indicate better agreement between the latency data and the theoretical function.

Table 9.9 shows that about half the subjects were fitted better by the min model than by the sum model. Values of R^2 are shown for latency data from each block of problems except the first, in which the children were getting used to the new apparatus. Subjects 2 and 4 were fitted

Table 9.9. Results of applying regression models to latency data
Source (Groen & Resnick 1977)

Subject	Block	Proportion errors	Proportion covert	R^2 Sum	R^2 Min	Slope of best fitting line (seconds)
1	2	15	.02	78*	65*	92
	3	09	00	45*	16	60
	4	.03	.04	79*	38	67
	5	03	08	.69*	57*	91
	6	.06	33	50*	59*	1 66
	7	.05	34	40*	63*	1 90
2	2	18	40	44*	65*	2.82
	3	14	57	51*	88*	2.30
	4	.11	57	.51*	.69*	2.06
	5	06	76	22	38*	1 40
	6	06	99	23	.54*	40
	7	.11	1 00	.17	43*	26
3	2	.04	.00	.14	.00	.
	3	.03	.09	.71*	.57*	.99
	4	.05	.05	.50*	.27	.73
	5	.11	.30	.06	.13	-
	6	.12	.92	.05	.30	-
	7	.07	.83	.03	.10	-
4	2	.25	.73	.23	.54*	1.77
	3	.12	.61	.38*	.41*	1.60
	4	.06	.94	.32	.65*	1.38
5	2	.04	.94	.47*	.43*	1.30
	3	.09	1.00	.55*	.49*	1.66
	4	.02	1.00	.25	.12	-
	5	.01	.99	.21	.17	-
	6	.06	1.00	.52*	.20	.64

Note Asterisks denote slope significantly different from zero at .01 level Italics denote maximum R^2.
Note From "Can Preschool Children Invent Algorithms?" by G. Groen and L.B. Resnick, 1977, *Journal of Educational Psychology, 69,* p 648 Copyright 1977 by the American Psychological Association. Reprinted by permission of the author.

better by the min model, subject 5 by the sum model, and subject 1 underwent a transition, being fitted better by the sum model in blocks 2 through 5, but by the min model in blocks 6 and 7 Another experiment, in which practice problems were presented in a systematic order, had similar results.

The important conclusion from these data is that the children must have discovered the procedure represented by the min model, since they were not taught how to add in that way. Neches (1981) developed an analysis of learning mechanisms that can produce modified procedures, and he used that system to simulate changes in counting procedures for addition problems The main ideas in the Neches model are that redundant components of the procedure can be removed, and when there are alternative ways of reaching the same result, the easier

method can be chosen. For example, in the sum procedure, the first addend is counted, and then later the process of counting the combined set includes counting the first addend as a part. Noticing this redundancy leads to removal of the initial count of the first addend from the procedure. Choice of the larger addend to initialize the procedure can be made if the subject notices that the same result is obtained with either addend, but that less effort is required when the larger addend is chosen. To produce modifications in its procedures, the Neches system requires a trace of its activity, including the goals that are active during the various stages of its performance.

The regression method has also been used in analyzing performance of adults in simple arithmetic tasks. Groen and Parkman (1972) found that college students' performance is

quite consistent with the min model. The slope of the best-fitting regression equation is far too small to correspond to verbal counting, but an analogue of a counting procedure might account for the result.

Performance in mental arithmetic has been studied recently by Ashcraft and his associates. Using a task in which subjects are shown a problem with a possible answer and are asked whether it is correct, Ashcraft and Battaglia (1978) found longer latencies for problems involving larger numbers, but this effect was not linear in the smaller addend, as required by the min model. A better predictor of latency was the square of the problem sum, an effect that seems inconsistent with a simple process of counting. Ashcraft and Battaglia also found shorter latencies for the rejection of wrong answers that differed more from the correct answer, than for wrong answers close to the correct answer. Another relevant finding by Winkelman and Schmidt (1974) was that latency increased for a false answer that could be correct for a different operation; for example 3 × 4 = 7. As Ashcraft and Stazyk (1981) have argued, these findings suggest a process of retrieval from memory, rather than a counting procedure, with effects on latency that result from the way in which information is stored and from processes of activation and search.

Problem Understanding; Representation

Before a problem can be solved, it must be understood. Many problems used in education are presented as natural-language texts that describe situations and ask questions, usually the values of some quantities. In laboratory studies, problems are often presented in the form of instructions that specify the goals and problem-solving operators that can be used in working on the problems. These texts or instructions must be interpreted, and some kind of representation of the problem must be generated before problem-solving processes can be put to work in seeking a solution.

The same problem may be represented in radically different ways, as is illustrated by the 'mutilated checkerboard' problem. The subject is given an ordinary 8 × 8 checkerboard, with alternating black and red squares, and a set of dominoes, each of which covers two squares. The

entire board can be covered by 32 dominoes, with no square left uncovered, and no domino hanging over the edge of the board Suppose now that the northeast square and the southwest square of the checkerboard are cut off, leaving 62 squares. Can the mutilated board now be covered neatly by 31 dominoes?

It is impossible for a human being or a computer to answer this question by exhaustive search in the obvious but enormous problem space in which the squares and dominoes are represented directly. Consider, however, an abstract problem space in which we represent only the number of dominoes that have been laid down, and the numbers of both black and red squares that remain uncovered. At the outset, because of the mutilation, there are 32 red squares, but only 30 black squares (or vice versa). Each domino covers exactly one red and one black square. Hence, no matter how the dominoes are placed on the board, after 30 have been placed, if that is possible, two red squares and no black squares will remain uncovered. But the final domino cannot cover two red squares, hence there is no way to complete the covering. Here, a change in problem representation changes the problem from one that is practically unsolvable to one that is quite easily solvable.

Another famous example of problem understanding, discussed by Wertheimer (1959), arises in finding the area of a parallelogram. Students are taught that the area of a parallelogram can be calculated with a formula $A = b \times h$, where b and h are the base and height, respectively. Wertheimer described two ways in which the formula may be understood. In one representation, b is the length of a horizontal side of the parallelogram, and h is the length of a vertical line drawn from a corner at the top of the figure to its base, as shown on the upper part of Figure 9.9. Many students, apparently using that representation, become confused if they are then asked to find the area of a parallelogram oriented differently, as in the lower part of Figure 9.9. Another way to understand the formula includes a relation between parallelograms and rectangles. A parallelogram can be transformed into a rectangle by removing a triangular piece from one end and attaching it to the other end. Then b and h are equal to the length and width, respectively, of the rectangle that the parallelogram can be transformed into. Children who
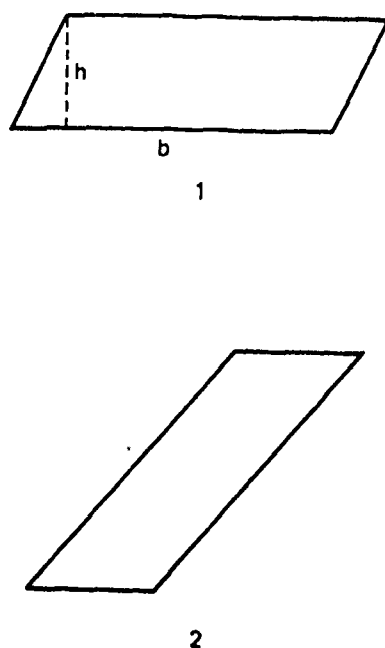
32 ~~dominoes.~~
~~and no domino~~
~~board. Suppose~~
~~and the squar~~
~~are cut off,~~
~~tilated board~~
~~in es?~~
~~being or a~~
~~by exhaustive~~
~~nous problem~~
~~dominoes are~~
~~however, an~~
~~we represent~~
~~have been laid~~
~~black and red~~
~~At the outset,~~
~~are 32 red~~
~~ares (or ice~~
~~tly one red and~~
~~latter how the~~
~~1, after 30 have~~
~~wo red squares~~
~~all uncovered.~~
~~cover two red~~
~~to complete the~~
~~oblem represen-~~
~~om one that is~~
~~it is quite easily~~

~~problem under-~~
~~(1959) arises~~
~~gram Students~~
~~rallelogram can~~
~~= b × h, where~~
~~it, respectively.~~
~~s in which the~~
~~n one represen-~~
~~ontal side of the~~
~~gth of a vertical~~
~~top of the figure~~
~~r part of Figure~~
~~using that repre-~~
~~they are then~~
~~elogram oriented~~
~~rt of Figure 9.9.~~
~~formula includes~~
~~as and rectangles.~~
~~isformed into a~~
~~gular piece from~~
~~other end. Then~~
~~ngth and width,~~
~~had the parallelo-~~
~~o. Children who~~

Figure 9.9. Parallelograms in two orientations. Some students who learn the formula $A = b \times h$ have difficulty applying it to a figure like the lower one (Wertheimer, 1945–1959).

understand the parallelogram problem in this way have no difficulty in solving problems where the figure is oriented differently and can frequently transfer their knowledge to solve more complex problems, such as finding the area of a trapezoid. The two representations involve different features of specific problems, one with b and h identified with specific locations in the figure, and the other with b and h defined in more general terms.

Understanding Problem Instructions

In most studies, consideration of subjects' behaviors in problem-solving tasks is begun after the subjects have received the definition of the problem with appropriate instructions, and have been tested by the experimenter for their understanding of the problem. A few studies investigate the processes required for assimilating the problem before attempting to solve it.

In the situations already studied, solution of the problem is likely to proceed by a form of means–ends analysis. Therefore, the information that subjects extract from instructions is probably similar to the information needed by

the General Problem Solver. When GPS is given a problem, it is provided with a list of the objects involved in the problem, the relevant properties of these objects, operators for legal moves, a description of the starting situation, and a set of tests to determine when the final goal has been reached. GPS may be provided with, or otherwise must acquire by learning, a set of tests for differences between situations and a set of productions that evoke, with certain differences, operators that are relevant to reducing these differences.

For example, in the Tower of Hanoi problem, the objects consist of N disks (where N = number) and three pegs. A legal move consists of transferring the smallest disk on one peg to another peg that holds no smaller disk. Hence, the size of a disk is its relevant property. Situations differ as to which disks are on a particular peg, or on which peg a particular disk is located. In one starting situation, all the disks are held on a single peg; the goal is to move the entire set of disks to another particular peg. The problem description must provide this information in English, and the subject (or computer program) must convert this English prose into an internal representation that permits situations, moves, and their consequences to be modeled. A disk, for instance, may be represented as a schema, one of whose attribut is its size, and a peg by a schema, one of whose attributes is the list of disks currently on that peg. A move operator is a process that changes a pair of the latter lists by moving the name of a particular disk from one list to the other.

Two central problems for psychological research on the understanding of problem instructions are: (1) how the verbal instructions are converted to an internal representation, and (2) what characteristics of the instructions cause the problem to be represented in one way, rather than other possible ways. The second question is especially important when alternative representations result in problem difficulty differences (as with the mutilated checkerboard example), or provide differing degrees of generality (as with the parallelogram problem). These questions have been addressed by Hayes and Simon (1974), who obtained information about internal representations by collecting extensive verbal protocols of problem-understanding processes. By using problems in which alternative representations were available, Hayes and

Simon also cast light on the question of which representations are formed.

The Understand program (Hayes & Simon, 1974) is a computer simulation of the problem-understanding process for puzzlelike problems like the Tower of Hanoi or Missionaries and Cannibals—that is, for problems that do not assume that the subject has any prior knowledge of the problem domain. The program matches human thinking-aloud protocols sufficiently well to lay claim to being a good first-approximation model of the process.

Understand operates in two principal phases. In the first, a language-parsing program extracts the deep structure from the language of the instructions. In the second phase, another set of processes constructs from this information a problem representation that is suitable as input to a GPS-like problem-solving program. This is accomplished by (1) identifying the objects and sets of objects that are mentioned in the parsed text, (2) identifying the descriptors of those objects and the relations among them, (3) identifying the descriptions of legal moves and constructing move operators that fit those descriptions, (4) identifying the description of the solution and constructing a test for attainment of the solution, and (5) constructing an organization of schemata that describes the initial problem situation.

For example, after parsing the written description of the Tower of Hanoi problem, Understand would identify pegs and disks as the relevant sets of objects, and would notice that disks are on pegs and that they move from one peg to another. It would extract the information that only the smallest disk on a peg may be moved, and only to a peg on which there is no smaller disk, and it would construct a test process for checking these conditions. It would determine that the problem is solved when all the disks are on, for example, the third peg, and would construct a test to determine when that condition is satisfied. Finally, it would generate a list structure showing that all the disks initially are on the first peg. From the evidence of protocols, and of subjects' subsequent problem-solving behavior, this is also the method that human solvers use.

PROBLEM ISOMORPHS

A powerful experimental manipulation for studying problem understanding is to use variant problem instructions, all of which describe *isomorphs* of a single problem. Two problems are isomorphic if the legal problem situations and the legal moves of the one can be mapped in one-to-one fashion on the situations and moves of the other. Then, if situation S' is the isomorph of S, and moves A', B', etc., are the isomorphs of A, B, etc., and if the succession of moves A, B, . . . takes the one system from S to T, then the succession of moves A', B', ... will take the other system from S' to T', where T' is the isomorph of T.

Using a number of isomorphs of the Tower of Hanoi problem, Hayes and Simon (1977) demonstrated that problem difficulty varied by a factor of two to one from one class of problem descriptions (transfer problems) to another (change problems). Moreover, protocols and diagrams produced by subjects showed that they consistently used different representations for the different classes of isomorphic problems. The Understand program behaved in the same way, constructing different representations for both the transfer and change problems. In only one case out of the nearly 100 that have been examined did a subject shift from the more difficult 'change' representation to the easier 'transfer' representation.

The reasons that the change problems take twice as long to solve as the isomorphic transfer problems are not yet fully elucidated. It can be shown, however, that the tests for legality of moves are a little more complex for change than for transfer and this complexity may increase the short-term memory load for the subject who is seeking to understand the problem instructions.

Problem isomorphs can be used to study transfer of training, as in the study conducted by Reed, Ernst, and Banerji (1974). They devised a variant of the Missionaries and Cannibals problem, called the Jealous Husbands problem. It differs from the Missionary–Cannibal problem in that specific husbands are paired with specific wives, and no woman may be left in the company of men unless her husband is present. Experimental results showed that subjects were not better at solving one of these problems if they had previously solved the other. We must conclude that, although subjects may use analogies to help solve problems, there is nothing automatic about the availability of an analogy, and subjects may fail to take advantage of analogies

unless their attention is drawn to them or they are made salient in some other way. Positive transfer between isomorphs in a different type of problem is discussed below in "Construction tasks and other insight problems."

Problem Representation in Mathematics and Physics

Typically a problem given in a mathematical or physics text describes a situation, including quantitative values of some variables, and asks for the value of another variable. The given quantities correspond to the initial state of a problem and the unknown quantity provides the goal. The problem is presented in a natural-language text, as are the instructions for novel problems discussed in the previous section. A physics or mathematics problem differs from a puzzle in that the instructions for the problem, do not provide a description of the problem-solving operators that can be used. It is assumed that the student already knows the operators, from class instruction or from reading the text. The interpretation of puzzle instructions is a representation that can be used by a general problem-solving system such as GPS, whereas the interpretation of a text problem in mathematics or physics is a representation that can be used only by domain-specific problem-solving procedures.

ALGEBRA WORD PROBLEMS

Word problems in algebra describe situations that can be translated into equations, which are then solved to find the values of unknown variables. An early model of solution to word problems, called Student (Bobrow, 1968), showed that the translation can be accomplished mainly by using the forms of sentences in the problem text, and the numerical quantities, with very little knowledge about the objects that are described. For example, in the sentence, "The number of customers Tom gets is twice the square of the number of advertisements he runs," Student does not need to know anything about what customers or advertisements are. but can form the equation $X = 2Y^2$ using the function words is and of in critical ways.

In an empirical study of the solving of algebra word problems, Paige and Simon (1966) found great similarities between human solutions and those given by Bobrow's Student program. Their more skillful subjects, however, used an

intermediate semantic representation in the translation of the English-language problem statements into algebraic equations. Some problems presented descriptions of situations that were contradicted implicitly by real-world knowledge (boards of negative length, nickels worth more than quarters, and so on). The weaker subjects often made accurate syntactic translations of English into equations, as Student does, even though the equations represented nonsense situations. The abler subjects either noticed the contradictions between the statements and their knowledge or translated the statements into equations that were not quite equivalent syntactically, but that represented physically realizable situations.

Another difference between subjects was that those who were more able, unlike the less able, generally drew diagrams of the problem situation that contained all the essential relations from which the equations could be derived.

Both kinds of evidence—the response to 'impossible' situations and the nature of the problem diagrams produced—indicate that the more competent subjects used an intermediate semantic representation of problem situations, rather than a direct translation from English to algebra.

ARITHMETIC WORD PROBLEMS

Detailed analyses of intermediate representations have been worked out for a class of word problems in elementary arithmetic. Riley et al. (1983) and Briars and Larkin (1984) have developed models of representation and solution of word problems that are solved by a single operation of addition or subtraction. Examples of the problems studied are: "Jay had eight books; he lost five of them; how many books does Jay have now?" or "Jay has some books; Kay has seven more books than Jay; Kay has eleven books; how many books does Jay have?"

In the Riley et al. (1983) model, problems are represented by three schemata that provide knowledge of basic quantitative relationships. One schema represents problems involving events that change the value of a quantity, either by increasing or decreasing it, as in the loss of five books; in the problems of the second schema two separate quantities are considered in combination; and in the third schema the problems involve comparison between two separate quantities. (This classification of problems is

not unique; Carpenter and Moser, 1982, Nesher, 1982, and Vergnaud, 1982, have offered similar, though distinct, characterizations.)

Arithmetic word problems are usually classi- fied according to the operations used in their solution, and children are often taught to look for certain key words to decide how to solve the problems This is inadequate, because choice of the correct operation depends on understanding the structure of quantities in the problem, rather than on a single feature corresponding to a key word. For example, 'altogether' is sometimes suggested as a key word for addition, but this is not a reliable cue, as in the problem, "Jay and Kay have nine books altogether; Jay has seven books; how many books does Kay have?"

The model by Riley et al. simulates children's solutions of word problems when small blocks are available for the children to use in solving the problems. The model forms representations of problem texts using the schemata of change, combination, and comparison. Based on the representation that is formed for a problem, the model performs quantitative actions, such as joining two sets of objects together or removing a specified number of objects from a set and counting how many remain. Different versions of the model were formed to correspond to different levels of skill that were observed in a study of children from kindergarten through third grade. The versions differ in the detail with which internal representations are formed (which affects their ability to retrieve infor- mation from earlier steps), and in their ability to perform transformations that provide infor- mation in a form needed to make inferences. The patterns of correct responses and errors observed in the performance of most of the children were consistent with the patterns obtained in the simulation models.

Briars and Larkin's (1984) model constructs less elaborate intermediate representations of problems, and thus relies more on procedures for inferences. Their model uses a schema for representing part-whole relations among sets for some relatively difficult problems.

Physics Problems

The knowledge structures used in simulating solutions to arithmetic word problems are quite general, involving relation between quantities that children probably learn about in their ordinary experience. In technical domains such

as physics, specific instruction is given to teach students the nature of theoretical quantities and the ways in which they combine

Novak (1976) constructed a program called Isaac that builds problem representations in a domain of physics (simple statics) from problem descriptions in English. Isaac uses schemata of physical subsystems (levers, masses, etc), assumed to be understood already by the solver in order to build a compound schema to fit the problem at hand. Thus, it may assemble a wall schema (surface), a floor schema (surface), a ladder schema (lever), and a man schema (mass) to represent a situation in which a man stands on a ladder that is leaning against a wall, assign- ing to each component appropriate numerical quantities and appropriate connections to the others.

Models such as Riley's for arithmetic word problems and Novak's for physics problems are based on the idea that understanding a problem requires schematic knowledge of the quantities in problem situations. The schemata provide knowledge of ways in which quantities are related to one another. These quantitative relations are not expressed adequately in the algebraic formulas that are taught in physics and other quantitative sciences, even though the formulas are based on quantitative relations and students must be able to choose formulas and assign values to variables correctly on the basis of the problem representations that they construct.

The distinction between knowledge of a formula and knowledge of quantities and their relations is illustrated in experiments conducted by Mayer (1974). The experiments were instruc- tional studies, concerned with different methods of teaching the formula for binomial probability One group of subjects received instruction that emphasized calculation, presenting components of the formula with explanations of the cal- culation steps, some practice exercises, and relatively brief explanations of the referents of terms in the formula. Another condition emphasized the information needed in order for students to acquire schematic knowledge. In it, definitions of terms and explanations of relevant concepts, such as the number of combinations and the probability of a single sequence of outcomes, were presented before calculation exercises were given. Tests given following instruction contained a variety of problems,

iv~n to teach
al quantities
in.
ogram called
ntations in a
from problem
ses schemata
masses, etc.),
by the solver
ema to fit the
semble a wall
a (surface), a
schema (mass)
a man stands
a wall, assign-
ate numerical
ections to the

ithmetic word
problems are
ding a problem
the quantities
emata provide
quantities are
e quantitative
quately in the
ght in physics
, even though
tative relations
hoose formulas
orrectly on the
tions that they

nowledge of a
ntities and their
ents conducted
ts were instruc-
fferent methods
mial probability
instruction that
ing components
ions of the cal-
exercises, and
of the referents
other condition
ded in order for
knowledge. In it,
tions of relevant
of combinations
gle sequence of
for calculation
given following
ty of problems.

including some that involved direct application of the formula, and others that required more interpretation. The latter group included word problems, problems that could not be solved because of inconsistent or insufficient information, and problems requiring use of a component of the formula rather than the whole formula The subjects whose instruction emphasized the formula excelled on the problems involving direct use of the formula, but the subjects given more conceptual instruction were more successful on the problems requiring more interpretation.

Several studies have compared the performance of physics students with that of expert physicists to identify some of the components of knowledge that characterize more advanced problem solvers. Three of the characteristics that differentiated the physicists were identified as (1) their use of abstract physics principles in representing problems as well as in providing methods of solution; (2) the strong organization of their knowledge of physics, including relations among principles and recognition of complex patterns of problem features; and (3) the integration of their physics knowledge with general concepts and reasoning processes.

The use of abstract physics concepts by experts was shown in experiments by Chi, Feltovich, and Glaser (1981), who gave subjects a set of 24 physics textbook problems and asked them to sort the problems into groups. Groupings formed by advanced graduate students were based primarily on abstract principles, such as conservation of energy, whereas subjects who had completed a single course in mechanics were much more likely to base their groupings on superficial features such as the kinds of objects (pulleys, levers, etc) that were mentioned in the problems. Chi et al. (1981) also found that experts used abstract physics principles in studies where they reported their thoughts and hunches while deciding on a 'basic approach' to solving the problem. Use of abstract principles as included in a computational model developed by McDermott and Larkin (1978) that simulates the representation of textbook problems by an expert The representation of a problem included a diagram with major components and relations, followed by an abstract description of the theoretical entities such as forces and energies and their interrelations, based on general principles.

Instructional materials designed by Reif and Heller (1981) provide training for beginning students in a procedure for constructing abstract representations of problems. Reif and Heller provided an explicit method for arriving at the kind of problem representation used by experts (although their method was not patterned after the experts' performance, since experts form a representation rapidly and apparently automatically, without easily discerned intermediate steps).

Larkin and Reif (1979) also designed instruction to strengthen students' knowledge of relations among physics principles and their ability to apply principles in solving problems. The instruction grouped principles on a chart and suggested to students that, in applying certain principles it was generally useful to consider the application of other related principles. Qualitative analogies were also used, such as a fluid–current analogy for electric current and a height analogy for potential. Students who received this instruction solved test problems more successfully than students who received instruction in the principles only, without the organization and qualitative analogies.

Experts in various domains have been shown to have superior skill in recognizing complex patterns of information in the domain of their expertise. This phenomenon has been demonstrated in chess (Chase & Simon, 1973), go (Reitman, 1976), electronics (Egan & Schwartz, 1979), computer programming (McKeithen, Reitman, Rueter, & Hirtle, 1981), and radiology (Lesgold, Feltovich, Glaser, & Wang, 1981). A highly developed skill in pattern recognition may provide an explanation for the finding obtained in several studies that expert problem solvers tend to work forward from the given information to the unknown, whereas novices work backward from the unknown, searching through a series of subgoals for formulas that can provide the needed quantities (e.g., Simon & Simon, 1978). Applying formulas involves using more complex patterns of known values of variables, which experts have probably learned to recognize directly, thus avoiding the more laborious searches that novices conduct (Larkin, 1981). This view is supported by Malin (1979), who found that subjects were more likely to adopt a forward-search strategy to solve problems if the formulas they were using had an obvious organization than if the formulas did not fit together in any evident way.

A third characteristic of experts' knowledge is that their domain-specific knowledge (e g , in physics) is integrated with powerful general concepts and procedures for making inferences. An example comes from Simon and Simon (1978) who obtained protocols from a novice and an expert on problems from a high school physics text. One problem was. "An object dropped from a balloon descending at four meters per second lands on the ground 10 seconds later. What was the altitude of the balloon at the moment the object was dropped?" The novice subject's solution had the properties of means ends analysis, using the formula $s = v_0 t + .5at^2$. In contrast, the expert calculated a quantity that he called the total additional velocity by multiplying the time by the gravitational constant (i.e., $10 \times 9.8 = 98$), he then added that to the initial velocity to obtain the final velocity ($98 + 4 = 102$), took the average velocity [($4 + 102)/2 = 53$], and found the distance by multiplying the average velocity by the time of 10 seconds ($53 \times 10 = 530$ meters). The expert apparently had a representation of the problem in terms of physical quantities that enabled him to apply general procedures, such as computing components of velocity and taking an average, whereas the novice was restricted to using the formulas that were provided in the text. Relations between technical knowledge and general concepts have been investigated theoretically by deKleer (1975) and Bundy (1978), who developed models of physics problem-solving that combine general knowledge about the motion of objects on surfaces with knowledge of formulas in kinematics, and by Larkin (1982) who studied the use of spatial information in the solving of hydrostatics problems.

UNDERSTANDING OF STRUCTURE AND PRINCIPLES

The integration of problem-solving knowledge with general conceptual structures has also been used to characterize structural understanding as discussed by Wertheimer (1945/1959), and the understanding of general principles, including the relation of abstract properties of number (cardinality, order, one-to-one correspondence) to children's cognitive procedures for counting.

The understanding of structure has been investigated theoretically by Greeno (1983) on a problem, discussed by Wertheimer (1945/1959),

of proving the congruence of vertical angles Wertheimer distinguished between a relatively mechanical process for generating the proof. involving the use of algebra without cognizance of spatial relations in the problem. and a more meaningful process based on part whole relations between pairs of angles and operations to remove a part that is included in each of two whole angles. Greeno's model simulates the more meaningful process by using a schema that represents part whole relations in a general way and applying problem-solving operators that make inferences based on the part whole structure. Data were available in the form of protocols from students working on the vertical angle problem after they had learned to solve other problems with similar part whole structure involving line segments. The model simulates learning in the line-segment situation. Once the learned problem-solving operators are integrated into the part-whole schema, the model can apply this knowledge when it encounters the vertical-angle problem. The model thus provides an explanation for transfer that occurs between problems in different domains, with a characterization of structural understanding based on schematic representation. An account of transfer based on acquisition of a schema in a different problem domain is discussed below in "Construction tasks and other insight problems "

A similar idea was used by Resnick, Greeno, and Rowland (described by Resnick, 1983) in analyzing children's understanding of a procedure for subtraction with multidigit numbers According to their analysis, children who understand the procedure have a representation that includes general relations—such as part–whole relations between quantities represented by individual digits and the quantities represented by combinations of digits and constraints such as the requirement that the total value of a number remain unchanged when borrowing is used. The analysis focused on knowledge acquired in meaningful instruction (cf. Brownell, 1935), in which children were shown the correspondence between subtraction with numerals and an analogous subtraction procedure using blocks Resnick et al. (in Resnick, 1983) hypothesized that the understanding was achieved through acquisition of a schema, involving part-whole relations, that was general enough to apply to both—the numerals and the blocks.

Efforts are being made to develop rigorous and explicit characterizations of knowledge that includes implicit understanding of general principles (cf. Judd. 1908; Piaget. 1941/1952). A representation of preschool children's understanding of the principles of counting has been formulated by Greeno et al. (1984). Their analysis was based on evidence presented by Gelman and Gallistel (1978) that young children have significant understanding of principles such as cardinality, order, and one-to-one correspondence, rather than a simple 'mechanical' knowledge of counting procedures. The evidence includes their performance in novel situations, such as being asked to evaluate counting performance by a puppet that sometimes makes errors, or counting with the novel constraint of associating a specified numeral with a particular object. Greeno et al. (1984) also proposed an analysis of conceptual competence to represent children's implicit understanding of principles. Conceptual principles are represented as schemata that incorporate constraints on correct counting and express general properties, such as the part–whole relation between the counted objects and the whole set. The conceptual principles are related to procedures of counting by a set of planning rules, which permit derivation of procedures from the schematic representations of the principles.

PROBLEMS OF DESIGN AND ARRANGEMENT

Problems discussed in this section require finding an arrangement of some objects that satisfies a problem criterion. Simple examples include puzzles in which the objects are given in the problem situation. For example, an anagram presents some letters, and the task is to find a sequence of those letters that forms a word. In more complex cases, the problem solvers must provide the materials based on their own knowledge. Examples are writing an essay or a computer program.

The problem space in a problem of design includes the objects that are given to or are known by the problem solver. The space of possible solutions is the set of arrangements that can be formed with the available objects. The problem goal is to construct an arrangement that meets a criterion, which may be either specific or nonspecific. An anagram problem has a specific criterion: the sequence of letters should form a word. A written composition has several less specific criteria, such as clear exposition, persuasive argument, and an entertaining style. Many problems of design have a mixture of specific and nonspecific criteria. For example, a problem in computer programming may combine a criterion of a specific function to be computed with less specific criteria, such as efficient computation and clarity of structure.

Satisfying constraints is an important factor in solving problems of design. The metaphor that best characterizes typical solution processes is 'narrowing the set of possibilities' rather than 'searching through the set of possibilities.' Although it is entirely possible to describe the solution process as a search, the main steps in this search lead to the acquisition of new knowledge that rules out a whole set of problem states as potential solutions—a wholesale approach to the reduction of uncertainty. The use of constraints is important because the set of possible arrangements is usually very large, compared to those that satisfy the problem criterion.

Problems of design are differentiated from the transformation problems discussed above in "Well-Specified Problems," in both the nature of the goal and the set of alternatives that are considered. In a transformation problem such as the Tower of Hanoi or in finding a proof for a theorem, the goal is a specific arrangement of the problem objects, such as a specific location of all the disks in the Tower of Hanoi or a specific expression to be proved in logic. Thus, the question is not what to construct, as it is in a design problem, but how the goal can be constructed with the limited set of operators that are available. The search for the solution of a transformation problem often examines one problem situation after another, uncovering knowledge that helps point the direction of the search toward the goal situation.

Viewed in another way, however, transformation problems and problems of design are very similar in structure. The solution of a transformation problem is a sequence of actions that changes the initial problem situation into the goal. The solution process can be considered as the construction of an appropriate sequence of actions, involving search in the very large space of possible sequences. This view emphasizes

similarities between problems of transformation and of design, which are especially apparent when the solution of transformation problems includes planning.

Problem solving in design is discussed in four parts: (1) Two simple problems of forming arrangements—cryptarithmetic and anagrams—provide paradigms for analyzing search among sets of possible arrangements; (2) problems in which an arrangement of objects is already presented, and the task is to modify the arrangement according to some criterion (e.g., Katona, 1940); (3) 'insight' problems that depend on finding a successful formulation or representation of the problem; and (4) more complex problems of composition and design, including the composition of essays and musical pieces, the design of procedures, and the formation of administrative policies.

Simple Problems of Forming Arrangements

Cryptarithmetic Problems

In cryptarithmetic problems, digits are arranged to form a correct addition problem, constrained by a set of letters for which the digits are to be substituted (Newell & Simon, 1972). One of the best known examples follows:

DONALD

+ GERALD

= ROBERT

The task is to replace each letter in the array with a distinct digit, from 0 to 9, the same digit replacing a given letter in all its occurrences (no digit being used for more than one letter). To make the problem easier, the solver is usually told that D = 5.

The cryptarithmetic task was apparently first studied by Bartlett (1958), who reported some retrospective protocols of subjects in his book on thinking. Subsequently, Newell and Simon (1972) carried out extensive analyses of thinking-aloud protocols for cryptarithmetic problems. From this work, we now have quite a clear picture of how human subjects approach such problems.

There are 10! = 3,628,800 ways of assigning ten digits to ten letters. Most subjects, without calculating this number, realize that it is very large, and do not even attempt to solve the problem by making random assignments and testing them. Instead, they look for information in the form of constraints that permit values to be assigned to particular letters at once If that can be done, the number of possibilities declines rapidly Simply giving the information that D = 5 already reduces the possible solutions by a factor of 10, that is, to 362,880—still a large number!

The constraints in cryptarithmetic problems that sometimes make systematic elimination possible derive from the fact that each column of the literal array must be translated into a correct example of addition (subject to carrying into and out of the column). Thus, as soon as it is known that D = 5, the sixth column can be processed to produce the inference that T necessarily equals 0, and that 1 is carried into the fifth column. This single inference reduces the remaining set of possible assignments by a factor of nine to 40,320.

Next, consideration of the second column allows the subject to infer that E is equal to 0 or 9. Since 0 has already been preempted by T, we have E = 9, reducing the possible assignments to 5,040. A few more steps of reasoning, based on information contained in columns 1 and 5, allow the subject to infer that R = 7, reducing the possible assignments to 720. An inference in column 4 gives A = 4 (120 possibilities remain); and an inference on column 5 gives L = 8 (leaving only 24 possibilities). From column 1, G = 1 (leaving 6 possibilities), and now the remaining digits must be assigned to N, O, and B, a task easily carried out by trial and error.

Newell and Simon (1972) obtained thinking-aloud protocols of subjects solving cryptarithmetic problems. Problem behavior graphs were constructed based on the protocols, and a detailed model of one subject's problem-solving processes was developed in the form of a production system. (This methodology is discussed above in "Discovering Proofs in Logic.") In the model several productions represent a problem-solving strategy. These productions set goals of examining a column or the occurrences of a variable; they make decisions on the assignment of a value to a variable or the testing of a candidate value, and they perform other general functions. There are also a few dozen productions that represent the operation of specific processes. One, called *Process Column*, contains 26 productions; others are considerably simpler.

The productions in this process examine the letters in a column and use any information that has been gathered about them to make further inferences The subject's performance, recorded in a problem behavior graph, was compared in detail with the model, and approximately 80 percent of the protocol units were explained by processes in the model.

Protocols obtained from five subjects were consistent in their general characteristics of problem-solving processes. They also revealed significant individual differences, and these can be interpreted as differences between the problem spaces of the individual problem solvers. All the subjects made use of their knowledge of arithmetic in order to make inferences, and all subdivided the problems into subproblems involving the columns. There were important differences among subjects in their strategies for selecting columns to work on and in their use of specific constraints for making inferences.

For an efficient solution of this problem, subjects must use the search heuristic of attacking the most constrained columns first, since most information can be extracted from a column in which the assignment of one or more letters has already been made, or in which the same letter occurs twice. Some subjects used this selection heuristic immediately; others began by attacking the columns systematically, from right to left, and only later abandoned that strategy for the more powerful one Subjects who did not use the heuristic usually failed to solve the problem.

Another factor that influenced success was the use of specific constraints. The problem spaces of some subjects included rules of parity. For example, one of the inferences needed in order to conclude that $R = 7$ is that, whatever R's exact value, it must be an odd number. This is inferred by processing column 5, containing two Ls whose sum must be even, and the carrying of 1, making the total an odd number. Subjects whose problem spaces did not include the parity constraints were generally unable to solve their problems.

Even subjects who used the available heuristics and constraints for efficient elimination found the DONALD + GERALD problem difficult. Most of their difficulties arose from one or both of two sources. One such source is the making of conditional assignments (e.g., "suppose that $L = 1$") Then, if the assignment was wrong and they arrived at a contradiction, they may

have been unable to remember which prior number assignments they had inferred definitely and which they had postulated conditionally. Another source of difficulty involved errors of inference, resulting in incorrect assignments. For example, from the fact that $R = 7$ some subjects concluded that $L = 3$ (with a carry from the sixth column), ignoring the possibility that L might be 8, with a carry into the fourth column. When $L = 3$ led to a contradiction, they found it difficult to discover the cause.

Errors of inference are forms of the errors of syllogistic reasoning discussed below in "Propositional and categorical syllogisms." In the example just cited, subjects appeared to infer from the premise, "if $L = 3$ then $R = 7$" and the premise "$R = 7$," the conclusion "$L = 3$," an example of the classical fallacy of inferring the antecedent from the consequent. They did not notice that $L = 8$ also implies $R = 7$. Thus, the cryptarithmetic task draws on reasoning processes as well as search processes.

Nothing in the behavior of subjects solving cryptarithmetic problems suggests that they decide consciously to treat it as a constraint problem rather than a search problem. In fact, their behavior can be described as a search through the space of possible assignments, and Newell and Simon's analysis took +h s point of view. What distinguishes it from search in many other problem spaces is that the problem is factored into 10 separate but interdependent searches for the individual assignments. Success in each of these searches constrains the problem space by reducing the number of alternative possibilities for the remaining assignments, and by providing additional information about some of the columns. Hence, it is not unlike an ordinary search in which each step of progress provides clear feedback of information that the right track is being followed.

Anagrams

Anagrams are strings of letters that can be rearranged to form words, for example, *thgli* → *light*. The problem space of an *N*-letter anagram contains *N!* possibilities, and therefore, increases rapidly with *N*. The solution process can be viewed as a search through this space of permutations of the letters, but most persons presented with an anagram use various heuristics to speed up the search. One of these is to pick out initial combinations of letters that are pronounceable

(e.g., *lt* or *lt* in the example above), and then try to complete a word with the remaining letters. Imposing the condition of pronounceability on solution attempts may restrict the search space considerably.

The course of the search is also much influenced by the structure of long-term memory. For example, if there are two possible solutions to an anagram, the one corresponding to the more frequent and familiar word is likely to be found by most of the subjects. Moreover, the solution can be primed by presenting the word to the subject, or a semantically related word, some time before the anagram task is taken up (Dominoswski & Ekstrand, 1967).

Perceptual factors may affect performance on anagram tasks. Anagrams that are already words (e.g., *forth → froth*) or are easily pronounced (e.g., *obave → above*) take longer to solve than those without such properties (Beilin & Horn, 1962). This finding is consistent with Gestalt principles that meaningful forms resist restructuring. Gavurin (1967) found a correlation of 54 between success in solving anagrams and scores on a standard test of spatial abilities. When the subject was provided with tiles that could be rearranged physically, the correlation disappeared, indicating that the original relation had to do with the perceptual ability to operate on visual or auditory images.

It is easy to induce a problem-solving set in anagram solving by presenting subjects with several anagrams that call for the same permutation (say, 5 4 1 2 3) of the letters. If an ambiguous anagram (one with several possible solutions) is then presented, most subjects will find the solution requiring the same permutation rather than the alternative solution (Rees & Israel, 1935).

Thus, subjects' behaviors on the anagram task combines search (generating possible solutions) with constraint satisfaction (rejecting unpronounceable initial segments). The process of alternative generation, in turn, is influenced by long-term memory organization and priming, and by the subject's skill in forming and holding in short-term memory the permutations of the stimulus.

Problems of Modifying Arrangements

Unlike the problems just discussed in which arrangements are formed from materials pro-

vided that the problem solver must put together to satisfy a specified criterion, we now turn to problems in which an arrangement of objects is presented, and the task is to modify the arrangement. Perceptual processes important to the solution of these problems involve recognition of general features and complex patterns.

These problems combine features of the transformation problems discussed above in "Well-Specified Problems" with features of design problems. Like design problems, a goal is specified as a general criterion rather than as a specific state that the problem solver tries to produce. At the same time, in these problems significant restrictions on the operators can be used to change the situation. Therefore, the problems can be conceptualized as search either in a space of possible arrangements or in a space of possible sequences of moves.

Matchstick Problems

Figure 9.10 shows a matchstick problem used by Katona (1940). The 16 matches form five squares; the task is to move exactly three matches in such a way that the matches form only four
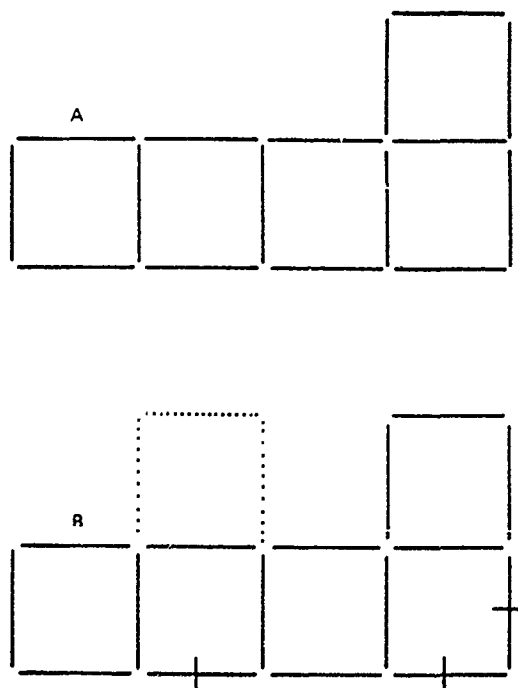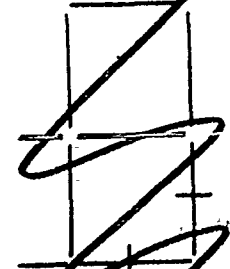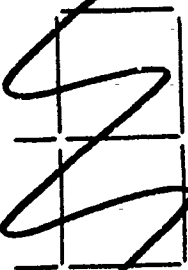

Figure 9.10. A matchstick problem used by Katona (1940).

squares. and all the matches serve as sides of squares. Katona tested subjects under three conditions: (1) in rote learning (subjects were shown and required to learn a specific solution), (2) with a logical condition for the solution (subjects were taught that in the solution, each match formed a side of one and only one square), and (3) with a heuristic for solving the problem (subjects were told "you need to open up the figure").

The subjects learned the solutions and then were tested on transfer tasks (different initial arrangements of the matches and different numbers of squares). Differences in the ease of learning the solution were minimal, with the rote solution being learned most rapidly. Two weeks later they were invited back and tested for their memory of the solution. In the test of transfer and retention, the logical and heuristic solutions far outshone the rote solution, and the heuristic solution scored slightly better than the logical. From this evidence Katona concluded that problem-solving knowledge and skills are better transferred and retained when the learning is meaningful than when it is rote.

The experimental manipulations leave implicit, however, the theoretical import of the term 'meaningful'. Why does meaningful learning facilitate retention and transfer, and why is the heuristic form of the instruction superior to the logical form?

With respect to transfer and retention, meaningful learning involves the same issues as structural understanding (discussed above in "Problem Representation in Mathematics and Physics"). Transfer is facilitated because, with more meaningful instruction, subjects acquire knowledge that can be applied more generally —in particular, to the new problems presented in the test as well as the problems used in training. It is easy to see why this occurs; the meaningful instruction can be applied to matchstick problems generally, while a specific solution sequence applies only to a single problem.

As for retention, meaningful forms of instruction may provide more redundancy, and hence, more opportunity to recover from partial forgetting. The general principles of single versus double function and of loosening or condensing the figure are constraints that can be used to limit search for information in memory, or to reconstruct solutions that are only partly remembered.

The difference between the two meaningful procedures appears to derive from the distinction between generators and tests. The instruction to 'open up the figure' provides a constraint on the selection of an operator – it suggests something to do, however vague, relative to a general property of the figure that can be perceived. The rule, 'each match must form a side of one and only one square,' constrains solution arrangements. It provides a test that can be applied to an attempted solution, but does not suggest what to move to produce the solution in the first place. In fact, the matches that have to be moved to solve the problem are not those with double function but rather those that already lie on the side of only one square. In this situation, at least, the knowledge that facilitates a solution most effectively increases the selectivity of the move generator rather than of the candidate solution states.

Katona noted that the heuristic of opening the figure or closing gaps uses a feature that is important in the perception of form, the Gestalt principle of good continuation. Attending to that feature and considering moves to adapt an arrangement to it constitutes a general strategy for solving matchstick problems.

Chess and Go

Board games offer problems of the same general form as matchstick problems. An arrangement of objects is presented—the current situation in the game—and a player has the task of selecting a move or move sequence. Some criteria for a good solution are quite specific (e.g., white to mate in four moves); more often they are general, involving a goal to achieve a stronger position. Recent experiments comparing the performances of individuals who differ in skill show the importance of knowledge in the recognition of large numbers of complex patterns that occur during games.

In complex games, as in other domains in which some people become expert, problems that would be difficult or impossible for novices are often solved 'instantly' by experts—that is, in a few seconds. For example, a chess grand master, who is presented with a position from an actual but unfamiliar game and asked to recommend a move, will usually be able to report a good move, often the best move, in five seconds or less (deGroot, 1965). In a 'blitz' game, the same player, required to move within 10

seconds, will probably be unable to play at the grand master level but will achieve master level. Players at this level will be able to play 50 or more opponents simultaneously, with a high level of success, taking only a few seconds for each move. When experts are asked how they solve problems so rapidly, they may reply, "I use intuition," or, "I use my judgment."

The nature of this intuition or judgment has been clarified by experiments on skill in chess by deGroot (1965) and Jongman (1968) and repeated and extended by Chase and Simon (1973), and on skill in the game Go by Reitman (1976). In the experiments on chess skill, a chessboard with a position from a game (containing perhaps 25 pieces) is shown to a subject for 5 to 10 sec. The subject is then asked to reconstruct the position. Chess grand masters and masters can perform this task with 90 percent accuracy. Ordinary players can replace only five or six pieces correctly (20 to 25 percent accuracy). In a second condition the task is the same, except that the pieces are now arranged on the chessboard at random, rather than in a pattern that could have arisen in a game. In this condition, the performance of masters falls to the level of ordinary players—both can replace, on average, only about six pieces. This second part of the experiment demonstrates that the chess masters do not have any special powers of visual imagery.

Reitman's (1976) study of skill in Go had similar results. Go is a game of territory played on a 19 × 19 grid. The pieces are round 'stones' differing only in color for the two players, black and white. An experienced subject (not as strong as a professional player), was able to reproduce 66 percent of the pieces of meaningful patterns, compared to 39 percent for a beginner who had played about 50 games. On random patterns the players replaced 30 percent and 25 percent, respectively or an average of five to seven stones.

This experimental procedure has been applied to the pattern-recognition abilities of experts in several other domains; see "More Complex Tasks of Composition and Design" and "Diagnostic Problem Solving," below.

The behavior of the chess and Go experts in the perception and memory task can best be explained as a function of their chess and Go experience. As a result of thousands of hours spent at game boards, they become familiar with many configurations of three, four, or more pieces that recur again and again in games. For example, a configuration known as a 'fianchet-toed castled Black King's position' occurs in perhaps one in ten games between expert chess players. This configuration is defined by the positions of six pieces. It has been estimated that a chess master has stored in long-term memory not fewer than 50,000 familiar patterns of this kind (Simon & Barenfeld, 1969; Simon & Gilmartin, 1973). This number is comparable to the 50,000 words in the vocabulary of a typical college graduate, or perhaps the total number of human faces a gregarious person learns to recognize over a lifetime.

When a chess master is confronted with a chessboard on which the pieces are arrayed in a 'reasonable' way, he can store this information in short-term memory in a half dozen or fewer 'chunks'—familiar configurations. The ordinary player, or the chess master confronted with a randomly arrayed chessboard, must store the information piece by piece, and hence, can hold the positions of only half a dozen or so pieces in short-term memory.

The skill that the expert acquires does not consist simply of being able to recognize familiar stimuli or configurations of stimuli. As deGroot showed, the recognition of perceptual features on the chessboard reminds the grand master of moves that are potentially good when those features are present. Indeed, we should expect the expert's knowledge for pattern recognition to be integrated with strategic knowledge so that the patterns the expert has learned to recognize are those relevant to the choices of moves and plans encountered in games.

The importance of game strategy in perception and representation of complex patterns was shown in an experiment by Eisenstadt and Kareev (1975). The games Go and Gomoku are played with entirely different rules, though on the same board and with the same kinds of pieces. Two groups of subjects, who knew how to play both games, were shown the same patterns of stones on boards. One group was told that the patterns were from a game of Go, and the other from a game of Gomoku. When they were subsequently asked to recall the patterns, the subjects in the first condition better recalled the pieces that were critical to selecting the correct move in the Go position, whereas the others recalled better those pieces that were critical to selecting a move in the Gomoku position. Thus,

in in games. For
n as a 'franchet-
ition occurs in
een expert chess
defined by the
been estimated
ed in long-term
familiar patterns
d 1969; Simon &
is comparable to
ilary of a typical
e total number of
person learns to

onfronted with a
s are arrayed in a
this information
if do en or fewer
ops. The ordinary
onfronted with a
must store the
id hence, can hold
zen or so pieces in

acquires does not
recognize familiar
stimuli. As deGroot
eptual features
ie grand master of
good when those
we should expect
attern recognition
gic knowledge so
rt has learned to
to the choices of
d in games.
strategy in percep-
complex patterns
by Eisenstadt and
o and Gomoku are
t rules, though on
the same kind of
s, who knew how to
n the same patterns
roup was told that
ame of Go, and the
v. When they were
ll the patterns, the
n better recalled the
selecting the correct
whereas the other
that were critical to
noku position. Thus,

in the face of a complex stimulus situation,
attention to a particular task determines the
sequence in which information is extracted from
the stimulus and the patterns in which it is
organized.

Studies of specific knowledge structures that
integrate strategic knowledge and knowledge
for recognizing patterns have been carried out
by Wilkins (1980) in a model of choosing moves
in chess, and by Reitman and Wilcox (1978) in a
model of playing Go.

Wilkins's (1980) model represents board
positions by recognizing concepts, such as
Attack and Safe, based on relations among
pieces. The model uses schemata that correspond
to the concepts in proposing and evaluating
plans. In formulating a plan, a concept such as
Safe or Defend-Threat can be set as a goal; the
schema for each concept includes conditions
that are required to satisfy the goal. The model's
strategy of using proposed plans to guide its
search restricts the set of moves it considers,
enabling relatively thorough evaluations. The
model is successful in solving problems of
choosing moves in middle game positions that
are sufficiently difficult to be used in a standard
chess textbook.

Reitman and Wilcox's (1978) model simulates
representation of board positions and changes
of board positions in Go. The model forms a
multilevel representation with low-level units
such as strings and chains of stones, and higher-
level units called groups and fields involving
collections of points and their surrounding
stones. The representations include features
that are relevant to Go tactics, such as the
stability of a group of stones. Perceptual activity
is organized according to several structures
including lenses, which monitor changes on the
board relevant to relations between groups of
stones, and webs, which monitor changes on
radii and circumferences around groups. The
model's capabilities for representation, combined
with some relatively low-level processes for
selecting moves, is similar to a human player
with the experience of playing 40 or 50 games.

The ability of experts to recognize complex
patterns of information related to a highly
integrated structure of actions has been found
in other domains in which expertise has been
analyzed. The importance of knowledge for
representing problems in physics was discussed
above in "Problem Understanding; Represen-

tation," and similar conclusions were found for
medical diagnosis and electronic troubleshoot-
ing ("Diagnostic Problem Solving," below). It is
reasonable to conjecture on present evidence
that high levels of expertise generally require
tens of thousands of perceptual 'chunks' relevant
to the domain. In domains where the minimal
time required to become a world-class master
has been measured, the estimate turns out to
be about a decade (Hayes, 1981; this finding is
discussed below for musical composition in
"Problems of Composition").

Construction Tasks and Other Insight Problems

Much attention in research has been given to
problems in which some physical device or
arrangement is required, often to satisfy a
functional criterion. An example is Duncker's
(1935/1945) famous 'tumor' problem in which
a patient has a stomach tumor that is to be
destroyed by radiation without damaging the
surrounding healthy tissue. How is it to be
done?

The source of difficulty in construction
problems differs from the problems discussed
above in "Simple Problems of Forming Arrange-
ments" and "Problems of Modifying Arrange-
ments" where difficulty arises from the large
number of possible solutions. The tumor problem
and other 'insight' problems are difficult, pri-
marily because most of the candidate solutions
considered are ruled out by the constraints of
the problem. In the tumor problem, for example,
simply directing the rays to the tumor would
destroy all the tissue along their path; to open a
path to the tumor by surgical procedures would
cause intolerable damage, and so on. The 'text-
book' solution to the tumor problem calls for
irradiating the tumor from many different angles,
and hence, via many different paths through the
surrounding tissue. By this means a large
quantity of radiation is concentrated on the
tumor, while each path of surrounding tissue is
subjected to only a small fraction of that amount

Solving the tumor problems and similar
insight problems often depends on finding a way
to represent the problem so that the solution
becomes obvious. Achievement of such a repre-
sentation, corresponding to a moment of insight,
is a phenomenon of great interest, especially in
relation to issues of cognitive organization

in Gestalt psychology. In problems such as cryptarithmetic and anagrams, the problem space is easily constructed, and problem-solving activity consists of searching in the set of possibilities that arise in that space. On the other hand, in insight problems such as the tumor problem, the problem solver's initial representation usually provides an inadequate problem space, one in which a solution will not be found. Problem solving involves a construction of several problem spaces—only to be discarded as factors are discovered that make each of them inadequate—until a successful representation is found. Processes of problem representation thus play a central role in the solution of these problems of construction. The process can be characterized as a search for alternative ways to represent the problem. However, the usefulness of such a characterization is limited unless the set of alternative representations can be specified more definitely than we are at present able to do.

Duncker (1935/1945) emphasized the *demand*, the condition to be met by the problem solution, as the chief source of solution proposals. The initial proposals are not unmotivated, but they are faulty in not attending to all the conditions a solution must meet. False analogies may produce inadequate solutions because of their failure to match the actual situation in crucial dimensions. At the same time, Duncker stressed that the proposals are not produced by simple association:

> In short, it is evident that such proposals are anything but completely meaningless associations. Merely in the factual situation, they are wrecked on certain components of the situation not yet known or not yet considered by the subject.

> Occasionally it is not so much the situation as the demand, whose distortion or simplification makes the proposal practically useless (p 3)

By constructing a taxonomy of correct and inadequate solutions to the tumor problem, Duncker showed how the solution-generating process can be understood as a process of means–ends analysis. His taxonomy can be depicted in outline form:

Treat tumor by rays without destroying healthy tissue

Avoid contact between rays and healthy tissue
    Use free path to stomach
      Use esophagus
    Remove healthy tissue from path of rays
      Insert a cannula
    Insert protective wall between rays and tissue
      Feed substance that protects
    Displace tumor toward surface
      Apply pressure
Desensitize the healthy tissue
    Inject desensitizing chemical
    Immunize by adaptation to weak rays
Lower intensity of rays through healthy tissue
    Postpone full intensity until tumor is reached
    Use weak intensity in periphery, strong near tumor
      Use a lens

Duncker described the solution process as the successive development or reformulation of the problem. Working both forward and backward may contribute to the process. Seeing a stick may give a chimpanzee the clue to obtaining a banana that is out of reach. Alternatively, the banana's being out of reach may lead the chimpanzee to look for an object that could be used to reach it (cf. Köhler, 1929). Mistakes may also call attention to features of the problem situation that must be incorporated in the solution, and hence, may lead to new solution attempts.

From the idea that the solution of a problem depends on an appropriate formulation, it would be expected that hints could be used to make problems significantly easier. One experiment on the effects of hints used a problem of constructing a hat rack, invented by Maier (1945). Two sticks and a clamp were given. The hat rack could be constructed by clamping the sticks together so that the assemblage was long enough to be wedged between the floor and the ceiling. Subjects usually began by either laying one stick on the floor and clamping the other stick to it vertically, or standing both sticks on the floor in an X or inverted V shape. Neither of these structures is stable. If the experimenter said, "In the correct solution, the clamp is used as a hanger," the solution was facilitated somewhat, mainly by reducing attempts made with one stick lying on the floor. If the experimenter said, "In the correct solution the ceiling is part of the

and healthy

om path of rays

tween rays and

protects
urface

ue
nica
to weak rays
ough healthy

until tumor is

periphery strong

lution process as
r reformulation of
orward and back-
process. Seeing a
the clue to obtain-
ich. Alternatively,
ach may lead the
ject that could be
29). Mistakes may
es of the problem
corporated in the
id to new solution

problem
Th... thought it would
l be used to make
c. One experiment
a problem of con-
d by Maier (1945).
given. The hat rack
amping the sticks
ge was long enough
or and the ceiling.
either laying one
the other stick to
sticks on the floor
e. Neither of these
perimenter said, "In
lamp is used as a
cilitated somewhat,
its made with one
experimenter said,
ceiling is part of the

construction." the solution was facilitated even more forcefully, by reducing attempts that used only the floor as support (Burke. Maier. & Hoffman, 1966).

A potential source of problem solutions is analogy with similar problems. Gick and Holyoak (1980) gave Duncker's tumor problem to subjects. some of whom had studied a story in which a fortress was taken by a converging attack The subjects who were familiar with the military problem were more successful than control subjects in solving the tumor problem. An important factor was the inclusion of an instruction that the fortress story might provide a useful hint for solving the problem. With the hint, most subjects found the convergence solution to the tumor problem, but without the hint only about half as many subjects found that solution, even though they had read the story and recalled it in a test.

In a subsequent study, Gick and Holyoak (1983) examined conditions favoring the spontaneous use of an analogy. Asking the subjects to summarize the military story, rather than recall it, had little effect, and giving them a verbal statement or diagram showing the convergence principle did not noticeably increase their use of the analogy. However, more solutions were proposed by subjects who read two stories involving convergence, summarized both of them, and discussed ways in which the stories were similar. Gick and Holyoak concluded that those subjects acquired a schema with the idea of convergence represented in a general way, and that such a schema is more likely to be used than is a specific analogous problem. (In "Problem Representation in Mathematics and Physics," above, a similar hypothesis was offered.)

Duncker (1935/1945) also studied problems that required subjects to construct some item out of potential components, including some inessential components, that were provided. He showed that the problems could be made difficult by presenting one of the components in such a way that it was conceptually 'unavailable' for its required function. For example, in one problem the building materials were a candle, matches, and a box full of thumbtacks. The task was to mount the candle on a wall so that it could burn without dripping wax on the floor. The problem could be solved by thumbtacking the box to the wall, then mounting the candle in it.

This problem was so difficult that fewer than half the subjects in one experiment were able to solve it in 20 minutes (Adamson, 1952). When the problem was presented to another group of subjects with the thumbtacks lying on a table, and the box empty, 86 percent solved it in less than 20 minutes. The phenomenon underlying this finding has been labeled 'functional fixity.' When an object is performing, or has recently been used to perform some function, subjects are less likely to recognize its potential use for another function.

Birch & Rabinowitz (1951) demonstrated a similar phenomenon, using another problem originally studied by Maier (1931). In a room where two strings were hanging from a ceiling, too far apart to be reached simultaneously, the task was to tie them together. This could be accomplished if a heavy object was tied to one string and the string was swung as a pendulum. This string could be grasped as it swung toward the subject, who meanwhile had the other string in hand. Two objects, an electric switch and a relay, were available for constructing the pendulum. The subjects had used either the switch or the relay (but not both) in a previous task. Of ten subjects who had used the relay previously, all used the switch to construct the pendulum; of nine who had used the switch, seven used the relay to construct the pendulum. Of six subjects who had used neither object previously, three used the switch and three the relay to construct the pendulum.

Several findings support a hypothesis that functional fixity results from a decrease in the likelihood of noticing certain critical features of objects in the situation, such as the flatness of a box (in use as a container), or the heaviness of a switch (after use in a circuit), or the features in functional fixity may be quite different in different situations, involving restrictive hypotheses about general classes of solutions in some cases, and simple competition between feature-recognition processes in others.

Some of the findings that support this explanation involve demonstrations that the solution of problems can be influenced even by very low-level perceptual factors. For example, in the pendulum task, the idea of making one string swing so that it would be reachable by someone holding the other one does not occur readily to most subjects, even in the presence of one or more heavy objects. Maier (1931) showed that

this idea occurred immediately to many subjects who had not previously thought of it, when the experimenter casually brushed against the string and set it swinging. Glucksberg and Weisberg (1966) presented pictures of the materials available for use in solving Duncker's candle problem, and found that solutions were markedly increased when the label 'Box' was included in the picture. A process of noticing features of objects that can be related to the problem goal (Dur er's 'suggestions from below') probably ple   significant role in the solution of construci   blems, as Weisberg and Suls (1973) conc.   in their theoretical analysis of solution   esses for the candle problem. Results consistent with that idea were obtained by Magone (1977), whose subjects produced a greater variety of solutions to Maier's two-string problem if they were initially prompted to consider features of objects than if they were prompted to seek a solution of a specified kind, such as extending one of the strings or causing a string to swing back and forth.

The Einstellung effect discussed above in "Einstellung" is similar in character to functional fixity in that in both effects, previous experience influences the availability of alternative steps toward solution. The processes responsible for the two effects are probably analogous in a subtle but significant way since in both, a form of search is made less likely than it would normally have been. With Einstellung, the previous use of a solution path suppresses a search for problem-solving operators. With functional fixity, the search for features of objects that could be useful to a solution is suppressed.

Another 'insight' problem that has been studied is the nine-dot problem. A three-by-three matrix of dots is given, and the task is to connect all the dots with four straight lines without any retracing. Several lines may pass through the same dot. The problem is difficult, most subjects do not think of drawing lines outside the space defined by the matrix of dots, as is required for the solution. The difficulty is apparently another instance of a restricted domain of search, but the obvious hypothesis of a restriction based on the spatial arrangement is not supported by data. Weisberg and Alba (1981) instructed their subjects to draw lines outside the square of dots, but that had little effect. However, when they

gave an easier problem requiring drawing lines beyond the region that contained dots to other subjects, subsequent solution of the nine-dot problem was facilitated A reasonable interpretation is that the easier problem led the subjects to consider problem-solving operators that were not in the problem space of subjects who had not solved the simpler problem first. This finding involves the same principles as the finding of Katona (1940; see "Matchstick Problems") that a heuristic for choosing operators is more effective than a test applicable to the results of operators.

More Complex Tasks of Composition and Design

Problems of Composition

Flower and Hayes (1980), who studied the task of writing an essay, noted that successful writing requires simultaneous compliance with a large number of constraints, operating at different levels. One set of constraints requires the selection and organization of ideas from the writer's knowledge into a coherent network of concepts and information for inclusion in the essay. Another involves the linguistic and discourse conventions of written language. A third is rhetorical, involving the need to arrange the essay so as to accomplish the writer's purpose for the intended audience.

Using protocols obtained from subjects working on writing tasks, Hayes and Flower (1980) found three general processes: planning, translating, and reviewing. These three processes allow the writer to attend to a subset of the constraints at any time. In planning, information relevant to the topic is generated from the problem solver's memory, and decisions are made about what to include. In translating, a text is produced using information that has been retrieved, consistent with a writing plan that has been formed. In reviewing, the generated text is evaluated and revised in accord with the constraints of rhetoric, text structure, and such detailed linguistic concerns as correct grammar Hayes and Flower found that writing involves a combination of these processes and postulated that the writing process includes a monitor that determines the sequence of subprocesses, depending on the nature of difficulties that arise.

To write successfully, an individual must

drawing lines
dots to other
the nine-dot
ble interpre-
ed the subjects
itors that were
t who had not
t. This finding
the finding of
Problems") that
rs is more effec-
the results of

omposition

tudied the task of
uccessful writing
ince with a large
iting at different
requires the selec-
; from the writer's
twork of concepts
in the essay.
stic and discourse
guage. A third is
ed to arrange the
e writer's purpose

from subjects work-
s and Flower (1980)
ses: planning, trans-
ese three processes
to a subset of the
lanning, information
generated from the
and decisions are
de In translating, a
mation that has been
a writing plan that
eving, the generated
sed in accord with the
xt structure, and such
ns as correct grammar.
that writing involves a
ocesses and postulated
s includes a monitor
uence of subprocesses,
ire of difficulties that

ly, an individual must

understand the constraints that apply at various levels to the text, must have effective methods for generating or revising text to conform to those constraints, and must actively engage in evaluation in light of the constraints. In studies of young writers, Bereiter and Scardamalia (1982) noted that inattention to constraints, especially global rhetorical concerns, characterizes the writing of many children When they revise a text that have produced, most children attend exclusively to low-level constraints, usually changing only single words or small phrases, rather than attempting to improve more significant general features of their essays. Bereiter and Scardamalia hypothesized that the difficulty lies in the process of evaluating the text, rather than in a failure to understand rhetorical goals or the lack of effective means to produce an improved text. They gave students a set of cue cards with evaluative comments, such as "I need another example here," "The reader won't be convinced by this," "Even I seem to be confused here," and "This is a good sentence." The children's task was to choose a card that seemed appropriate for each sentence in their texts and to make appropriate changes. The technique was effective and consistent with the idea that the children's problem lay in the difficulty of evaluating their texts and applying global constraints, rather than in ignorance of the constraints or methods for complying with them.

Multiple interacting constraints also characterize composition of music, as Reitman (1965) showed in an analysis based on a protocol obtained from a professional composer as he wrote a fugue. Reitman noted that schematic structures that he called transformational formulas played an important role, these included knowledge of the main components of the musical form being composed (exposition, development, and conclusion) as well as subcomponents of those units (exposition → thematic material + countermaterial, thematic material → motive + development, etc.). Reitman found that much problem-solving activity was concerned with constraints. Some constraints were generated by properties of the instrument (piano) chosen for the piece, requiring musical material suited to the instrument. Other constraints were produced by material already included in the piece, such as a requirement that counter-material should be compatible with thematic

material, but sufficiently different to elicit interest. The composer characterized patterns that he developed as conventions, producing melodic, rhythmic, and instrumental properties that were then "used to carry on the movement of the music" (Reitman, 1965, p. 169), with variations introduced to maintain interest.

A substantial knowledge base is required to solve problems of composition, and an important question is how much experience and training a person needs to make substantial creative contributions to a field such as musical composition. Using data from biographies and a standard catalogue of recordings, Hayes (1981) determined the time between a composer's beginning serious musical training and the first composition that had five independent recordings in the catalogue. In almost every case, at least ten years of virtually full-time training occurred before a composer produced a work of sufficient quality to appear commonly in the recorded repertoire.

Recognition and Knowledge of Constraints

In problems that impose constraints, a problem solver must recognize the constraints in order to perform successfully. "Cryptarithmetic Problems" discussed Newell and Simon's (1972) finding that individual differences in cryptarithmetic depended on inclusion in the subjects' problem spaces of significant constraints, such as odd-even parity. Two studies have investigated this factor. one on examination questions by Bloom and Broder (1950) and one on administrative policy by Voss, Greene, Post, and Penner (1983).

In comprehensive college examination questions studied by Bloom and Broder (1950), students were often required to make inferences or deal with information presented in an unusual form. For students who performed poorly, a significant factor was their inattention to constraints in the statements of some questions. For example, when the task was to choose the best explanation for a situation, some students would ignore the relation of alternative answers to the situation and would pick the answers that seemed most nearly true in itself. For such students, the activity of problem solving occurred in a problem space that lacked some of the information that was required for good performance. Bloom and Broder developed an instructional method in which students compared their own problem-solving process, recorded in a thinking-

aloud protocol, with the process of another student whose performance was more successful. This training was effective for many students, teaching them to attend more carefully to constraints in questions as well as such other helpful strategies as increasing their efforts to infer plausible answers from information they could retrieve from memory.

Voss et al. (1983) obtained thinking-aloud protocols on problems involving the design of an administrative policy. For example, problem solvers were asked to develop a policy for improving agricultural productivity in a region of the Soviet Union. Subjects with different amounts of knowledge about Soviet government and history worked on the problem, including students in an introductory course in Soviet politics, experts in political science (some of whom specialized in the Soviet Union and some with other specialties), and experts in another field altogether (chemistry). The solving process of experts was primarily to formulate the problem, and then, after a long initial period devoted to considering historical and political factors, to make successive reformulations based on evaluations of proposed solutions against known constraints. The inexpert student subjects offered problem formulations that failed to include important constraints. Experts in chemistry worked more systematically than the political science students, sometimes using general knowledge about administrative systems to provide useful conjectures, but they too lacked the rich formulations that characterized the problem solver with specialized knowledge.

Design of Procedures

Another type of problem involves tasks in which the materials consist of a set of actions that can be performed, and the problem solver constructs a procedure from these components. These problems are similar to problems of transformation, discussed in "Well-Specified Problems" especially when planning is used to construct a sequence of actions to reach the problem goal.

Hayes-Roth and Hayes-Roth (1978) gave subjects a map of a fictitious town, showing the locations of several stores and other businesses. The subjects were also given a list of errands, such as buying fresh vegetables at the grocery, picking up medicine for a dog at the vet, and seeing a movie. The subjects' task was to plan a schedule that included as many of the errands

as possible. The task presented some general constraints, in particular, a limit on the amount of time available. It also presented local constraints and interactions. For example, it is better to buy groceries late in the day, so they will still be fresh when the shopper returns home; and it is best to go to the movie at one of the times when the feature is starting. Interactions include the proximity of shops, making it more efficient to group together in the sequence errands that involve shops that are near one another.

The Hayes-Roth's simulated performance on their planning task with a model that contained several planning specialists and a blackboard control structure, a design similar to one used earlier in a speech understanding system called Hearsay (Reddy, Erman, Fennell, & Neely, 1973). The specialists are designed to make suggestions about different kinds of planning decisions: They all have access to inferences, suggestions, and other information, which is located in the system's blackboard. This system design supports a feature called opportunistic planning, which has been found in the performance of human problem solvers. Opportunities arise in the form of conditions that make it easy to include an errand, such as the proximity of a store to a place that is already included in the plan, and an appropriate specialist can be activated by that condition.

In the writing of a computer program, the procedure is designed to perform a designated function. Studies of computer programmers and designers have revealed important characteristics of the knowledge required for the solution of these design problems.

Soloway, Ehrlich, Bonar, and Greenspan (1982) gave three problems, typical of elementary programming courses, to students in the first and second introductory courses in programming They identified schematic cognitive structures that they called plans, needed for successful problem solving. The required schemata are quite basic, involving the construction of iterative loops and the use of variables. The schemata provide knowledge of requirements for performing significant program functions, such as the interactions between processing and testing a variable within a loop and between the loop processing and initialization. Students who lacked adequate versions of these schemata made significant errors, for example, by failing

aloud protocol, with the process of another student whose performance was more successful. This training was effective for many students, teaching them to attend more carefully to constraints in questions as well as such other helpful strategies as increasing their efforts to infer plausible answers from information they could retrieve from memory.

Voss et al. (1983) obtained thinking-aloud protocols on problems involving the design of an administrative policy. For example, problem solvers were asked to develop a policy for improving agricultural productivity in a region of the Soviet Union. Subjects with different amounts of knowledge about Soviet government and history worked on the problem, including students in an introductory course in Soviet politics, experts in political science (some of whom specialized in the Soviet Union and some with other specialties), and experts in another field altogether (chemistry). The solving process of experts was primarily to formulate the problem, and then, after a long initial period devoted to considering historical and political factors, to make successive reformulations based on evaluations of proposed solutions against known constraints. The inexpert student subjects offered problem formulations that failed to include important constraints. Experts in chemistry worked more systematically than the political science students, sometimes using general knowledge about administrative systems to provide useful conjectures, but they too lacked the rich formulations that characterized the problem solver with specialized knowledge.

Design of Procedures

Another type of problem involves tasks in which the materials consist of a set of actions that can be performed, and the problem solver constructs a procedure from these components. These problems are similar to problems of transformation, discussed in "Well-Specified Problems" especially when planning is used to construct a sequence of actions to reach the problem goal.

Hayes-Roth and Hayes-Roth (1978) gave subjects a map of a fictitious town, showing the locations of several stores and other businesses. The subjects were also given a list of errands, such as buying fresh vegetables at the grocery, picking up medicine for a dog at the vet, and seeing a movie. The subjects' task was to plan a schedule that included as many of the errands

as possible The task presented some general constraints, in particular, a limit on the amount of time available It also presented local constraints and interactions For example, it is better to buy groceries late in the day, so they will still be fresh when the shopper returns home; and it is best to go to the movie at one of the times when the feature is starting. Interactions include the proximity of shops, making it more efficient to group together in the sequence errands that involve shops that are near one another.

The Hayes-Roth's simulated performance on their planning task with a model that contained several planning specialists and a blackboard control structure, a design similar to one used earlier in a speech understanding system called Hearsay (Reddy, Erman, Fennell, & Neely, 1973). The specialists are designed to make suggestions about different kinds of planning decisions: They all have access to inferences, suggestions, and other information, which is located in the system's blackboard This system design supports a feature called opportunistic planning, which has been found in the performance of human problem solvers Opportunities arise in the form of conditions that make it easy to include an errand, such as the proximity of a store to a place that is already included in the plan, and an appropriate specialist can be activated by that condition.

In the writing of a computer program, the procedure is designed to perform a designated function. Studies of computer programmers and designers have revealed important characteristics of the knowledge required for the solution of these design problems.

Soloway, Ehrlich, Bonar, and Greenspan (1982) gave three problems, typical of elementary programming courses, to students in the first and second introductory courses in programming They identified schematic cognitive structures that they called plans, needed for successful problem solving. The required schemata are quite basic, involving the construction of iterative loops and the use of variables. The schemata provide knowledge of requirements for performing significant program functions, such as the interactions between processing and testing a variable within a loop and between the loop processing and initialization. Students who lacked adequate versions of these schemata made significant errors, for example, by failing

...more general
n the amount
...local con-
example, it is
a day, so they
opper returns
...e at one of
...ing. Inter-
shops making
...er in the
shops that are

performance on
that contained
a blackboard
...lar to one used
ng system called
...ell, & Newly,
signed to make
...nds of planning
...as to inferences,
...ation, which is
...ard. This system
led opportunistic
...ind in the perfor-
...ers. Opportunities
...s that make it easy
the proximity of a
...ready included in
...specialist can be

...uter program, the
...rform a designated
...programmers and
...rtant characteristics
for the solution of
...and Greenspan
typical of elementary
students in the first
...urses in programming.
cognitive structures
eeded for successful
...ed schemata are quite
...struction of iterative
...bles. The schemata
...irements for perform-
...unctions such as the
...cessing and testing a
and between the loop
...zation. Students who
ns of these schemata
for example, by failing

to recognize distinctions between different looping structures Experiments on memory for program texts have shown that experienced programmers can recall more successfully than beginners (Adelson, 1981, McKeithen et al., 1981; see also "Einstellung"). The acquisition of plan schemata as hypothesized by Soloway et al. (1982) provides a natural explanation of this finding.

More advanced problems, involving software design, were studied by Polson, Atwood, Jeffries, and Turner (1981). A task in software design involves planning a complex program, actual writing of the program is performed separately. Polson et al studied the design of a program for compiling an index for a text, given a set of key words to be included in the index. Both professional software designers and students gave solutions with thinking-aloud protocols. The experts recognized functions that had to be included in the solution, such as defining a data structure for the text and searching the key word set for a word that would match each word encountered in the text. Polson et al. concluded that experts' knowledge includes general design schemata that enable decomposition of problems and the progressive forming of more well-defined subproblems, with specific techniques available for some of the subproblems encountered. These schemata provide another example of knowledge for action organized hierarchically like that developed by Sacerdoti (1977, see "Domain Specific Knowledge for Familiar Problems with Specified Goals.")

In the domain of microbiology, two versions of a program that solves problems of experimental design, called Molgen, have been developed. One program by Stefik (1981) designs procedures for modifying the genetic structure of micro organisms. An important issue considered by Stefik is the handling of constraints that arise from interactions between components of a procedure. Molgen designs procedures in a top-down manner, in which abstract plan schemata are gradually made more specific. A method of constraint posting was developed in which requirements for one of the design components could be taken into account in the decisions made about other components.

The second version of Molgen, by Friedland (1979), designs analytic experiments, such as the determination of the sequence of base molecules in a DNA strand or the location of a set of restriction sites on a molecule. In this model schemata called skeletal plans incorporate information about experimental procedures that, through a process of filling in details, develop specific experimental plans based on the specific problem requirements.

INDUCTION

In a problem of induction, some material is presented and the problem solver tries to find a general principle or structure that is consistent with the material. Important examples include (1) scientific induction, including situations in which the material is a set of numerical data and the task is to induce a formula or a molecular structure, (2) language acquisition, where the material is a set of sentences and the task is to induce the rules of grammar for the language, and (3) diagnosis, in which the material is a set of symptoms and the task is to induce the cause of the symptoms. Problems of analogy and extrapolating sequences are inductive tasks that are widely used in intelligence tests. The task of inducing a rule for classifying stimuli into categories has been used in a larger and significant body of experimental study.

An induction problem presents a dual problem space that includes a space of stimuli or data and a space of possible structures, such as rules, principles, or patterns of relations (cf. Simon & Lea, 1974). The task can be conceptualized as a search, within the space of structures, to find a structure that satisfies a criterion of agreement with the stimuli or data. An experimental subject can be tested by being required to use the structure for stimuli that have not yet been shown. When the task is to induce a rule for classifying stimuli, new stimuli may be presented to test whether the subject can classify them correctly. When the task is to induce a pattern in a sequence, the subject may be required to extend the sequence by producing additional elements that fit the same pattern as those that are given.

Solving an induction problem can proceed in two ways, and most tasks use a combination of the methods. The first, a top-down method, involves generating hypotheses about the structure and evaluating them with information about the stimulus instances. The second, a bottom-up method, involves storing information about the

individual stimuli and making judgments about new stimuli on the basis of similarity or analogy to the stored information. Use of the top-down method requires a procedure for generating or selecting hypotheses, a procedure for evaluating hypotheses, and then a way of using the hypothesis generator to modify or replace hypotheses that are found to be incorrect Use of the bottom-up method requires a method of extrapolating from stored information, either by judging the similarity of new stimuli to stimuli stored in memory, or by forming analogical correspondences with stored information.

Induction involves a form of understanding in which a representation is found that provides an integrated structure for diverse stimuli. This general feature also characterizes processes of representing problems such as the textbook physics problems discussed above in "Problem Understanding; Representation." There the space of stimuli is the information in the problem situation—often a problem text or instructions—and the space of structures is a set of possible representations that can be constructed. To be successful, a problem representation must provide the information needed to achieve the problem goal. Thus, in representing transformation problems, the inductive search is constrained by the requirements of problem-solving operators that are available. In some problems of induction, such constraints are not present, and one does not have to do anything with the pattern that is found in the information. However, in some inductive problems, such as medical diagnosis, there are strong constraints related to available operators. The goal is to restore the ailing person to proper functioning, and the effort to induce a cause serves the goal of determining an effective remedy.

In some task domains, the possible structures are represented explicitly as formulae. Examples include induction of quantitative formulas from numerical data in physics, or induction of the molecular structure of a chemical compound. Patterns induced in letter-sequence problems also consist of explicit formula-like rules. These tasks share important properties with problems of design and arrangement (discussed above in "Problems of Design and Arrangement"). The goals of these induction tasks can be considered as the design of a formula that agrees with the data. The solution of design problems generally requires use of strong constraints to limit the space of possibilities for search and this important property is also found in tasks that involve induction of formulas.

The discussion of inductive problem solving will cover. (1) induction of categorical concepts. (2) induction of more complex concepts involving sequential stimuli, (3) induction of relational structure, and (4) diagnostic problem solving.

Categorical Concepts

Of the various inductive tasks that have been studied, by far the most attention has been given to the induction of categorical concepts. This is partly in recognition of their practical importance. Our human capability of organizing experience using conceptual categories undoubtedly contributes much to making our cognitive lives manageable.

In an experiment on concept induction, the experimenter constructs a set of stimuli (e.g., diagrams with figures that vary in shape, size, color, and other attributes) and decides on a rule to classify the stimuli (e.g., "the red circles are positive, all other stimuli are negative"). The subjects are given information about several individual stimuli—that is, they are told whether each stimulus is positive or negative. The subject's task is to induce the rule of classification. Usually the experimenter tests whether the subjects have induced the concept by presenting new stimuli to determine whether they can classify them correctly.

In an early discussion, Woodworth (1938) distinguished between processes of concept induction involving bottom-up and top-down methods. In a bottom-up process, knowledge of the concept is analogous to a composite photograph, consisting of an impression summed over the various stimuli in the category, with the common features emphasized and the variable characteristics 'washed out.' In a top-down process, the problem-solver actively constructs hypotheses about features that define the concept and tests these hypotheses with additional information about examples.

The following discussion deals first with two studies of top-down processes, and then with studies of bottom-up processes of inducing concepts.

Multifeature Concepts
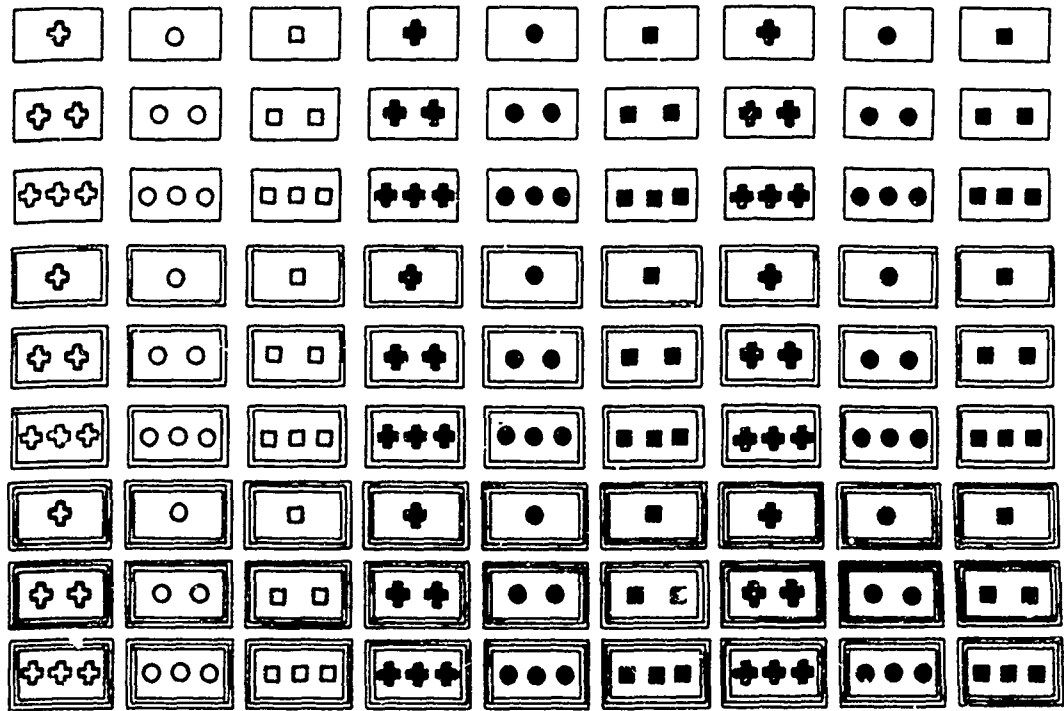
When two or more stimulus features are

Figure 9.11 An array of instances comprising combinations of four attributes, each exhibiting three values. Open figures are in green, striped figures in red, solid figures in black. From *A Study of Thinking* (p. 42) by J.S. Bruner, J J Goodnow, and G.A. Austin, 1956. New York. Wiley. Copyright 1956 by Jerome Bruner. Reprinted by permission.

combined to form a categorical concept, they are combined in some logical formula, such as 'A and B,' or 'If A, then B.' A stimulus is a positive example of the concept if the formula truly describes the stimulus. In the set of stimuli shown in Figure 9.11 the concept 'Green and Circle' specifies the stimuli in column 2; the concept 'Green or Circle' specifies the stimuli in columns 1, 2, 3, 5, and 8.

Consider the requirements for performance of this task, assuming that it is done in a top-down, hypothesis-testing manner. First, the stimulus features must be discriminated, the problem solver must have processes for *recognition* of the features that are used to define concepts. Second, there must be a process for *hypothesis formation*, which constructs candidate hypotheses to be considered. Third, a process of *hypothesis evaluation* is needed to test the hypotheses that have been formed. Fourth, a process for *hypothesis modification* is required in order to use the results of the tests to elimi-

nate incorrect hypothesis, to change existing hypotheses, or to form new ones.

In a landmark study of multifeature concept induction, Bruner, Goodnow, and Austin (1956) observed subjects who, as they worked on concept induction problems, made oral reports about their hypotheses. In certain of these experiments subjects were instructed that concepts were conjunctions of features, and that their task was to induce how many features were relevant and what the features were. Two experiments are considered here.

In one experiment subjects were required to solve two problems with the array shown in Figure 9.11 and a third problem of the same kind from memory—that is, with the stimuli not available. Each of the problems began with the experimenter providing a positive instance—a stimulus that was a member of the concept category. The subject could then choose any stimulus in the display and ask whether it was a positive or negative instance of the concept. The

subject could offer a hypothesis after the choice of stimulus, but this was not required. The subject continued choosing stimuli and receiving information until the correct concept was induced.

The results obtained by Bruner et al. (1956) included characterizations of a variety of strategies used by subjects in selecting stimuli. Strategies of one kind, called focusing strategies, involved finding a positive instance of the concept, then determining which of its features were relevant. For example, suppose the concept was 'Red and Circle.' The subject might be told that the stimulus with three red circles and two borders was a positive instance. The subject could then choose a stimulus that differed from the focal stimulus in the number of circles, say, two red circles with two borders. This would be a positive instance, and the subject would infer that the number of figures was not a relevant attribute. The subject might then vary the color of the figures, choosing the stimulus with three green circles and two borders. This would be a negative instance, and the subject would infer that the color of the figures was relevant, that is, that 'Red' was part of the definition of the concept. With further choices and information, the concept's definition would be inferred.

Other strategies called scanning strategies, involve consideration of specific hypotheses and the use of information to narrow down the set of possible hypotheses. For example, a subject might consider as distinct possibilities the hypotheses 'three figures,' 'red,' 'three and red,' 'circle,' 'three circles,' and 'red circles.' Then finding that a stimulus with two red circles and two borders is a positive instance, all the hypotheses with the property 'three' could be eliminated. Use of a scanning strategy places severe demands on memory. It is impossible to consider all the possible hypotheses simultaneously (there are 255 of them), but it is desirable to consider as many as one can, since information can be used to evaluate hypotheses only in the sample being considered.

The focusing strategies and the scanning strategies differ primarily in the processes they use to form hypotheses. In the focusing strategies, information about instances is used to constrain hypothesis formation. Tests are performed to see whether an attribute is relevant, and when the attribute is eliminated, no hypothesis using it will be formed. If the focusing strategy is used successfully, all but the correct attributes can be eliminated, and the correct hypothesis can be formed directly In the scanning strategies, less use is made of problem information in forming hypotheses, and hypotheses that are in the sample are tested directly with information about instances. Information is used somewhat more directly in evaluating hypotheses in the scanning strategies, but there is consequently a greater need to keep in memory a large set of hypotheses.

Bruner et al. (1956) used 12 subjects whose performance was used to classify them as either focusers or scanners. Seven subjects were classified as focusers and the rest were treated as scanners. The focusing strategy was advantageous for the subjects who used it. They required about half as many choices as the scanners to solve a problem with the stimulus array present (medians of 5 and 10 choices, respectively). In addition, the scanners had noticeably greater difficulty in solving a problem 'in their heads' than they did when the stimuli were present (median of 13 choices), except for one scanner who discovered the focusing strategy while working on the third problem. The focusers' performance without stimuli present did not differ from their performance on the second problem with stimuli present.

Bruner et al. (1956) conducted two experiments to investigate situational factors that influenced subjects' choices of strategies. One experiment compared the effect of an orderly arrangement of stimuli with the same stimuli presented haphazardly. The stimuli used abstract forms, differing on six dimensions, with two values on each dimension. With the 64 stimuli arranged systematically, similar to the arrangement in Figure 9.11, almost all subjects used focusing strategies. When stimuli were not arranged systematically, subjects typically used scanning strategies. There was also a tendency to use scanning strategies when concrete stimuli were used, such as drawings of persons who varied in sex, size, and clothing.

Analyses by Hunt (1962) and Hunt et al. (1966) provided a hypothesis on how to represent categorical concepts in cognitive structure. Hunt proposed that the knowledge of a categorical concept is a cognitive procedure for deciding whether a stimulus is a member of the category. The form of the procedure that Hunt investigated was a decision network, a structure

ll but the correct
and the correc-
tly. In the scan-
nade of problem
·ses and hypoth-
e tested directly
ices. Information
tly in evaluating
strategies, but
r need to keep in
ses.

2 subjects whose
fy them as either
)jects were classi-
were treated as
was advantageous
t. They required
s the scanners to
ilus array present
. respectively). In
oticably greater
n in their heads'
iull were present
t for one scanner
g strategy while
:m. The focusers'
present did not
on the second

icted two experi-
onal factors that
of strategies. One
ect of an orderly
the same stimuli
.muli used abstract
·nsions, with two
ith the 64 stimuli
lar to the arrange-
al! subjects used
stimuli were not
ects typically used
a also a tendency
en concrete stimuli
s of persons who
ing.

and Hunt et al.
in how to represent
gnitive structure
)wledge of a cate-
ive procedure for
is a member of the
ocedure that Hunt
etwork, a structure

of perceptual tests organized in a way that reflected the logical structure of the concept. (This same form was used by Feigenbaum, 1963, for the Elementary Perceiver and Memorizer, used in simulations of rote verbal memorizing. Examples of such decision networks, for recognizing some concepts in geometry problems, were shown in Figures 9 5 and 9 6.) Experiments conducted by Trabasso, Rollins, and Schaughnessy (1971) provided evidence that supports Hunt's characterization Trabasso et al. measured latencies for categorical decisions about stimuli and obtained results that agreed with Hunt's model· Longer times were required for decisions in which the model specified a larger number of perceptual tests A model that simulates acquisition of conjunctive concepts was developed by Williams (1971) who used Hunt's representational hypothesis together with assumptions about limited short-term memory capacity and changes in the salience of dimensions.

An important aspect of the acquisition of complex concepts is induction of the logical relation between the stimulus features in the definition This has been studied by Bourne and his associates in experiments in which subjects are informed of the features that the rules include. For example, a subject may be told that the rule includes 'Red' and 'Circle,' but the subject would then have to discover from examples whether the combination is conjunctive, disjunctive, conditional, or biconditional. When subjects are not experienced in this rule-learning task, there are substantial differences in the difficulty of inducing the various kinds of rules, and these correspond to differences among the types of rules found in standard concept induction tasks (Haygood & Bourne, 1965).

One possible explanation for differences in difficulty is that the rules differ in familiarity to the subjects, with conjunction being the most familiar way to combine features. Overuse of conjunction would lead to a bias in the process of forming hypotheses, with the less familiar forms of hypothesis generated later, if at all, and consequent delays in problem solutions. Evidence in support of this interpretation was obtained by Bourne (1970), who found that differences among the rule forms decreased when subjects were given a series of rule-induction problems. A more specific hypothesis, proposed by Bourne (1974), is that, with experience, subjects acquire a strategy for representing information about

stimuli in terms of truth-table values based on the features known to be relevant. For example, if 'Red' and 'Circle' are the features, then a red circle has the value T-T (true on both attributes), a green circle has the value F-T, and so on. This is an efficient representation for solving concept-induction problems, because each of the alternative rule forms corresponds to a distinctive subset of truth-table values. A conjunctive rule is satisfied only by T-T, a disjunctive rule is satisfied by T-F, F-T, and T-T, a conditional rule is satisfied by T-T, F-T, and F-F, and a biconditional rule is satisfied by T-T and F-F. The truth-table hypothesis is supported by Dodd, Kinsman, Klipp, and Bourne's finding (1971), that training on a task of sorting stimuli into the four categories of the truth table facilitated subsequent performance on rule-induction problems.

Single-Feature Concepts

Induction of conceptual rules may also consist of single features, such as 'all the red pictures,' or, 'the circles.' The task of inducing such a concept is simpler, of course, than inducing a multifeature concept.

EVIDENCE FOR TOP-DOWN INDUCTION

Single-feature concept induction has been studied extensively by H.H. and T.S. Kendler and their associates. One question addressed in their experiments is whether concepts are acquired in the form of a verbalized rule or in the form of an aggregation of individual stimulus–response connections. It is likely that a verbalized rule would result from a top-down hypothesis-testing process of induction, and an aggregation of stimulus-response connections from a bottom-up process.

Evidence has been obtained in experiments in which the conceptual category is changed without informing the subject. A subject is given an initial concept-induction problem involving a single stimulus feature (e.g., "respond positively to red stimuli"). After the subject meets a criterion of correct responses, the rule is changed, either by changing the positive value of the same attribute (e.g., from red to green), called a reversal shift, or by changing to a different attribute (e.g., from red color to large size), a nonreversal shift. It was found that both adult human subjects, and kindergarten children who solved the initial problem quickly, adjusted

more easily to the reversal than to the nonrever-
sal shift (Buss. 1953. Kendler & D'Amato. 1955.
Kendler & Kendler, 1959), whereas rats and
slower-learning kindergarten children adjusted
more quickly to the nonreversal shift (Kelleher.
1956) An interpretation is that adults and
school aged children use a hypothesis such as 'it
depends on color,' which does not have to be
changed to adjust to the reversal shift, while
nonhuman subjects and preschool children
learn specific stimulus–response associations,
for which the reversal shift requires a greater
change In a later study, Erickson (1971) found
that college students adjusted more rapidly to
nonreversal shifts if they had been carefully
instructed about the nature of the concept
induction task, suggesting that when subjects
have more complete information about the task
they tend to remove stimulus attributes from
consideration when their hypotheses are not
confirmed.

Further evidence that adult human perfor-
mance in concept induction is based on definite
hypotheses was obtained by Levine (1963) who
showed that on a series of test trials with no
feedback given, nearly all the sequences of
responses given by college students were con-
sistent with a systematic hypothesis about the
conceptual rule.

PROCESSES OF SAMPLING HYPOTHESES

The processes of forming and evaluating hypothe-
ses in single-feature concept induction are quite
straightforward. Any stimulus feature that is
noticed can be the basis of a rule, and a rule that
links a feature with a response is confirmed
or refuted directly by information about the
category of any example Because the hypotheses
are simple, and many hypotheses are possible, it
is efficient for subjects to consider samples of
hypotheses rather than one hypothesis at a time.
When a sample of hypotheses is considered, the
subject can on each trial eliminate hypotheses
that are inconsistent with the information
given about that trial's stimulus. If the sample
includes the correct hypothesis, the process of
elimination can narrow the sample down to that
hypothesis, which solves the problem. If the
sample does not include the correct hypothesis,
all the hypotheses in the sample will eventually
be eliminated and the subject will have to
generate another sample. Note that this method
is similar to the strategies that Bruner et al.

(1956) called scanning Like the scanning strate-
gies. the strategy of testing samples of hypotheses
is demanding on memory

Proposals about the processes of choosing
hypotheses to be considered. the eliminating of
hypotheses on the basis of stimulus information,
and the recall of previously eliminated hypotheses
have been discussed in theoretical papers by
Gregg and Simon (1967) and by Millward and
Wickens (1974).

Wickens and Millward (1971) provided support
for the assumption that experienced subjects
remember stimulus attributes after eliminating
them. According to their model, if the sample
of hypotheses is exhausted, the attributes of
eliminated hypotheses may still be stored in
memory. Limitations of memory apply to both
the size of the sample that can be considered
and the number of previously eliminated attri-
butes that can be remembered. In Wickens and
Millward's experiment, subjects received exten-
sive training in concept induction, solving many
problems with the same set of stimuli, with
different attributes used to define the concept in
the successive problems. Performance improved
sharply after the first problem or two, and
stabilized within 10 to 20 problems. The model of
attribute elimination was supported by statistical
data as well as by the subjects' responses to
a retrospective questionnaire. Differences in
performance among the individual subjects can
be explained by assuming that they all performed
in accord with the model's assumptions, but that
they differed in the size of hypothesis sample
that they considered and in their capacity to
remember previously eliminated hypotheses.

When performance of inexperienced subjects
has been analyzed using stochastic models, the
results have revealed a problem-solving process
of suprisingly simple structure. Restle (1962)
investigated the mathematical properties of a
process in which a subject considered a sample
of hypotheses and on each trial chose a response
based on one of the hypotheses. In Restle's
model it is assumed that the way subjects
process information differs, depending on
whether the response on a trial happens to be
correct. After each correct response, hypotheses
that are inconsistent with the information
about that trial's stimulus are eliminated from
the sample. After an error, the subject considers
a new sample of hypotheses. A simple stochastic
process results if it is assumed that sampling

occurs with replacement If this assumption is correct. solution of the problem is an all or none event. the probability of solving the problem with no more errors after a new sample is taken is a constant. independent of the number of trials or errors that have occurred previously This implication is counterintuitive If we assume that the subject is sampling and testing hypotheses. the assumption of sampling with replacement says that there is no accumulation of information over trials that makes sampling of the correct hypothesis more likely The all-or-none property is also incompatible with almost any assumption that learned stimulus–response associations are strengthened gradually over trials. or that there is a summative or 'composite photograph' process as Woodworth (1938) proposed.

The counterintuitive all-or-none property of Restle's model received strong empirical support in experiments by Bower and Trabasso (1964). Their experiments with college students as subjects included conditions in which the categorical rule was changed before the subject solved the problem. using either a reversal or a nonreversal shift The assumption of resampling with replacement after errors predicts that shifts prior to solution should not delay the solution of the problem, and this surprising result was obtained.

Computer simulation models of the concept-induction task. using different hypothesis-generating strategies, have been proposed by Gregg and Simon (1967). They showed that when these process models are aggregated (approximately) into simple stochastic models like Restle's (1962), they provide an information-processing explanation for the simple statistical regularities implied by the stochastic models and found in Bower and Trabasso's (1964) data. Gregg and Simon found that a range of different models is required to account for the set of experiments reported by Bower and Trabasso. According to these models, the nature of sampling depends primarily on how much information the subjects can retain about the classification of previous instances and about which hypotheses have already been refuted by the evidence. In general. the process models that fitted the data best were those that implied severe restrictions on short-term memory for previous instances and their classification. Given this retriction on memory, the models are consistent with the all-

or none property —that is. the expected number of trials to solve the problem is independent of the time the subject has already spent on it

Bottom-Up Induction of Concepts

In addition to inducing categorical concepts in a top-down, hypothesis-based manner. induction also can be a bottom-up process. involving gradual emergence of the concept from the features of individual stimuli. This idea has received less attention in psychological research. but it has not been totally missing from the discussion.

Hull (1920) conducted a study of learning in which the materials were pseudo-Chinese ideograms paired with nonsense syllables. The stimuli paired with the same response syllable from list to list all shared a stimulus component. a radical that was part of each of the stimuli. Hull's subjects showed positive transfer on the later lists in the experiment, indicating that they had induced the concepts to some extent. However, most of them were not aware of the feature or features that were shared, indicating that they were not actively testing hypotheses about the categorical rules. It seems likely that the subjects stored information about the individual stimulus–response pairs and gradually built up impressions that included the shared components.

A result similar to Hull's was obtained by Reber (1967), who studied induction of rules for an artificial language. Reber constructed sequences of letters using a set of grammatical rules: for example, "Start with a T or a V," or "After an initial T, use a P or another T," or "After a V that is not at the beginning. use a P or end the sequence." The sequences, from six to eight letters long, were used in a learning task in which subjects were shown the sequence and had to recall them. Subjects working on the grammatical sequences learned faster than subjects who worked on a comparable set of random letter sequences. After learning a set of grammatical sequences, subjects were able to discriminate, with greater than 75 percent accuracy, between new grammatical sequences and sequences that violated the grammar. Even so, subjects were not aware of the rules that were used to form the grammatical sequences. and showed little awareness of their shared features.

Rosch (1978) recently argued persuasively that little of our conceptual knowledge is organized on the basis of definite feature structures, like those used in most experiments on induction of categorical rules. First, Rosch, Mervis, Gray, Johnson, and Boyes-Braem (1976) proposed, with empirical support, that concepts at different levels of generality are not equally salient, but that there are basic categories whose members share features that are not shared by members of other categories, including characteristic patterns by which we interact with them physically. For example, chair, table, and hammer refer to basic categories, while their superordinates, furniture and tool, and their subordinates, such as picnic table and claw hammer, are less fundamental. Data supporting this distinction were obtained by Rosch et al., whose subjects were given a series of 90 terms and were asked to write all the attributes that came to mind. Another group of subjects was given the same terms and were asked to write descriptions of muscle movements that they would make in interacting with the objects. Many more attributes and movements were associated with the basic terms than with their superordinates, and few additional attributes beyond those for the basic terms were given for the subordinate terms.

Rosch (1973, 1975) has also contended that natural concepts are represented as prototypes, rather than as sets of features. A prototype may be thought of as a kind of schema for recognition of members of a category, which is activated more readily by typical representatives than by atypical ones. For example, in the category of birds, robins and canaries are judged to be more typical than penguins or peacocks; in the category of tools, hammers and saws are judged more typical than anvils or scissors. Rosch (1975) found that there is firm agreement among subjects in ratings of typicality. Evidence that typicality influences cognitive processes has been obtained when subjects are asked to judge whether statements such as "A robin is a bird" or "An anvil is a tool" are true. In these experiments, judgments are made more quickly for the statements involving more typical examples (Rosch, 1973; Rips, Shoben, & Smith, 1973).

Acquisition of prototypical concepts has been studied experimentally (Posner, Goldsmith, & Welton, 1967; Franks & Bransford, 1971;

Reed (1972), and others) For these experiments, a set of stimuli is constructed by varying a single stimulus, the prototype The stimuli, which may be geometric forms, patterns of dots, or schematic faces, are shown to subjects, after which a recognition test is given. Subjects' confidence in recognition is a function of the similarity of stimuli to the prototype. When the prototype itself is shown, subjects respond positively and confidently, even if the prototype was not included in the set of stimuli they saw. Several investigators have shown that this performance can be explained by considering the frequencies with which various stimulus features occur during the learning trials; for example, the features of the prototype appear with great frequency, even if the prototype itself is not presented (Reitman & Bower, 1973; Neumann, 1974).

A model that simulates bottom-up acquisition of a prototypical concept has been formulated by Anderson, Kline, and Beasley (1979), using general principles of learning in the context of a production-system model of performance. The Anderson et al. system stores cognitive representations of the patterns seen in individual stimuli, and additional representations are stored by processes of generalization and discrimination. Representations are strengthened when they provide a basis for recognizing stimuli that are presented. The Anderson et al. simulation accurately mimics subjects' performance on recognition tests, including false recognition of prototypes that have not been presented during learning.

A reasonable expectation is that many learning processes are not strictly top-down or bottom-up, but a combination of the two. Such combinations were analyzed by Greeno and Scandura (1966) and by Polson (1972) in studies of concept induction involving verbal items. In an experimental setup like that used by Hull (1920), lists of paired associates were presented to be memorized, and in successive lists the same response term was paired with different but interrelated stimuli. Greeno and Scandura found that transfer to individual items occurred in an all-or-none manner; different sets of items had differing proportions of items with no errors, but for items with any errors performances in the transfer conditions could not be distinguished from each other or from performance on control items. The finding of all-or-none transfer suggests

a top-down conceptual process in which any individual item either is or is not recognized as a member of a definite category Polson (1972) studied acquisition of the conceptual categories and found that it was not an all-or-none process. His findings suggest a two-stage process. For some subjects, there is an initial stage of bottom-up learning, in which associations of responses with patterns of features are stored, with transfer depending on features that are shared by similar items. In the initial phase, the subject may notice by chance the shared features of members of a concept category. Once the shared feature of a category is recognized, the second stage of learning occurs, involving an active, top-down process in which the subject searches actively for features to use in classifying the stimuli.

It is likely that both the top-down and the bottom-up methods of learning about categories are available to human learners, and the question arises as to what circumstances make it more likely for one rather than the other to occur. Brooks (1978) compared a condition in which subjects were asked to learn names for individual stimuli with one in which subjects induced a rule for classifying stimuli. Explicit rule induction led to better knowledge of relevant features, reflected in better performance on classification of new stimuli, as would be expected from learning by top-down induction. Subjects who learned individual names showed superior performance in recognition of specific stimuli from the learning set, but also recognized new stimuli at an above-chance level, as would be expected from bottom-up acquisition of a concept involving a summation of instances.

Sequential Concepts

We now turn to two more complex tasks involving induction of concepts, in which the materials are sequences of elements organized in patterns, and the subject's task is to induce the patterns. In the first task which concerns extrapolating sequences of letters, the subject's task is to identify patterns in the sequences presented and to use the patterns to extend the sequences. The second task concerns induction of grammatical rules of a language from example sentences that are consistent with the grammar.

In these tasks the problem space includes a set of stimuli and a space of possible structures, as in all induction problems. However, compared to the space of possible rules for classifying stimuli, the spaces of possible pattern descriptions for sequences and of possible grammatical rules are extremely large. To solve these problems, substantial reductions of the search spaces are required. These reductions are accomplished by constraints on the generation of hypotheses. In sequence extrapolation, a limited set of relations and sequence forms are considered; in grammar induction, hypotheses about the structures of sentences are constrained by the structures of situations that the sentences describe.

Sequence Extrapolation

An example of a sequence extrapolation problem follows: $mabmbcmcdm \ldots$, where the task is to extend the sequence. In a model of sequence extrapolation formulated by Simon and Kotovsky (1963), a pattern is induced from basic relations between the letters in the problem string. The pattern is a kind of formula for producing the sequence; once discovered, the formula can be used to extend the sequence, as required.

For example, for the problem $mabmbcmcdm \ldots$, the formula that is induced is the following: $[s_1 : m; s_2 : a], [s_1, s_2, (N(s_2)), s_2]$. The first part of the formula is initialization. There are two subsequences, denoted s_1 and s_2. s_1 starts with m, and s_2 starts with a. The second part of the formula gives instructions for producing the sequence. The instructions are interpreted as follows s_1—write the current symbol of s_1; s_2—write the current symbol of s_2; $(N(s_2))$—change the symbol in s_2 to the successor (N for next) of the current symbol; finally, s_2—write the (new) current symbol of s_2. The entire sequence is generated by repeating this routine as many times as necessary.

The problem solver constructs a formula as a hypothesis, based on the first letters of the given sequence, and tests the hypothesis with more letters. Since there are many different ways to form a sequence of letters, the number of possible formulas is, in principle, extremely large. To make the task manageable, some constraints have to be imposed. In Simon and Kotovsky's (1963) model, constraints are imposed on the generation of hypotheses. As in the focusing strategies about the structure of a pattern are based on features of the stimulus, rather than being generated a priori. Furthermore, only a few of the possible hypotheses are ever generated.

because the model considers only a small set of relations between elements and it is assumed that the sequence fits a specific form.

The model knows the alphabet of letters, both forward and backward. The relations that are recognized are identity, I, and successor N. The problem solver assumes that the sequence is periodic, an important structural characteristic.

The model begins by determining the period of the sequence. Periodicity can be discovered either by noting that a relation is repeated every nth symbol, or noting that a relation is interrupted at every nth position. In the problem *mabmbcmcdm* . . . the periodicity is identified by noting that the relation I occurs at every third symbol. Then the problem solver produces a description of the symbols that occur within the periods and relations between corresponding symbols in successive periods. For *mabmbcmcdm* ... the description requires two subsequences, one of which is just repetition of m; the other starts with a and moves incrementally to produce the final term in the set of three symbols. The result of the process is a formula for producing the sequence, such as the one described earlier for the example problem.

Because the product of the inductive process is an explicit formula, sequence extrapolation can be considered as a problem of design as well as of induction. Viewed in this way, the problem solver has available a set of symbols—s_1, s_2, s_3, (perhaps more), N, and the letters of the alphabet—and has the task of constructing from these symbols. The feature of sequence extrapolation that makes it an inductive task is the criterion that the construction must satisfy, the criterion that the formula should produce the sequence of letters that is given in the problem. In ordinary problems of design, such as anagram or cryptarithmetic, the criterion is a general property rather than agreement with an arrangement of stimuli.

Simon and Kotovsky (1963) reported data on the difficulty of solving 15 different sequence-extrapolation problems by two groups of subjects and found that the solvers agreed fairly well with their program on the relative difficulty of the 15 problems. In a more thorough empirical study, Kotovsky and Simon (1973) collected thinking-aloud protocols on problems with sequences so presented that the subjects had to lift a panel to see the individual letters. The data were consistent with the model in important

respects Like the model. the subjects determined the periodicity of sequences and looked for relations between successive elements or between elements separated by a regular period. Representations of sequences induced by the subjects agreed with those induced by the model in a majority of cases.

There were also discrepancies, some of which involved relatively minor details of programming, but two of which revealed significant processes in humans not represented in the model. First, the subjects' performance showed closer integration than did the program between discovery of the period of the sequence and induction of the pattern description. These are distinct phases in the model, whereas the human problem solvers used information in forming the pattern description that they had picked up during the phase of finding the period. Another discrepancy between human data and Simon and Kotovsky's simplest model was that in some problems, human solvers induced patterns with hierarchical structure, involving a single low-level description and a higher-level switch that transited between versions of the low-level structure. A hierarchical relation between levels of pattern description is a basic structural feature of sequential patterns that can play a dominant role in the induction process, as Restle (1970) has shown.

Grammatical Rules

In considering the induction of the grammar of a language we limit the discussion to those aspects of language acquisition that relate directly to general issues in the theory of problem solving.

In acquiring the grammar of a language, learners are presented materials that include sentences in the language. Their task is to infer a set of rules that can be used to parse sentences they hear and to produce sentences that are grammatical in the language. Thus, the problem solving involves a search in a space of possible syntactic rules. The space of stimuli includes the grammatical sentences that the learners hear and the task is to induce the rules that characterize the structure of those sentences.

Human knowledge of the rules of grammar is *implicit*, in contrast to the explicit formulas that are induced in the sequence extrapolation task. This can be seen in the fact that very young children have a significant knowledge of

grammar (e g., Brown, 1973), whereas adults know grammatical rules explicitly only if they have had special training. Because of the implicit nature of grammatical knowledge. the product of language learning is characterized as a set of *procedures*, rather than explicit formulas or other descriptions of structure. The procedures acquired by learners of a language enable them to produce and understand sentences that are in accord with the grammar of the language, and to distinguish between grammatical and ungrammatical sequences of words. Such a set of procedures is referred to as knowledge of the grammatical rules, because the rules are built into the procedures. As with much procedural knowledge, an individual's knowledge of the rules in the form of procedures does not imply the ability to state what the rules are.

There is evidence from both empirical studies (e.g., Moeser & Bregman, 1972) and theoretical analyses (e.g., Wexler & Cullicover, 1980) that grammatical rules are learned more easily if reference is provided for terms in the language. This indicates that, in the space of stimuli for inducing a grammar, each sentence is paired with a situation that the sentence describes. The functions of situations in facilitating induction of grammatical rules probably include assisting in determining which words belong together in constituent units (cf. Morgan & Newport, 1981).

An analysis of language acquisition by Anderson (1975, 1977) serves as an example of a definite information-processing mechanism for acquiring knowledge of grammatical rules in the form of procedures. Anderson's system includes learning processes that show how semantic reference can facilitate the acquisition of grammar. His learning program called LAS for *Language Acquisition System*, induces rules of grammar when it is given sentences in a language accompanied by the semantic objects to which the sentences refer. For example, if the sentence, "The red square is above the small circle," was presented to LAS, there would also be a semantic network that represented an object with the properties *red* and *square*, and another with the properties *small* and *circle*, and the relation *above* between the two objects.

LAS has a procedure, used in its learning of grammar, that identifies the objects in the semantic network that correspond to words in the sentence and forms a structure showing the relations among those concepts. This structure

is used to determine constituent units of the sentence In the example sentence, the words *red* and *square* are bracketed together, because they are properties of the same object, as are *small* and *circle* The relational term *above* is at a higher level in the bracketing formed by LAS. The procedures that LAS acquires include rules for parsing noun phrases (NP) such as *the red square* and *the small circle*, and sentences of the form NP Relation NP There is also a mechanism for generalization in LAS that makes it eventually parse similar structures with a single rule, and some of these generalizations produce recursive parsing rules. The generalization process sometimes produces incorrect rules that are too general, and LAS also includes a discrimination mechanism that restricts the application of its language-processing procedures.

Viewed as a problem-solving system, LAS conducts search in a space of procedures for producing and understanding sentences. Note that LAS can also be viewed as designing or constructing these procedures. The system's use of the structure of situations provides significant constraints that are needed for the search. As in Simon and Kotovsky's (1963) model of sequence extrapolation, the constraints are applied to the generation of hypotheses. Processes for modifying the induced procedures are available; LAS can generalize its procedures, which makes its performance more efficient, and it can add restrictions to the application of procedures when it is informed that the use of a procedure has produced an error.

Nonsequential Patterns

Our discussion of induction of patterns that are not sequential in character begins with a simple case; an analogy problem in which one or two pairs of items are presented that are related in some way. The task is to form another pair with the same relation. The solving of simple analogy problems has been analyzed both empirically and theoretically. For more complicated problems, involving induction of concepts in mathematics and of quantitative regularities and structures in scientific domains, the available analyses are primarily theoretical.

Analogy Problems
The form of an analogy problem is A : B :: C : D, where D is often a set of alternative items that

can complete the analogy, with the subject required to choose one from the set. A and B are related in some way, and the correct choice is a D item with the same relation to C as B has to A Solution of an analogy problem involves search in a space of relations for a relation that can be applied to both the A . B and the C . D pairs, or to one of the C Di alternatives more successfully than any of the others. Analogy problems are commonly used in tests of intellectual ability. In factor-analytic studies, analogy problems contribute most to the factor of induction, the single best predictor of academic achievement (Snow, 1980).

Solutions of analogy problems requires (1) a process for recognizing or analyzing relations between pairs of stimuli—that is, between the A and B stimuli and between C and each of the Di alternatives, and (2) a process that compares relations found for the A . B pair with relations found for the various C . Di alternatives and chooses the C . Di relation that best matches an A . B relation. In the simplest case, the relation for A and B that comes to mind first also applies to one and only one of the C . Di pairs. When this does not occur—because the relations found for A . B apply either to more than one C . Di pair or to none of them—some further analysis of the A . B pair is required. In such cases, other A . B relations may be suggested by relations that are found in considering the C . Di pairs.

Two processes for solving analogy problems have been described. In one, relations between pairs of items are based on information stored in the problem solver's memory. Memory-based analogy problems include most verbal analogies, where solutions use relations between words that are stored in memory or are inferred from word meanings. In the other process, relations are determined by analysis of features of stimuli. For example, in analogy problems composed of geometric diagrams, the relations between pairs of terms are found by comparing pairs of diagrams and identifying differences between the members of each pair.

Relations Based on Semantic Memory

Solutions to many verbal analogies are based on their meanings stored in semantic memory. Reitman (1965) formulated a model for verbal analogies based on the activation of concepts in a semantic network. Reitman's model, called Argus, solves problems such as bear . pig .

chair (foot. table. coffee strawberry) Argus has knowledge of words in a network of relational connections, for example bear and pig are both connected to animal through the relation *superordinate* Activation and inhibition are transmitted through connections between units

Argus can perform according to different strategies. In one strategy, the A and B terms are activated, and relations that become active are noted, then C becomes active, and the Di alternatives are activated in turn A goal is set for relations that are the same as the ones activated by the A B pair When a C Di pair activates those relations, that Di alternative is chosen. In the example, after bear and pig are activated, their superordinate relations to animal become active, because these relations lie on a path between the activated terms Then chair is activated along with the Di alternatives taken in turn, with the goal of finding active superordinate relations. This goal is achieved when table is activated, because both chair and table are connected by superordinate relations to furniture.

Strategic factors in analogy problems were demonstrated in an experiment by Grudin (1980) Grudin presented two kinds of analogy items: standard items, where a salient relation between A and B can be matched with one of the C : Di pairs. and nonstandard items, where there is no salient relation between A and B, but a relation between A and C matches one between B and a Di alternative. An example is the item bird . air .. fish . (breath, water, swim) in standard form, which in its nonstandard version is bird . fish . air . (breathe, water, swim) The nonstandard problems are more difficult, as measured by the time required for a solution However. if subjects can adapt their strategies to look for relations between A C and B Di pairs, the difficulty of nonstandard problems may be reduced Grudin's sequence of problems included five-item sets that were either all standard or all nonstandard. followed by either a standard or a nonstandard problem. During solution of a set of nonstandard items, a shift in strategy could occur, involving more attention to the A . C and B Di pairs. This produced shorter times for nonstandard problems following nonstandard sets than for nonstandard problems following standard sets, and that result was obtained.

Thinking-aloud protocols in solution of verbal

analogies were obtained in a study by Heller (1979) and were also described by Pellegrino and Glaser (1982) Heller first presented the three terms of an analogy stem and asked the subject to think aloud and to include a statement of any A B relations and expectations about the answer that came to mind Then four alternative answers were presented individually and the subject judged whether each alternative was an acceptable answer. and why The complete problem was then presented for a final choice

Heller's findings were consistent with the general features of Reitman's (1965) hypotheses of solution strategies and of finding relations by the activation of a semantic network. In Heller's experiment strategic factors provide an interpretation of individual differences in performance, and the activation hypothesis is supported by the variability in solution sequences.

Heller's major finding was a striking difference between the degree to which groups of subjects adhered to the task constraints of analogy problems The main constraint of an analogy is the requirement that the relation A . B and C Di correspond. If a subject chooses a Di response on the basis of a relation to C without regard to whether that relation corresponds to the A B relation, then the analogy constraint has not been applied Subjects who had good overall performance mentioned the similarity or difference between an A . B relation and at least one of the C Di relations on nearly all the problems In contrast, subjects with poorer overall performance were inconsistent in applying the constraint of matching the A . B and C Di relations and frequently accepted answers based on a vague relation between Di and C, or with other terms in the analogy. To account for the differences among subjects in this adherence to the task constraints, Heller proposed that individuals differ in the strengths of the goals that require different solution strategies In Reitman's model, this would be analogous to the better subjects' having more strongly activated strategic goals or to differ ences in the degree to which other processes interfered with goals.

Heller's protocols also revealed considerable variability in the sequence of steps taken to solve the problems In most cases, subjects identified an A B relation and then thought about C Di alternatives in the context of that relation There were also cases in which a

relation between A and B came to mind as a subject thought about one or more of the C Di relations Such solution sequences occurred in about 20 percent of the problems on which subjects adhered to the analogical constraints. Reitman's assumption that relations are found by activation of a semantic network provides an interpretation of the variability of solution sequences, since activation of a relation in the context of a C . Di pair would facilitate its recognition for A . B in some cases where A B did not elicit it.

Further information relevant to individual differences was obtained in a study by Pellegrino and Glaser (1982). Analogy items with single D alternatives were presented and subjects judged the items as true or false. Pellegrino and Glaser used an experimental and statistical method introduced by Sternberg (1977), in which the four terms are presented in sequence, with the subject making a response to request presentation of the each succeeding term. The latencies of the responses are used to estimate the time for various components of the solution process, according to a general model. Each latency includes time to encode the new item. When B is presented, the latency includes time to infer one or more relations between A and B. When C is presented. the latency includes time to map A . B relations into the C term. When D is presented, C . D relations are inferred and compared with the A . B relations. It was assumed that the comparison process could have three outcomes. The relations could correspond well, leading to a response of *true*. The lack of correspondence could be so great that the subject would immediately reject the analogy and respond *false*. The subject could also judge that the correspondence was indeterminate and requires a more extended analysis. possibly including review of the A and B terms to find new relations.

Pellegrino and Glaser used four sets of items in this study, positive items, which were judged to be appropriate analogies, and negative items, which were judged to be inappropriate. Within each of these sets. there were items in which the C and D terms were closely associated and other items in which they were not associated. A weak C–D association for a positive item, or a strong C–D association for a negative item, was expected to make the item more ambiguous and increase the frequency of extended analyses in the final component of the solution process. The results

supported this expectation. estimates of the proportion of problems with extended analyses were higher for weakly associated than for strongly associated positive items (.55 and 23, respectively), and also higher for strongly associated than for weakly associated negative items (.19 and .07, respectively). A similar correlation of item difficulty with time spent in the final stage of solution was obtained by Barnes and Whitely (1981).

Pellegrino and Glaser's major finding was that the frequencies of an extended analysis were correlated with the subject's overall ability in the analogies task. The subjects were college students divided into two groups on the basis of their scores on a standard analogies test. The estimates of time for the various information-processing components were generally longer for the low than for the high-ability subjects. But the most striking difference was in the frequency of engaging in an extended analysis, which was more than twice as high for the low-than for the high-ability subjects. Pellegrino and Glaser concluded that, since the low-ability subjects often arrived at the final stage of processing an analogy with an inadequate representation of the relations among the other terms, they had to reconsider the A, B, and C terms more frequently than the high-ability subjects. A similar difference in the solution process was found by Snow (1980), for spatial reasoning tasks in which the items are diagrams, and the subjects' reexaminations of terms could be observed by recording eye movements. In verbal analogies, this difference in processing could be due to differences in the information in semantic memory, differences in the activation process, or differences in strategy, with the low-ability subjects more likely to want to see the final term to facilitate recognition of A:B relations. This conclusion is consistent with Heller's finding that students with low ability in analogies often choose responses that violate the constraint of an analogy problem. When they lack a response that satisfies the constraints, they are likely to choose a response on some other basis.

In Reitman's (1965) model of verbal analogy solution, relations are relatively discrete components of semantic memory. This characterization is probably correct for most verbal analogies, though not for all. An example is provided by Rumelhart and Abrahamson (1973), who studied the solution of verbal analogy problems in a single semantic domain. the names of animals

Analogies composed of animal names have two properties that are different from most verbal analogies. first. they depend on more than one relation, and the relations are combined somehow in solving the problem. and second. the relations differ in degree. rather than just being present or absent.

An example that illustrates multiple relations is the following: rabbit·sheep..beaver (tiger, donkey). Donkey seems the better answer. perhaps because while a relation involving size is similar for beaver:tiger and beaver donkey. and both are similar to the size relation for rabbit:sheep, there also is an additional difference for beaver:tiger—tigers are ferocious while beavers are not, and thus the beaver:donkey pair matches the rabbit:sheep pair better, which also lacks a difference in ferocity. The graded nature of relations is illustrated by rabbit:beaver::sheep:(donkey, elephant). Donkey seems the better answer. The judgment seems to depend mainly on the sizes of the animals, and beavers are larger than rabbits, but the difference is not large enough to make sheep:elephant seem appropriate.

It is convenient to use a spatial representation to represent differences of graded magnitudes that can be combined easily. In such a representation, the dimensions of the space correspond to salient ways in which items differ from each other. Each item is located at a point in the space. The coordinates of the point correspond to the values that the item has on each of the dimensions.

A spatial representation of a set of items can be obtained by presenting pairs of the items to subjects and asking them to judge how similar the members of each pair are to each other. These judgments of similarity are used as estimates of the distances between pairs of items, and items are located in the space so that the distances between points are as close as possible to the estimates obtained in the experiment. In the method of choosing the spatial representation, called *multidimensional scaling*, an attempt is made to represent the items in one dimension; if that is unsuccessful two dimensions are used, and so on until a space is found in which the points are located so that interpoint distances agree satisfactorily with the similarity judgments given by subjects.

~~problems in a es of animals il names have nt from most bend on more is are combined n and second. ather than just ultiple relations :: beaver: tiger, better answer. in involving size beaver: donkey size relation for litional difference ferocious with e beaver: donkey eep pair better. if ferocity. The is illustrated by elephant). Donkey judgment seem to f the animals, and s, but the difference ke sheep: elephant atial representation graded magnitudes In such a represen- ie space correspond as differ from each d at a point in the he point correspond has on each of the of a set of items can pairs of the items to to judge how similar r are to each other ilarity are used as es between pairs of ed in the space so that oints are as close as obtained in the experi- choosing the spatial multidimensional scaling. resent the items in one ccessful two dimensions til a space is found in ated so that interpoint only with the similarity jects.~~

Henley (1969) obtained judgments of similarity for pairs of animal names, and obtained a spatial representation with three dimensions: size, ferocity, and a third dimension that probably involves a mixture of attributes, including similarity to humans. These results were used by Rumelhart and Abrahamson (1973) in their study of analogy problem solving. The relation between two items A and B corresponds to the vector that connects the points for A and B in the spatial representation. The vector represents the combination of differences in the three dimensions between the two items: for example, the vector from *beaver* to *tiger* represents a moderate increase in ferocity, a large increase in size, and very little difference in 'humanness.' In Rumelhart and Abrahamson's model, to solve an analogy A . B :: C : (D1, D2, D3, D4), the A → B vector is translated to C, and the probability of choosing each of the Di alternatives is a function of its distance from the ideal point defined by the end of the vector. In one experiment, the model provided accurate predictions of the frequencies of subjects' rankings of the various response alternatives in analogy problems. In another experiment, fictitious animal names were assigned to locations in the spatial represen- tation. These fictitious names were used in analogy problems for which subjects received feedback, and the subjects induced features of the fictitious animals, responding appropriately to new analogies involving their names.

RELATIONS BASED ON FEATURE ANALYSIS

In a geometric analogy problem, the terms are diagrams that differ in various ways. In the example given in Figure 9.12, the best answer is apparently D2. Diagrams A and B are related by deletion of the dot and moving the rectangle from inside the triangle to a position at the left of the triangle. Diagrams C and D2 are related similarly: the dot in C is also deleted, and the Z is moved from inside the segment of the circle to the left of the segment.

As Figure 9 12 illustrates, the relation between two diagrams can involve several aspects, corresponding to components of the diagrams that differ Some of the differences may be quantitative, for example, the amount of rotation of a component or the amount by which the size of a component is increased or decreased. In analogy problems involving animal names, these characteristics of composite and quantitative
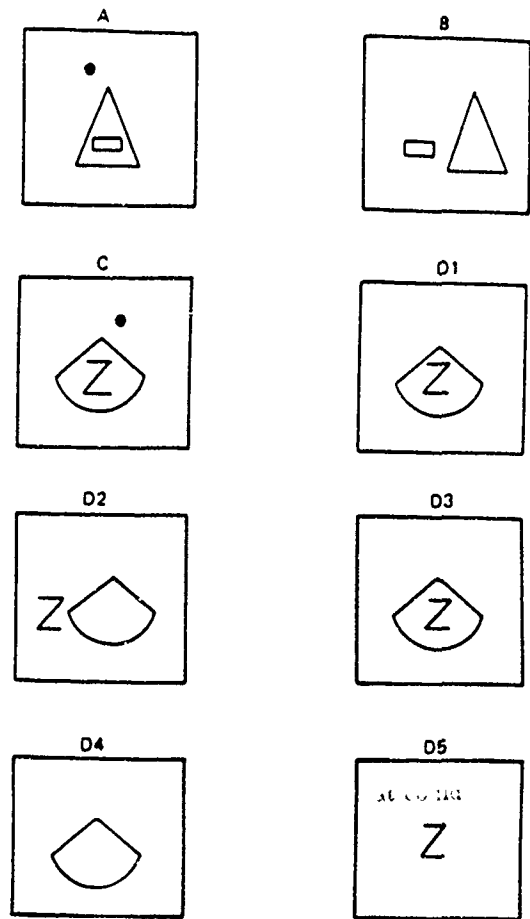

Figure 9.12. A geometric analogy problem. From "A Program for the Solution of Geometric-Analogy Intelligence Test Questions" by T.G. Evans in *Semantic Information Processing* (p. 273), M. Minsky (Ed.), 1968. Cambridge, MA: MIT Press. Copyright 1968 by the MIT Press. Reprinted by permission.

relations make a spatial representation of items a reasonable one. On the other hand, spatial representation is not economical for geometric analogies, because there are too many ways in which diagrams can differ. For animal names, a satisfactory approximation can be reached by characterizing all pairwise relations by dif- ferences on three dimensions, but geometric diagrams do not have so simple a structure.

In geometric analogies, relations are found by examining features of the diagrams, rather than by retrieving information from memory, as with verbal analogies. Therefore, a model for solving geometric analogy problems has two components, one that analyzes diagrams and

identifies relations between them. and another that compares the relation of A B with relations of the C. Di alternatives and chooses the best match.

Evans (1968) developed a model that solves geometric analogy problems. The program is given descriptions of some diagrams that specify the locations of straight lines, curved lines, and closed figures. From these descriptions, relations among components are derived: for example, that one figural component lies inside another, or above it in a diagram.

The model then compares its representations of the diagrams in pairs and forms descriptions of the relations between the members of the pairs. These relations are in the form of transformations—that is, changes in one diagram that would make it the same as the other diagram in the pair. For example, in one diagram a component might be removed or added, or one might be changed in size or rotated, or the relative positions of two components might be changed, for example, by moving one from inside the other to above the other.

The relation between A and B is then compared with the relations between C and each of the Di alternatives, by matching components of A with components of C and determining which of the transformations in the A:B relation also occur in the C:Di transformation. The Di alternative chosen is the one for which the greatest number of transformations can be made to correspond.

Evans (1968) developed his model as a project in artificial intelligence, rather than as a simulation of human problem solving, but the model nevertheless has features that seem plausible as psychological hypotheses. One such feature is a suggestion that problems with more complex diagrams or relations between diagrams should be more difficult for human subjects to solve. In the model, diagrams are more complex if they have more components, and relations are more complex if there are more transformations—that is, if there are more changes in components between related diagrams. These two factors were varied in an experiment by Mulholland, Pellegrino, and Glaser (1980). and both had significant effects. Problems whose diagrams had more components and problems with more transformations both required longer times for solution.

In the human solution of geometric analogy problems. we should expect some of the same characteristics of performance that have been observed in the solution of other analogy problems. In verbal analogy problems when the subject's representation of the A B relation and the C. Di relations are not sufficient to provide a determinate answer. additional further processing is necessary Findings by Sternberg (1977) show that this factor is important in geometric analogy problems as well. Sternberg measured the time to solve problems presented after part of the problem had been shown. enabling part of the processing to occur. He used the differences between conditions as estimates of the times for components of the solution process. In comparing subjects with differing levels of general reasoning ability. Sternberg found a large difference in the time required to process the C:Di alternatives in geometric analogy problems, with much of the difference attributable to a process of comparing alternatives when prior processing had not provided a unique solution.

Inductive Problems in Mathematics and Science

Cognitive analyses have been developed in the form of computer programs that invent new mathematical concepts, based on properties of examples, and that induce formulae and structures from data in scientific domains. Three models are discussed: one that invents new mathematical concepts, one that induces formulae from sets of quantitative data, and one that induces molecular structure from data of mass spectroscopy.

INVENTION OF CONCEPTS IN MATHEMATICS

A program called AM (Lenat, 1982) generates examples of concepts that it knows, and develops new concepts based on properties of the examples. The main domain in which AM was run was elementary mathematics. The AM program was given initial concepts involving sets and developed a variety of concepts involving numbers. For example, AM developed concepts of addition and multiplication. developed the concept of primes. and arrived at a conjecture that every number is the product of a unique combination of prime numbers.

It is useful to compare AM's task to the standard experimental task of concept induction, for example, that of Bruner et al. (1955). In

standard concept induction, a set of examples is provided by the experimenter, with some positive examples and some negative examples determined by a rule, and the subject's task is to induce the rule. Hypotheses are generated by the subject and tested with information about further examples until the correct concept has been found Each hypothesis that is generated is itself a concept, in the sense that it provides a rule for classifying the stimuli. The main problem-solving work is to determine which rule is correct.

The task of AM is not defined as well, in two respects First, the examples are not provided by an experimenter, but rather are produced by AM Second, AM does not have a specified criterion of correctness for the concepts that it generates Instead, AM evaluates its concepts by some criteria of importance, based in part on how easy it is to generate examples.

In AM the knowledge of concepts is organized as a set of *facets*, including some that are standard for semantic networks, such as generalizations, specializations, and examples, and others that are especially useful in mathematics, such as objects that are in the domain or range of a function. Facets also hold procedural information, such as ways to test whether an object is an example of the concept. Reasoning activity in AM is organized as a set of tasks, each involving a concept and one of its facets. Examples of tasks include filling in examples of a concept or forming a generalization or a canonical representation of a concept. Tasks that are proposed are placed on an agenda, and the choice of a task to perform is based on an evaluation of the reasons for the task, including the importance of concepts for which the task would contribute new information. Heuristics that contribute to the developments of new concepts include efforts to form a more general concept if an existing concept has very few examples, and to form new representations that clarify the relations between concepts.

Table 9.10. **Data for a simple induction problem**

Time	Distance
1	0.98
2	3.92
3	8.82
4	15.68
5	24.50

We note that AM does not really do mathematics in the usual sense It has no concept of deductive consequence and thus, does not develop a body of concepts and principles with a formal structure Even so, it provides an example of a system that goes well beyond the knowledge that it is given initially, moving into a conceptual domain that is quite different from that of its initial concepts.

INDUCING QUANTITATIVE REGULARITIES

A system called Bacon induces formulas from numerical data (Langley, 1981, Langley, Bradshaw, & Simon, 1983). The data are values of some variables that are controlled and other variables that are measured, a simple example is in Table 9.10. The goal is to find a formula that describes the relation between the variables, in this case distance and time. The two components of the problem space are the subspace of stimuli (the set of data) and the space of structures (the set of formulas constructable with the variables that are included in the data).

A simpler approach than Bacon's is adequate for relatively simple induction problems. This simpler approach tries to fit alternative formulas that are known in advance. For example, for Table 9.10, a linear function can be tried, and the discrepancy that is noted shows that there is positive acceleration. This suggests trying a quadratic formula, which fits the data. Generate-and-test methods of this kind have been analyzed by Huesmann and Cheng (1973) and by Gerwin (1974), with supporting experimental data.

The task of inducing formulas can become unmanageable for a simple generate-and-test method if there are several variables that can be related in complex ways. For example, Bacon is able to induce Coulomb's Law, $f = q_1 q_2 / d^2$, which relates electrical force to the charges on two bodies and the distance between them, and a formula for the electrical current in a wire connected to a battery and a metal rod, $I = T/(R - L/D^2)$ which depends on the temperature differential of the bar and the internal resistance of the battery, and the length and diameter of the wire. The set of formulae that includes these is extremely large, and it seems unlikely that simple equation fitting would be an effective method for inducing formulae of this complexity.

Bacon's search method uses properties of the data to guide the formation of hypotheses. Other

induction systems have this capability including the concept-induction strategy of focusing described by Bruner et al. (1956), the method for inducing patterns in letter sequences studied by Simon and Kotovsky (1963), and AM's heuristics for generating new concepts based on properties of examples Bacon s heuristics involve properties of quantitative data and thus differ, as one would expect, from the heuristics of other systems such as AM, where the data involve categories of examples and sets of defining features. Bacon's use of data has the further interesting feature of creating new data in the process of evaluation hypotheses. In evaluating a hypothesis, Bacon calculates values of a new function of available data, and if the hypothesis does not succeed, those values become part of the data available to Bacon for further problem solving. Thus, though an attempt to solve the problem may fail, it leaves new results that may be instrumental in a later successful attempt.

Bacon's basic method is to search for a function of data that gives constant values across experimental conditions. As an example, the formula for the data in Table 9.10 is $d = kt^2$, where k is a constant, the form in which Bacon discovers the law is $d/t^2 = k$.

Bacon uses heuristic rules to form hypotheses, consisting of functions of variables in its data base that might give constant values. For example, if two quantities increase or decrease together, Bacon forms their ratio as a new quantity to be considered. If one variable decreases as another increases, Bacon forms their product as a new quantity. These heuristics, and another that forms linear functions of variables, enable Bacon to induce relatively complex functions. The first two are sufficient for the problem in Table 9.10. First, note that t and d increase together, and form the ratio t/d. Since this ratio decreases with t, Bacon forms the product t^2/d, which quantity is constant across the observations.

Some other heuristic methods are also used, including the definition of 'intrinsic variables' as properties of objects that are associated with constant values of quantities, and attempts to find a common divisor for values of intrinsic variables that have been induced. These heuristics enable induction of properties such as the resistances of different wires from measurements of current, and the atomic and molecular weights of chemical elements from data about the weights and volumes of elements and compounds involved in chemical reactions

As previously noted induction problems can also be understood as problems of design, especially when the structures that are induced are expressed explicitly as formulae This view is particularly appropriate to Bacon s induction of formulae. Consider the task as the construction of a formula using symbols for the variables in the problem. Bacon's heuristics then are rules for forming combinations of the symbols that may satisfy the problem criterion. Even if a formula does not solve the problem, it may provide part of the formula that is needed Thus, the process of search through the construction of partial solutions, which is characteristic of design problems, provides an appropriate characterization of Bacon's process of induction.

Bacon is not intended as a complete simulation of cognitive processes in scientific research, where hypotheses about causal mechanisms often play a critical role in the decision to measure variables or to examine quantitative relation. Even so, it demonstrates that quite simple heuristics are sufficient to produce quite complex inductive conclusions from quantitative data, and it is reasonable to suppose that these heuristics correspond to significant components of complex scientific reasoning.

INDUCING MOLECULAR STRUCTURE

Another scientific task that has been investigated is induction of the molecular structure of organic compounds. A system called Dendral induces molecular structure from data in the form of mass spectra (Lindsay, Buchanan, Feigenbaum, & Lederberg, 1980). A mass spectrum is a set of quantities of the fragments of various sizes that are produced when molecules of a substance are bombarded by electrons

Like AM and Bacon, Dendral performs induction using heuristic search. An important difference is that Dendral uses search heuristics that are based on principles that are specific to organic chemistry, whereas AM's methods apply to any structure of categorical concepts, and Bacon's methods can be applied to any quantitative data.

Dendral's method of induction has three main stages. First, the chemical formula of the compound is inferred from features of the mass spectrum Then hypotheses about molecular structures are generated with constraints based

...nd compounds

...iction problems
...blems of design.
...that are induced
...ulae. This view
...Bacon's induction
...as the construc-
...for the variables
...ies then are rules
the symbols that
...erion. Even if a
...problem, it may
...is needed. Thus
the construction
is characteristic
...an appropriate
...cess of induction.
a complete simu-
...cientific research,
...sal mechanisms
...the decision to
nine quantitative
...trates that quite
...to produce quite
from quantitative
...uppose that these
...icant components
...g.

...ture
...have been investi-
...cular structure of
...m called Dendral
...from data in the
...dsay, Buchanan,
...)). A mass spectrum
...ments of various
...molecules of a
electrons
...al performs induc-
...An important dif-
...search heuristics
...that are specific to
...s AM's methods
...egorical concepts,
...applied to any

...luction has three
...ical formula of the
...atures of the mass
...about molecular
...n constraints based

on knowledge of the class of compounds that the substance belongs to Finally, the hypotheses are tested by comparing their implications with the quantitative details of the mass spectrum, and the hypothesis that fits the data best is chosen.

The data used to infer the chemical formula are the peaks in the mass spectrum The largest mass represented is probably either the mass of the molecular ion or a mass smaller than the molecular ion by one fragment Differences between peaks usually correspond to the masses of fragments that are broken off in the bombardment. Dendral uses the value of the largest peak and the interpeak distances, along with knowledge of chemistry, to infer one or more chemical formulas that are consistent with the spectrum.

Dendral's next task is to generate possible molecular structures, with the ions in the formula arranged in ways consistent with known possible arrangements. There are many millions of possibilities for most problems, so Dendral formulates constraints based on knowledge of the class of compounds that the sample belongs to. With the constraints, Dendral constructs hypotheses about molecular structure with a method that first determines the maximum number of rings in the structure, then constructs the possible partitions of ions into rings and remaining components, and finally constructs the possible structures for each possible partition.

Dendral now tests its many hypotheses, using the quantitative details of the mass spectra. In the different hypothesized structures, since different components are separated by different numbers of bonds, there are differences in the likelihood of this occurring together in a fragment. Assuming that fragments are produced by breaking one or two bonds at once, predictions are made about the relative amounts of material to be found at each peak in the spectrum, and the structure that fits the data best is chosen.

Note that Dendral's task, like Bacon's, involves constructing an explicit formula to represent the structure it induces. Thus, its method can also be considered to solve problems of design, where the materials for the construction are symbols that represent the atomic components of chemical compounds, and the chemical knowledge that it uses constrains the search to arrangements of those materials that agrees with the mass spectra.

Diagnostic Problem Solving

In the problem solving tasks of troubleshooting in electronics and diagnosis in medicine, the problem solver has a space of stimuli consisting of one or more symptons and further information that can be obtained by performing tests. The space of structures is a set of possible causes of the symptoms—faulty components in electrical circuits or disease states in medical diagnoses.

In addition to its characteristics of inductive problem solving, diagnostic problem solving also has components of operational thinking, because it is based on the goal of curing a patient's illness or repairing a device. Thus the information and conclusions in the diagnosis are directed toward making a decision about a remedial treatment that should be applied.

Troubleshooting

The task in troubleshooting is to determine which of the many components of an electronic system is causing the system to function improperly. There may be more than one fault, but it simplifies the problem greatly to assume that there is only a single fault in the system.

In a general way, troubleshooting resembles the task of inducing categorical concepts when the subject chooses the stimuli for which information is given. In concept induction the problem solver obtains information by asking whether a specific stimulus is positive or negative. In troubleshooting, information is obtained by taking readings of voltage or current at specific locations in the circuit. In both tasks there are many possible hypotheses to be considered, but the set of possibilities can be specified: in concept induction it is the set of logical combinations of the stimulus attributes, and in troubleshooting it is the set of possible faults of components. These similarities in the tasks are correlated with an important resemblance in effective methods for working on the problems. The focusing strategy in concept induction uses information obtained about instances in order to eliminate classes of hypotheses, rather than considering each hypothesis individually as is done in the less effective scanning strategy (Bruner et al., 1956). Similarly, in troubleshooting an important component of strategy is to conduct tests that permit elimination of sets of possible faults from consideration. Use of this strategy is made possible by both a general knowledge of

electronic components and a knowledge of the specific circuit in the problem. This requirement of knowledge to support the process of induction is analogous to the role played in concept induction by knowledge of the alternative logical forms (conjunction, disjunction, etc.) and the truth-table combinations that correspond to them (Dodd et al., 1971), although the knowledge required in troubleshooting is considerably more elaborate.

A model of troubleshooting is included in a system called Sophie that provides computer-based instruction for trainees in electronics maintenance (Brown, Burton, & deKleer, 1983). The troubleshooting system provides a model for the student to observe in learning how to diagnose faults in a circuit. If the student specifies a fault in the circuit, Sophie can diagnose the fault, perform a series of tests to obtain readings of current or voltage at various points in the circuit, form hypotheses about the fault, and eventually arrive at a decision about it. Sophie has a store of general knowledge about electronics and an explicit representation of strategy that enables it to provide explanations of both the principles of electronics and the strategic purposes of its activity for tests that it is performing. Sophie's troubleshooting knowledge is also used to evaluate the problem-solving performance of students, by providing a series of problem-solving steps that can be compared with the steps taken by students.

Sophie's knowledge for troubleshooting has four main components: two components of electronics knowledge, a component of knowledge for making specific inferences, and a component of strategic knowledge. The component for specific inferences includes general knowledge in the form of 'experts' that have information about characteristics of different kinds of electronic components such as resistors and diodes. These experts can use data obtained from readings to calculate values for other variables, assuming normal functioning of components of the circuit; the inferred values can then be compared with actual readings of those variables.

A second component of Sophie's knowledge is information about the specific circuit that is used for instruction. The circuit is represented hierarchically as a set of modules with sub-modules and components. Possible functional states of each module and component are repre-

sented, including normal functioning and possible fault states. Experimental evidence obtained by Egan and Schwartz (1979) is consistent with a hypothesis that human electronics experts represent circuits in ways similar to Sophie's. Egan and Schwartz showed that experts encode information from circuit diagrams rapidly, much the way experts perform in other domains such as chess (see "Problems of Modifying Arrangements"), and that functional modules made up of components that are spatially contiguous in the diagram play an important role in the performance.

A third part of Sophie's knowledge involves specific actions that occur during troubleshooting. This knowledge is in the form of rules for making inferences about the states of modules and components of the circuit. Readings are used to eliminate hypotheses about faults by showing that a module is functioning normally, and for propagating inferences in the hierarchical representation; for example, if a component is faulted, then all the modules that contain that component must also be faulted.

The fourth component of knowledge is Sophie's strategy, a breadth-first search method with backtracking. Sophie considers all the possible states that can occur, according to its representation of the circuit, and eliminates possible faulty states on the basis of readings that are consistent with normal functioning. It assumes normal functioning of components until there is a reading that conflicts with that assumption; however, it keeps a record of the assumptions used in its inferences, and if information contradicts an assumption made earlier, inferences based on that assumption are revised.

Medical Diagnosis

In medical diagnosis, as in troubleshooting, a system—in this case, a human body—is functioning improperly, and the inductive task is to infer the cause of the malfunction. Also, as in troubleshooting, the purpose of the diagnosis is to determine a treatment that can remedy the malfunction, and the diagnostic activity is conducted in a way that provides information relevant to choosing a treatment.

Several systems have been developed that solve diagnostic problems in various domains of medicine, including diagnosis of infectious agents and prescription of antibiotics (Shortliffe, 1976), prescription of digitalis therapy for cardiac

patients (Silverman. 1975). and diagnosing and prescribing treatment for varieties of glaucoma (Weiss. Kulikowski. & Safir. 1977). (For a review. see Ciesielski. Bennett. & Cohen. 1977.) One system. Caduceus. which performs general diagnosis. is discussed here. along with empirical studies of diagnostic problem solving by physicians with varying amounts of training and experience.

A MODEL OF KNOWLEDGE FOR GENERAL DIAGNOSIS

Knowledge used in general medical diagnosis has been investigated in the context of a model named Caduceus (Miller. Pople. & Myers. 1982; Pople. 1982). The knowledge with which Caduceus diagnoses diseases is similar in important ways to the knowledge used by Sophie for diagnosing faults in electronic circuits. Its hierarchical form enables systematic search in the space of hypotheses. The Caduceus system also has rules that infer hypotheses from symptoms and test results. and that propagate the inferred information using the hierarchical structure of its knowledge.

Caduceus's knowledge about diseases is of two kinds, organized in separate but related graph structures. One of these, called a nosological graph, provides a taxonomy of diseases based on the organs of the body involved and on etiological factors. This graph groups diseases according to their manifestations. The other knowledge structure, called a causal graph, contains information about disease states and processes. The causal graph contains technical concepts of pathology that refer to states of disease, such as cardiogenic shock.

Caduceus has the goal of identifying one or more disease entities that provide a complete explanation of a set of symptoms and findings in the case. Subproblems are formulated from findings that are not yet integrated into an explanatory network; these constitute diagnostic tasks that are generated by the system. Identification of the disease depends mainly on the nosological graph. this hierarchical structure is used in a top-down search to narrow the possible disease entities. The information about the states and processes of disease in the causal graph provides links between hypothesized disease entities and the specific symptoms and test results that are available. Caduceus concludes its diagnostic analysis when an explanatory network has been developed that includes all the available symptoms and findings.

EMPIRICAL STUDIES OF DIAGNOSTIC PERFORMANCE

An extensive study of performance in diagnostic problems was conducted by Feltovich (1981. also described in Johnson, Duran. Hassebrock. Moller, Prietula. Feltovich. & Swanson, 1981). The results were consistent with the general properties of the Caduceus model. They also provide information about characteristics of knowledge for diagnosis at different levels of experience and expertise. Feltovich obtained problem-solving protocols for cases in pediatric cardiology from individuals varying in experience from fourth-year medical students who had just completed a six-week course in pediatric cardiology to two professors who had more than 20 years of experience in that subspecialty. Information from five cases was presented serially and the physicians gave their hypotheses and other thoughts about the cases, attempting to arrive at a correct diagnosis.

The performance of experts indicated that their knowledge differed from that of novices in several ways, consistent with the general features of expert knowledge in chess and Go discussed above in "Problems of Modifying Arrangements." The major difference was that experts had more integrated knowledge about diseases —more detailed knowledge of variation in disease states and more precise knowledge of relation between diseases and symptoms. For example. one advanced expert mentioned groups of hypotheses that were supported by the findings presented first and then used later information to narrow the range of possibilities. The other advanced expert used more of a depth-first strategy—proposing a likely hypothesis based on preliminary findings, but modifying the hypothesis in a flexible way when later evidence provided counterindications. The knowledge of novices was primarily in the form of a few specific disease forms used in textbook cases. The novices responded to early evidence by proposing reasonable hypotheses but were less likely to recognize the significance of later evidence and change their hypotheses when necessary. The sets of hypotheses mentioned by novices during problem solving were significantly smaller than those of the experts. In a study of expert and novice radiologists, Lesgold

et al (1981) came to similar conclusions regarding expert knowledge for diagnosis They found that in reading x-ray films experts generated representations in a three-dimensional system and used salient features to generate initial hypotheses that were refined or modified on the basis of more detailed features The knowledge necessary for recognizing features associated with abnormalities appeared to be well integrated with a general knowledge of anatomy. The integration of experts' knowledge was indicated by their ability to use features noted early as constraints on later interpretations (cf. Stefik. 1981) Novices—in this case. first-year residents in radiology—depended more on finding an explanation for a few features and to let other details be assimilated to the initial hypothesis rather than used to generate alternative hypotheses or modifications.

Conclusions from these studies of expert diagnosticians in medicine show close similarity to the studies of expert performance in other problem-solving domains, especially physics and chess. According to current findings, a major source of expert performance is the expert's ability to represent problems successfully. This results from the expert's having a well integrated structure of knowledge in which patterns of features in the problem are associated with concepts at varying levels of generality, enabling efficient search for hypotheses about the salient features of the problem that cannot be observed directly, as well as for methods and operations to be used in solving the problem.

EVALUATION OF DEDUCTIVE ARGUMENTS

The relation between human reasoning and formal logic has long been a subject of discussion and debate and, for some decades, a subject for experiment as well. It is generally agreed that human 'logical reasoning' does not always conform to the laws of formal logic. Formal logic is a normative theory of how people ought to reason, rather than a description of how they do reason. It is important, then, to develop a descriptive theory of human reasoning to compare and contrast with the logic norms.

Experiments aimed at developing a theory of human reasoning have mostly set tasks of judging the correctness or incorrectness of formal syllogisms. These tasks require application of the rules of deductive argument that are special in some ways. and correct performance depends on the subject s knowledge and use of the technical rules of formal deductive inference However. the processes used in these tasks do not differ in any fundamental way from those involved in problem solving in other domains. Psychological analyses provide no basis for a belief in deductive reasoning as a category of thinking processes different from other thinking processes. other than in the special set of operators that are permitted in rigorous deductive arguments. As Woodworth put the matter. "Induction and deduction are distinguished as problems rather than processes" (1938. p. 801).

Two tasks are discussed: First, we discuss propositional and categorical syllogisms, which present arguments in the sentential and predicate calculus: subjects frequently make errors in evaluating these syllogisms, and research has focused on why the reasoning process differs from correct logical inference. Second, we discuss linear syllogisms, which present arguments that depend on transitivity of order relations. Subjects make the transitive inferences in these tasks without difficulty, and psychological analyses have focused on the cognitive representation of information in the syllogisms.

Propositional and Categorical Syllogisms

Subjects in experiments on propositional or categorical syllogisms are asked to judge the validity of arguments such as the following (invalid) propositional syllogism:

> If I push the left-hand button. the letter T appears.
> I did not push the left-hand button.
> Therefore, the letter T did not appear.

The major premise states what will happen if the button is pushed. It says nothing about what will or will not happen if the button is not pushed. Hence the conclusion does not follow from the premises. Yet in a typical experiment (Rips & Marcus, 1977) a fifth of the subjects accepted this as a valid syllogism.

Categorical syllogisms in the predicate

calculus involve statements containing the terms *some*, *all*, and *no*. An example of a (valid) categorical syllogism is

> Some jewels are diamonds.
> All diamonds are valuable.
> Therefore, some jewels are valuable

Again, human subjects make frequent mistakes in judging whether certain kinds of categorical syllogisms are valid. For example, many subjects judge mistakenly that the following argument is a valid syllogism (Johnson-Laird & Steedman, 1978):

> Some As are Bs.
> Some Bs are Cs.
> Therefore, some As are Cs.

In experiments on syllogistic reasoning, the type of syllogism presented is most commonly taken as the independent variable, and the numbers of subjects that make errors on syllogisms of different kinds are measured. By comparing the error rates for different kinds of syllogisms, the experimenter seeks to formulate and test hypotheses about the cognitive processes that subjects use to make such syllogistic judgments.

For example, though many subjects will accept, "No As are Bs and no Bs are Cs, therefore no As are Cs," almost all will reject, "No As are Bs and no Bs are Cs, therefore all As are Cs." Yet both syllogisms are equally invalid. Such errors of reasoning have sometimes been attributed to an 'atmosphere effect.' In the example above, since *no* is present in both premises, it appears to be more acceptable than *all* in the conclusion (Woodworth & Sells, 1935). Alternatively, some investigators have claimed that the reason for these errors is that the quantifiers and connectives, *all*, *some*, *no*, *if . . . then*, *and*, *or*, do not have the same meanings in natural language as they do in formal logic (Braine, 1978). According to this hypothesis, since the experimenter judges the correctness of answers by their conformity to the rules of formal logic, whereas the subjects use the natural language meanings, errors follow when the two kinds of meaning diverge.

Errors and latencies in reasoning tasks depend not only on the form of the syllogism, but also on whether it has meaningful content (Wilkins, 1928). Thus, subjects may respond differently to the syllogism. "If some As are Bs and some Bs are Cs, then some As are Cs," and the syllogism "If some birds have blue eyes and some blue-eyed creatures are human, then some birds are human."

In general, subjects' error rates are lower when syllogisms have meaningful content, but there is an important class of exceptions. Subjects often reject valid syllogisms when the conclusions are contrary to facts known to them. "If all horses have four feet and all fish are horses, then all fish have four feet," may be rejected by subjects who know that fish are footless. The rate of rejection rises when subjects react emotionally to the conclusion. "If drug addiction is a disease and diseases should not be punished, then drug addiction should not be punished," is more likely to be rejected by subjects who support strong measures against drug usage then by those who do not (Janis & Frick, 1943; Lefford, 1946). Conversely, subjects often accept invalid syllogisms when the conclusions are consistent with their knowledge about the world or their preferences.

All these findings must be stated as 'tendencies,' since many subjects who make errors on some syllogisms of a certain form do not make such errors consistently. Moreover, there are large individual differences among subjects. For example, subjects trained in formal logic generally make fewer errors—not surprisingly—than do subjects without such training.

While human syllogistic reasoning conforms to some broad generalizations of the sorts that have been mentioned already, the findings derived from experiments are complex and confusing. In recent years, a few investigators have sought to cut through the confusion by creating models of the inference process or some components of it. The attempt to create such models has revealed features of the reasoning task that had not been entirely obvious.

Subjects may use any one of a wide range of strategies to solve the problems, and there is no reason to believe that all subjects use the same strategies. Subjects who reason by vague verbal analogies may succumb to the atmosphere effect, whereas subjects who create semantic images of the propositions and reason by operating on those images may make quite different errors. (Certain syllogisms may require the creation of images more complex than a subject can handle in memory.) Subjects' knowledge of logical

inference can be embedded in formal axioms or in-inference rules, with different consequences for the likelihood of error. The axioms that define connectives or the inference rules may conform to some natural logic that deviates from the formal-logic of the textbooks.

Several quite successful recent efforts at modeling have used the idea that evaluation of syllogisms is a form of problem solving similar to that discussed above in "General Knowledge for Novel Problems with Specific Goals." Using a set of inferential operators, the subject attempts to confirm the conclusion working from the premises, and accepts the conclusion if this problem-solving effort succeeds. The process typically used by subjects differs from the task of finding explicit proofs in that the inferential operators are not expressed overtly and need not, of course, correspond completely to the rules of-formal logic.

· Models of evaluating propositional syllogisms have been formulated by Osherson (1975), Braine (1978), and Rips (1983). These models are based on the concept of natural deduction, discussed by Gentzen (1935/1969). A system of natural deduction is a form of production system. Rules for making inferences specify conditions in the form of patterns of propositions, and when a pattern is matched in premises the inference is made. The models account for performance by postulating sets of inference rules assumed to be used implicitly by subjects. Rips also formulated a specific process of applying the rules and forming representations of the derivation. An interesting feature of Rips's formulation is the inclusion of suppositions that provide a backward-chaining component in the search process. A syllogism is judged valid if the system can generate a derivation of the conclusion from its inference rules.

The idea that sentential syllogisms are evaluated by natural deduction provides an interpretation of many of the kinds of errors that occur in syllogistic reasoning. Because it is an informal reasoning system, it is not surprising that it is susceptible to influence by general knowledge and affect. Performance would be expected to improve if subjects were taught a more explicit procedure for verifying the applicability of inference rules in evaluating syllogisms, and this result was obtained in the domain of geometry proofs in a study by Greeno and Magone (described in Greeno, 1983).

Models of reasoning for categorical syllogisms have been formulated by Guyote and Sternberg (1981) and by Johnson-Laird and Steedman (1978). These models use the idea that the information in premises is represented in the form of examples; for example. "Some jewels are diamonds" might be represented by a symbol for a jewel that is a diamond and another symbol for a jewel that is not a diamond. A representation based on the premises is formed and is used to evaluate the conclusion. Errors occur because the representations are incomplete: the examples generated by the system often fail to exhaust the possibilities, leading to incorrect conclusions.

Linear Syllogisms

In a linear syllogism, premises specify ordered relations between pairs of objects, and questions are asked about pairs for which the order was not specified. An example from Egan and Grimes-Farrow (1982) is:

> Circle is darker than square.
> Square is darker than triangle.
> Is triangle darker than circle?

(An alternative is to ask, "Which is darkest?" or. "Which is lightest?") Problems are presented with relations expressed differently, such as "Triangle is lighter than square," or "Triangle is not as dark as square," with the premise information given in different orders, and with different questions.

To answer the question, the information in the premises must be encoded in some representation that enables the answer to be derived. Three hypotheses about representation have been considered.

According to a spatial hypothesis (DeSoto, London, & Handel, 1965; Huttenlocher, 1968) information in the premises is integrated into an ordered list, possibly using an image in which symbols are spatially aligned. A representation for the example would be an ordering with circle first, square second, and triangle third, perhaps imagined in a vertical line with the circle at the top. Then a question such as, "Is circle darker than triangle?" would be answered by comparing the positions of the circle and the triangle in the ordered representation.

A second hypothesis (Clark, 1969) is that the representation consists of propositions in which

individual objects are associated with values of attributes For the example. circle would be associated with a large degree of darkness. square with a medium degree. and triangle with a small degree A question would be answered by retrieving representations of the objects in the question and comparing the properties associated with them.

The third hypothesis is that representation of binary relations are stored in memory This hypothesis assumes the simplest process of representation. since information in memory corresponds directly to the information in the premises To answer a quest.on, however, a sequence of propositions has to be retrieved, for example. to answer "Is circle darker than triangle" both "Circle darker than square" and "Square darker than triangle" have to be retrieved.

The hypothesis that binary relations are represented is ruled out by data obtained by Potts (1974). who had subjects study paragraphs containing series with six terms and asked questions (e g . "Does B precede D?") involving pairs that varied in their separation: with the ordering $A > B > C > D > E > F$, the pair $C > D$ has a separation of 0; the pair $B > D$ has a separation of 1; the pair $B > E$ has a separation of 2. and so on. If binary relations are represented in memory, questions about pairs with greater separation should take longer, since answers to these questions require more inferential steps. The finding was the opposite: it took less time to respond to items with greater separation. This finding has also been obtained with comparisons involving general knowledge, such as the relative sizes of animals (Banks, 1977).

The question whether premises are represented by an integrated spatial array or by propositions associating properties with individual objects has been harder to resolve. Huttenlocher (1968) provided an argument for the spatial hypothesis, including the finding that latency is shorter when the second premise has the third individual as the subject of the sentence (e.g., $A > B$, $C < B$ rather than $A > B$, $B > C$). The interpretation is that the subject imagines placement of the new object in a spatial array, and this is easier if the object is mentioned as the sentence subject than the sentence object. Clark (1969) argued for a propositional representation, presenting evidence that performance is influenced

by linguistic factors such as the congruence of questions with premises (e g . $A > B$, which is greater?" is easier than $B < A$ which is greater?")

Sternberg (1980) formulated models that specify stages of processing based on assumptions of a spatial or a propositional representation of premises. He also formulated a model that combines those assumptions. so that linguistic factors influence an initial encoding of premises. and relations among propositions influence conversion of the information into an integrated spatial array. The combined linguistic-spatial model provided a more accurate account of latency data than did either of the simpler models.

Several investigators have provided evidence that subjects do not all solve linear syllogisms in one way; rather, different subjects use different representations (Mayer, 1979; Sternberg & Weil. 1980). Egan and Grimes-Farrow's (1982) evidence was particularly direct. They used retrospective protocols obtained after solutions of individual problems. The protocols indicated that some subjects used spatial representations consistently, and others sometimes formed representations by associating certain objects in the problem with different quantitative values of attributes. The protocol evidence was substantiated by analyses showing that subjects differed in their performance. according to the representations they reported using. The order in which objects were mentioned was significant for subjects who used spatial representations, and the linguistic factor of consistency of the relational term used was significant for those subjects who sometimes used individual object propositions.

Conclusions

Until recently, little attempt has been made to establish a relation between research on reasoning and research on problem solving of the sorts discussed earlier in this chapter. Sometimes this separation has been justified on the grounds that syllogistic reasoning is 'deductive' whereas problem solving is 'inductive,' but we have seen that this distinction does not hold. Although a syllogism is a deductive structure, neither finding valid steps nor testing whether proposed steps are valid is a deductive process. Indeed, the major process in the evaluation of a propositional or categorical syllogism

is to seek a proof of the conclusion, the process discussed above in "General Knowledge for Novel Problems with Specific Goals" as the prototypical example of goal based problem solving. For linear syllogism problems, the major process is an example of inductive problem solving, as defined in "Induction," in which the subject forms an integrated representation of the premises using the structure of an ordered list induced from the order relations that the premises state.

Although all reasoning involves problem solving, it does not follow that there is no need for a special theory of syllogistic reasoning. To understand human reasoning, we must understand the meanings that people attach to words and the rules of inference that constitute their systems of 'natural logic' as well as the structure of the control system that guides their problem-solving search. Recent investigations show progress on these questions.

CONCLUSIONS

The literature reviewed in this chapter includes analyses of problem solving on a few dozen tasks One way to express the important general characteristics that have emerged here is to apply problem-solving analyses to a new domain. The analyses shown have provided strong guidance about the kinds of processes and knowledge structures that one should look for in an investigation of problem solving.

First, it is important to investigate the subjects' knowledge and processes for representing the problem If the subjects do not have special training in the problem domain, they must construct a problem space that includes representations of the problem materials, the goal, operators, and constraints. If subjects have special training or experience in the domain, their prior knowledge includes general characteristics of the problem space. and their representations of individual problems are based on that general knowledge. Experts in various domains are cognizant of the general methods that can be used for solving problems, and their representations include use of problem information relevant to the choice of a solution method.

A second major task is to characterize the problem representations that subjects form in their understanding of the problem In relatively unfamiliar domains. the problem solving is primarily a process of search, and the problem representation determines the space of possibilities in which the search will occur Some basic features of the problem space depend on the problem itself. A problem may present constraints on the operators that the subjects are permitted to use in trying to achieve a goal, or on the arrangement of materials that is acceptable as a solution. The problem may also require induction of a pattern or rule from materials presented. These alternatives lead to differences in the problem space. a space of possible sequences of actions. of possible solution arrangements, of possible structures. or some combination of these.

The problem space constructed by an individual subject is also determined by the method of search that the subject uses, the features of the problem that are used, and the general knowledge that is applied. In a problem of transforming a situation by a sequence of actions, subjects typically use some form of means-ends analysis. They may distinguish between features of the situation that are more-or-less essential for the solution, and they may organize their search by a process of planning that focuses on the more essential features. Searching in a space of possible solution arrangements typically involves generating partial solutions on a trial basis, and the search is influenced by the subjects' knowledge of constraints that can be used to limit the candidate arrangements that are considered. Similarly, solution of induction problems is influenced by the subjects' knowledge of general constraints on possible solutions, which may be used in generating and testing hypotheses. or in synthesizing or abstracting structures from the features of individual objects that are provided.

In solving problems for which subjects have special training or experience, the problem space of operators and constraints is provided by the subjects' existing knowledge. Experts have highly organized knowledge that includes solution methods and concepts for representing problems at varying degrees of generality and abstraction. For simple problems, experts' knowledge often provides a basis for immediate recognition of methods as well as detailed features relevant to the solution. Their knowledge of relations among methods and operators

is to seek a proof of the conclusion. the process discussed above in "General Knowledge for Novel Problems with Specific Goals" as the prototypical example of goal-based problem solving For linear syllogism problems. the major process is an example of inductive problem solving, as defined in "Induction." in which the subject forms an integrated representation of the premises using the structure of an ordered list induced from the order relations that the premises state.

Although all reasoning involves problem solving, it does not follow that there is no need for a special theory of syllogistic reasoning. To understand human reasoning, we must understand the meanings that people attach to words and the rules of inference that constitute their systems of 'natural logic' as well as the structure of the control system that guides their problem-solving search. Recent investigations show progress on these questions.

CONCLUSIONS

The literature reviewed in this chapter includes analyses of problem solving on a few dozen tasks One way to express the important general characteristics that have emerged here is to apply problem-solving analyses to a new domain. The analyses shown have provided strong guidance about the kinds of processes and knowledge structures that one should look for in an investigation of problem solving.

First, it is important to investigate the subjects' knowledge and processes for representing the problem If the subjects do not have special training in the problem domain, they must construct a problem space that includes representations of the problem materials, the goal, operators, and constraints. If subjects have special training or experience in the domain, their prior knowledge includes general characteristics of the problem space. and their representations of individual problems are based on that general knowledge. Experts in various domains are cognizant of the general methods that can be used for solving problems, and their representations include use of problem information relevant to the choice of a solution method.

A second major task is to characterize the problem representations that subjects form in their understanding of the problem In relatively unfamiliar domains. the problem solving is primarily a process of search. and the problem representation determines the space of possibilities in which the search will occur Some basic features of the problem space depend on the problem itself. A problem may present constraints on the operators that the subjects are permitted to use in trying to achieve a goal. or on the arrangement of materials that is acceptable as a solution. The problem may also require induction of a pattern or rule from materials presented. These alternatives lead to differences in the problem space. a space of possible sequences of actions. of possible solution arrangements, of possible structures, or some combination of these.

The problem space constructed by an individual subject is also determined by the method of search that the subject uses, the features of the problem that are used, and the general knowledge that is applied. In a problem of transforming a situation by a sequence of actions, subjects typically use some form of means–ends analysis. They may distinguish between features of the situation that are more-or-less essential for the solution, and they may organize their search by a process of planning that focuses on the more essential features. Searching in a space of possible solution arrangements typically involves generating partial solutions on a trial basis, and the search is influenced by the subjects' knowledge of constraints that can be used to limit the candidate arrangements that are considered. Similarly, solution of induction problems is influenced by the subjects' knowledge of general constraints on possible solutions, which may be used in generating and testing hypotheses, or in synthesizing or abstracting structures from the features of individual objects that are provided.

In solving problems for which subjects have special training or experience, the problem space of operators and constraints is provided by the subjects' existing knowledge. Experts have highly organized knowledge that includes solution methods and concepts for representing problems at varying degrees of generality and abstraction. For simple problems, experts' knowledge often provides a basis for immediate recognition of methods as well as detailed features relevant to the solution. Their knowledge of relations among methods and operators

and of constraints in the domain enables them to solve problems in a highly organized and planned manner.

The study of problem solving and reasoning has progressed to a substantial level of knowledge and theory, however, several questions remain unanswered.

First, while we are beginning to understand the performance of experts on simple problems, little is known about their performance on problems that are difficult and deep. When confronted by problems for which their knowledge does not provide a ready method of solution, do experts resort to weak methods of search and analysis fundamentally similar to those used by novices? Or do experts who have acquired powerful processes of reasoning in one domain apply those processes to solving problems in areas where specific solution methods have not been worked out and stored in memory?

A second question, closely related to the first, involves the general nature of problem solving in its more powerful and productive forms. In their discussions of productive thinking, Duncker (1935, 1945) and Wertheimer (1945, 1959) raised a critical issue that has not been dealt with in the recent literature, namely, the process of constructing more powerful representations of problems by analysis of problem components. The initial representation of a problem frequently fails to include important relations that are required for meaningful solution, although the problem solver is able to construct a reformulation that includes its important structural features.

A third question concerns learning. How is problem-solving skill learned? To analyze acquisition requires an understanding of the skills and knowledge to be acquired, and promising results in characterizing skill and knowledge in problem solving could provide a basis for the investigation of learning. New approaches to the acquisition of cognitive skill such as those of Anderson (1982), Anzai and Simon (1979), Neches (1981), and Neves (1981), may provide some keys to the analysis of learning processes.

A fourth question concerns the theoretical power of general principles in the analysis of problem solving and reasoning. The literature discussed in this chapter offers detailed hypotheses about performance on specific tasks that are testable at the level of their assumptions about specific processes. The more general assump-

tions are more heuristic These general concepts and principles provide guidance in constructing hypotheses about specific cognitive structures and processes, but they rarely constrain those hypotheses in wholly specifiable ways. It is an open question whether complex processes of problem solving and reasoning can be defined solely by underlying formal principles. Some investigators (Keil, 1981; VanLehn, Brown, & Greeno, 1984) have urged that research should seek general principles with deductive power that would determine characteristics of process models. Others (e.g., Newell & Simon, 1976) assert that there are good reasons for expecting that complex cognition is constrained only by relatively weak structural principles, of the kind that are characteristic of current theoretical analyses.

A review of any body of scientific research can be closed with the remark that much has been accomplished, and more remains to be done, and the psychology of problem solving and reasoning is no exception. The progress of the 1960s and 1970s has provided concepts and methods that future investigators may use as the basis for further advances.

REFERENCES

Adamson. R E. (1952). Functional fixedness as related to problem solving: A repetition of three experiments. *Journal of Experimental Psychology*, *44*, 288-291

Adelson, B. (1981). Problem solving and the development of abstract categories in programming languages. *Memory and Cognition*, *9*, 422-433.

Anderson, J.R. (1975). Computer simulation of a language acquisition system: A first report. In R.L. Solso (Ed.), *Information processing and cognition: The Loyola symposium*. Hillsdale, NJ: Erlbaum.

Anderson, J.R. (1976). *Language, memory, and thought*. Hillsdale, NJ: Erlbaum.

Anderson, J.R. (1977). Computer simulation of a language acquisition system. A second report. In D. LaBerge & S.J. Samuels (Eds.), *Perception and comprehension*. Hillsdale, NJ: Erlbaum.

Anderson, J.R. (1982). Acquisition of cognitive skill. *Psychological Review*, *89*(4), 396-406.

Anderson, J.R. (1983). *The architecture of cognition*. Cambridge, MA: Harvard University Press.

Anderson, J.R., Farrell, R., & Sauers, R. (1984).

Learning to program in LISP *Cognitive Science. 3.* 87-129.

Anderson. J.R.. Greeno. J G . Kline. P J . & Neves. D M. (1981) Acquisition of problem-solving skill. In J R. Anderson (Ed.), *Cognitive skills and their acquisition* Hillsdale. NJ Erlbaum.

Anderson. J R.. Kline. P J . & Beasley. C M. (1979). A general learning theory and its application to schema abstraction. In G.H Bower (Ed.). *The psychology of learning and motivation* (Vol. 13. pp 277-318). New York Academic Press

Anzai. Y & Simon. H.A (1979) The theory of learning by doing. *Psychological Review. 86.* 124-140.

Ashcraft. M.H. & Battaglia, J (1978). Cognitive arithmetic· Evidence for retrieval and decision processes in mental addition. *Journal of Experimental Psychology: Human Learning and Memory, 5,* 527-538

Ashcraft. M.H. & Stazyk. E.H (1981). Mental addition: A test of three verification models. *Memory and Cognition, 9,* 185-196.

Atwood, M.E. & Polson, P.G. (1976). A process model for water jug problems. *Cognitive Psychology, 8,* 191-216.

Banks, W.P (1977). Encoding and processing of semantic information in comparative judgments. In G.H. Bower (Ed.), *The psychology of learning and motivation: Advances in research and theory:* Vol. 11. New York: Academic Press.

Barnes, G.M. & Whitely, S.E. (1981). Problem restructuring processes for ill-structured verbal analogies. *Memory and Cognition, 9,* 411-421.

Bartlett, F.C. (1932). *Remembering: A study in experimental and social psychology.* Cambridge, England· Cambridge University Press.

Bartlett, F.C. (1958). *Thinking.* New York: Basic Books.

Beilin. H. & Horn. R. (1962). Transition probability effects in anagram problem solving. *Journal of Experimental Psychology, 63,* 514-518.

Bereiter. C. & Scardamalia, M. (1982). From conversation to compositions: The role of instruction in a developmental process. In R. Glaser (Ed.), *Advances in instructional psychology:* Vol. 2. Hillsdale. NJ. Erlbaum.

Birch, H.G. & Rabinowitz, H.S. (1951). The negative effect of previous experience on productive thinking *Journal of Experimental Psychology, 41,* 121-125.

Bloom, B.S. & Broder, L.J. (1950). *Problem solving processes of college students.* Chicago. University of Chicago Press.

Bobrow. D.G. (1968). Natural language input for a computer problem-solving system. In M. Minsky (Ed.), *Semantic information processing.* Cambridge, MA: MIT Press.

Bourne. L.E . Jr (1970) Knowing and using concepts *Psychological Review. 77.* 546 556

Bourne. L.E . Jr (1974) An inference model of conceptual rule learning In R Solso (Ed) *Theories in cognitive psychology* Washington DC Erlbaum

Bower. G H & Trabasso. T R. (1964) Concept identification. In R.C Atkinson (Ed). *Studies in mathematical psychology* (pp. 32-94) Stanford. CA Stanford University Press.

Braine. M.D S. (1978). On the relation between the natural logic of reasoning and standard logic *Psychological Review. 85,* 1-21.

Briars, D.J. & Larkin. J.H. (1984). An integrated model of skill in solving elementary word problems *Cognition and Instruction, 1.* 245-296.

Brooks, L. (1978). Nonanalytic concept formation and memory for instances. In E Rosch & B.B Lloyd (Eds.), *Cognition and categorization.* Hillsdale. NJ: Erlbaum.

Brown, J.S. & Burton, R.B. (1980). Diagnostic models for procedural bugs in basic mathematical skills. *Cognitive Science. 4,* 379-426.

Brown, J.S., Burton, R.R., & de Kleer, J. (1983). Pedagogical. natural language and knowledge engineering techniques in SOPHIE I, II, and III. In D. Sleeman & J.S. Brown (Eds.), *Intelligent tutoring systems.* New York: Academic Press.

Brown, J.S. & VanLehn, K. (1980). Repair theory: A generative theory of bugs in procedural skills. *Cognitive Science, 4,* 379-426.

Brown. R. (1973). *A first language: The early stages.* Cambridge. MA: Harvard University Press.

Brownell, W.A. (1935). Psychological considerations in the learning and teaching of arithmetic. In *The teaching of arithmetic: Tenth yearbook of the National Council of Teachers of Mathematics.* New York: Columbia University Press.

Brownell. W.A. (1941). *Arithmetic in grades I and II: A critical summary of new and previously reported research.* (Duke University Research Studies in Education. No. 6). Durham, NC. Duke University Press.

Bruner. J.S.. Goodnow, J.J . & Austin. G.A. (1956) *A study of thinking.* New York: Wiley.

Bundy, A. (1978). Will it reach the top? Prediction in the mechanics world. *Artificial Intelligence, 10.* 129-146.

Burke. R.J.. Maier. N.R.F., & Hoffman. L.R. (1966). Function of hints in individual problem solving *American Journal of Psychology. 79.* 389-399

Buss, A.H. (1953). Rigidity as a function of reversal and nonreversal shifts in the learning of successive discriminations. *Journal of Experimental Psychology, 45.* 75-81.

Carpenter. J A.. Moore. O K.. Snyder. C R., & Lisansky. E.S (1961). Alcohol and higher-order problem solving Quarterly Journal of Studies on Alcohol, 22. 183–222.

Carpenter. T P & Moser. J M (1982). The development of addition and subtraction problem-solving skills In T P Carpenter. J M Moser & T Romberg (Eds). Addition and subtraction. A cognitive perspective Hillsdale. NJ Erlbaum.

Chase. W G & Simon. H A (1973). Perception in chess. Cognitive Psychology, 4, 55–81.

Chi. M.T H.. Feltovich. P . & Glaser. R. (1981) Categorization and representation of physics problems by experts and novices Cognitive Science, 5, 121–152.

Ciesielski, V B.. Bennett. J.S , & Cohen, P R. (1977). Applications-oriented AI research. Medicine. In A. Barr and E.A. Feigenbaum (Eds.), The handbook of artificial intelligence. Vol. 2. Stanford, CA Heuristech Press.

Clark. H.H. (1969). Linguistic processes in deductive reasoning Psychological Review, 76, 387–404.

deGroot, A.D. (1965). Thought and choice in chess. The Hague: Mouton.

deKleer, J. (1975). Qualitative and quantitative knowledge in classical mechanics. (AI Lab Tech. Rep. No. AI-TR-352). Cambridge. MA. Massachusetts Institute of Technology

DeSoto. C.B.. London. M., & Handel, S. (1965). Social reasoning and spatial paralogic. Journal of Personality and Social Psychology, 2, 513–521.

Dodd, D.H.. Kinsman. R.. Klipp, R.. & Bourne, L.E.. Jr. (1971). Effects of logic pretraining on conceptual rule learning. Journal of Experimental Psychology, 88. 119–122.

Dominowski, R.L. & Ekstrand, B.R. (1967). Direct and associative priming in anagram solving. Journal of Experimental Psychology, 74, 84–86.

Duncker, K. (1945). On problem solving. Psychological Monographs, 58:5 (Whole No. 270). (Original version published in 1935, in German.)

Egan. D.E. & Grimes-Farrow, D.D. (198.). Differences in mental representations spontaneously adopted for reasoning. Memory and Cognition, 10, 297–307

Egan, D.E. & Schwartz, B.J. (1979). Chunking in recall of symbolic drawings. Memory and Cognition, 7, 149–158.

Eisenstadt, M. & Kareev, Y. (1975). Aspects of human problem solving: The use of internal representations. In D.A. Norman & D.E. Rumelhart (Eds), Exploration in cognition (pp. 308–346). San Francisco: Freeman.

Erickson, J. (1971). Problem shifts and hypothesis behavior in concept identification. American Journal of Psychology, 84(1), 101–111.

Ericsson. K A & Simon. H A (1980) Verbal reports as data Psychological Review 37. 255 251

Ernst, G W & Newell A (1969) GPS A case study in generality and problem solving New York Academic Press.

Evans. T G (1968) A program for the solution of geometric-analogy intelligence test questions In M Minsky (Ed.). Semantic information processing (pp. 271–353). Cambridge. MA. MIT Press

Feigenbaum. E.A. (1963). The simulation of verbal learning behavior In E.A. Feigenbaum & J Feldman (Eds.), Computers and thought New York McGraw Hill.

Feltovich. P.J. (1981) Knowledge based components of expertise in medical diagnosis. (Report No PDS-2) Pittsburgh: University of Pittsburgh. Learning Research and Development Center

Flower. L.S. & Hayes, J.R. (1980). The dynamics of composing: Making plans and juggling constraints In L.W. Gregg & E.R. Steinberg (Eds.), Cognitive processes in writing. Hillsdale, NJ: Erlbaum.

Franks, J.J. & Bransford, J.D. (1971). Abstraction of visual patterns. Journal of Experimental Psychology, 90, 65–74.

Friedland. P.E. (1979). Knowledge-based experiment design in molecular genetics. (Report No. 79-771). Stanford, CA: Stanford University, Computer Science Dept.

Gavunn, E.I. (1967). Anagram solution and spatial aptitude. Journal of Psychology, 65, 65–68.

Gelman. R. & Gallistel, C.R. (1978). The child's understanding of number. Cambridge, MA: Harvard University Press.

Gentzen, G. (1969). Investigations into logical deduction. In M.E. Szabo (Ed. and Trans.), The collected papers of Gerhard Gentzen. Amsterdam: North-Holland. (Original work published 1935.)

Gerwin, D. (1974). Information processing, data inferences, and scientific generalization Behavioral Science, 19, 314–325.

Gick, M.L. & Holyoak. K.J. (1980). Analogical problem solving. Cognitive Psychology, 12, 306–355.

Gick, M.L. & Holyoak. K.J. (1983) Schema induction and analogical transfer. Cognitive Psychology, 15, 1–38.

Glucksberg, S. & Weisberg, R.W. (1966). Verbal behavior and problem solving: Some effects of labeling in a functional fixedness problem. Journal of Experimental Psychology, 71, 659–664.

Greeno, J.G. (1976). Indefinite goals in well-structured problems. Psychological Review, 83, 479–491

Greeno, J.G. (1978). A study of problem solving. In R. Glaser (Ed.), Advances in instructional psychology: Vol. 1. Hillsdale, NJ: Erlbaum.

Greeno, J.G. (1983). Forms of understanding in mathematical problem solving. In S.G. Paris, G.M. Olson, & H.W. Stevenson (Eds.), Learning and motivation in the classroom. Hillsdale, NJ: Erlbaum.

Greeno, J.G., Magone, M.E., & Chaiklin, S. (1979). Theory of constructions and set in problem solving. Memory and Cognition, 7, 445–461.

Greeno, J.G., Riley, M.S., & Gelman, R. (1984). Conceptual competence and young children's counting. Cognitive Psychology, 16, 44–143.

Greeno, J.G. & Scandura, J.M. (1966). All-or-none transfer based on verbally mediated concepts. Journal of Mathematical Psychology, 3, 388–411.

Gregg, L.W. & Simon, H.A. (1967). Process models and stochastic theories of simple concept formation. Journal of Mathematical Psychology, 4, 246–276.

Groen, G.J. & Parkman, J.M. (1972). A chronometric analysis of simple addition. Psychological Review, 79, 329–343.

Groen, G.J. & Resnick, L.B. (1977). Can preschool children invent addition algorithms? Journal of Educational Psychology, 69, 645–652.

Grudin, J. (1980). Processes in verbal analogy solution. Journal of Experimental Psychology: Human Perception and Performance, 6, 67–74.

Guyote, M.J. & Sternberg, R.J. (1981). A transitive-chain theory of syllogistic reasoning. Cognitive Psychology, 13, 461–525.

Hayes, J.R. (1981). The complete problem solver. Philadelphia: Franklin Institute Press.

Hayes, J.R. & Flower, L.S. (1980). Identifying the organization of writing processes. In L.W. Gregg & E.R. Steinberg (Eds.), Cognitive processes in writing. Hillsdale, NJ: Erlbaum.

Hayes, J.R. & Simon, H.A. (1974). Understanding problem instructions. In L.W. Gregg (Ed.), Knowledge and cognition. Hillsdale, NJ: Erlbaum.

Hayes, J.R. & Simon, H.A. (1977). Psychological differences among problem isomorphs. In N.J. Castellan, P.B. Pisoni, & G.R. Potts (Eds.), Cognitive theory: Vol. 2. Hillsdale, NJ: Erlbaum.

Hayes-Roth, B. & Hayes-Roth, F. (1978). Cognitive processes in planning. (Report No. R-2366-ONR). Santa Monica, CA: Rand Corporation.

Haygood, R.C. & Bourne, L.E., Jr. (1965). Attribute- and rule-learning aspects of conceptual behavior. Psychological Review, 72, 175–195.

Heller, J.I. (1979). Cognitive processing in verbal analogy solution. Unpublished doctoral dissertation, University of Pittsburgh.

Henley, H.M. (1969). A psychological study of the semantics of animal terms. Journal of Verbal Learning and Verbal Behavior, 8, 176–184.

Huesmann, L.R. & Cheng, C. (1973). A theory for the induction of mathematical functions. Psychological Review, 80, 126–138

Hull, C.L. (1920) Quantitative aspects of the evolution of concepts. Psychological Monographs, 28 (Whole No. 20).

Hull, C.L. (1930). Knowledge and purpose as habit mechanisms. Psychological Review, 37, 241–256.

Hull, C.L. (1943). Principles of behavior: An introduction to behavior theory. New York: Appleton-Century-Crofts.

Hull, C.L. (1952). A behavior system. New Haven: Yale University Press.

Hunt, E.B. (1962). Concept learning: An information processing problem. New York: Wiley.

Hunt, E.B., Martin, J., & Stone, P.I. (1966). Experiments in induction. New York: Academic Press.

Huttenlocher, J. (1968). Constructing spatial images: A strategy in reasoning. Psychological Review, 75, 550–560.

Janis, I.L. & Frick, F. (1943). The relationship between attitudes toward conclusions and errors in judging logical validity of syllogisms. Journal of Experimental Psychology, 33, 73–77.

Jeffries, R., Polson, P.G., Razran, L., & Atwood, M.E. (1977). A process model for missionaries–cannibals and other river crossing problems. Cognitive Psychology, 9, 412–440.

Jongman, R.W. (1968). Het oog van de Meester. Amsterdam: van Gorcum.

Johnson, P E. Duran, A.S., Hassebrock, F., Moller, J., Prietula, M., Feltovich, P.J., & Swanson, D.B. (1981). Expertise and error in diagnostic reasoning. Cognitive Science, 5, 235–283.

Johnson-Laird, P.N. & Steedman, M. (1978). The psychology of syllogisms. Cognitive Psychology, 10, 64–99.

Judd, C.M. (1908). The relation of special training to general intelligence. Educational Review, 36, 28–42.

Katona, G. (1940). Organizing and memorizing. New York: Columbia University Press.

Keil, F.C. (1981). Constraints on knowledge and cognitive development. Psychological Review, 88, 197–227

Kelleher, R T. (1956). Discrimination learning as a function of reversal and nonreversal shifts. Journal of Experimental Psychology, 51, 379–384.

Kendler, H.H. & D'Amato, M.F. (1955) A comparison of reversal and nonreversal shifts in human concept formation behavior. Journal of Experimental Psychology, 49, 165–174.

Kendler, H.H. & Kendler, T.S. (1959). Reversal and nonreversal shifts in kindergarten children. Journal of Experimental Psychology, 58, 56–60.

Kintsch. W (1974) *The representation of meaning in memory* Hillsdale. NJ Erlbaum

Kohler W (1929) *The mentality of apes* New York Harcourt Brace

Kotovsky K & Simon, H.A (1973) Empirical tests of a theory of human acquisition of concepts for sequential events. *Cognitive Psychology. 4,* 399-424

Langley P (1981) Data driven discovery of physical laws *Cognitive Science. 5*(1). 31-54

Langley. P. Bradshaw, G L., & Simon, H.A (1983) Rediscovering chemistry with the BACON system. In R Michalski. J Carbonell. & T Mitchell (Eds.), *Machine learning An artificial intelligence approach* Palo Alto. CA Tioga Press.

Larkin. J H (1981) Enriching formal knowledge. A model for learning to solve problems in physics. In J R Anderson (Ed.), *Cognitive skills and their acquisition* Hillsdale. NJ· Erlbaum.

Larkin. J H (1982) *Spatial knowledge in solving physics problems.* (C.I.P.434). Carnegie-Mellon University. Department of Psychology.

Larkin. J H. McDermott. J. Simon, D.P . & Simon, H A. (1979) Models of competence in solving physics problems *Cognitive Sciences. 4,* 317-345

Larkin. J H & Reif. F. (1979). Understanding and teaching problem-solving in physics. *Journal of Science Education. 1*(2), 191-203.

Lefford. A (1946) The influence of emotional subject matter on logical reasoning. *Journal of Genetic Psychology. 34,* 127-151.

Lenat. D.B. (1982). *AM: Discovery in mathematics and heuristic search.* Stanford. CA. Stanford University Press.

Lesgold, A.M., Feltovich, P.J., Glaser, R., & Wang, Y. (1981). *The acquisition of perceptual diagnostic skill in radiology.* (Report No. PDS-1). University of Pittsburgh. Learning Research and Development Center.

Levine. M (1963) Mediating processes in humans at the outset of discrimination learning. *Psychological Review. 70,* 254-276.

Lewis. C.H. (1978). *Problem system models of practice effects* Unpublished doctoral dissertation. University of Michigan, Ann Arbor.

Lindsay, R K., Buchanan, B.G., Feigenbaum, E.A., & Lederberg, J. (1980). *Applications of artificial intelligence for organic chemistry.* New York: McGraw Hill.

Luchins, A.S (1942). Mechanization in problem solving. *Psychological Monographs, 54.* 6 (Whole No. 248).

Magone. M.E. (1977). *Goal analysis and feature detection as processes in the solution of an insight problem* Unpublished master's thesis, University of Pittsburgh.

Maier N R.F (1931) Reasoning and learning *Psychological Review 38* 332 346

Maier N R F (1945) Reasoning in humans III The mechanisms of equivalent stimuli and of reasoning *Journal of Experimental Psychology 35* 349 360

Malin, J T (1979) Information-processing load in problem solving by network search *Journal of Experimental Psychology Human Perception and Performance. 5,* 379-390.

Maltzman I. (1955) Thinking From a behavioristic point of view *Psychological Review. 62,* 275-286

Mayer, R.E. (1974) Acquisition processes and resilience under varying testing conditions for structurally different problem-solving procedures *Journal of Educational Psychology. 66,* 644-656.

Mayer, R.E. (1979). Qualitatively different encoding strategies for linear reasoning premises Evidence for single association and distance theories. *Journal of Experimental Psychology. Human Learning and Memory, 5,* 1-10.

McDermott. J. & Larkin, J.H. (1978). Representing textbook physics problems. In *Proceedings of the Second National Conference, Canadian Society for Computational Studies of Intelligence.* Toronto, Canada: University of Toronto.

McKeithen, K.B., Reitman, J.R., Rueter, H.H., & Hirtle, S.C. (1981). Knowledge organization and skill differences in computer programmers. *Cognitive Psychology, 13,* 307-325.

Millenson, J.R. (1967). *Principles of behavioral analysis.* New York: MacMillan.

Miller G A. Galanter, E.H., & Pribram, K.H. (1960). *Plans and the structure of behavior.* New York: Holt Rinehart, & Winston.

Miller R A., Pople, H.E., & Myers, M.D. (1982). Internist-1. An experimental computer-based diagnostic consultant for general medicine. *New England Journal of Medicine, 307,* 468-476.

Millward, R.P. & Wickens, T.D (1974). Concept-identification models. In D.H. Krantz, R.C. Atkinson, R. D. Luce, & P Suppes (Eds.), *Contemporary developments in mathematical psychology* Vol. 1. (pp 45 100) San Francisco W H. Freeman and Company.

Moeser, S D & Bregman, A S. (1972). The role of reference in the acquisition of a miniature artificial language *Journal of Verbal Learning and Verbal Behavior. 11,* 759-769.

Moore, O K & Anderson, S.B (1954). Modern logic and tasks for experiments on problem solving behavior *Journal of Psychology, 38,* 151-160.

Morgan, J L & Newport, E.L. (1981). The role of constituent structure in the induction of artificial language. *Journal of Learning and Verbal Behavior, 20,* 67-85

Mulholland, T.M., Pellegrino. J.W., & Glaser. R. (1980). Components of geometric analogy solution *Cognitive Psychology, 12,* 252-284.

Neches. R. (1981). *Models of heuristic procedure modification.* Unpublished doctoral dissertation. Carnegie-Mellon University, Pittsburgh.

Nesher. P. (1982). Levels of description in the analysis of addition and subtraction. In T.P. Carpenter. J.M. Moser. & T. Romberg (Eds.). *Addition and subtraction: A cognitive perspective.* Hillsdale, NJ: Erlbaum.

Neumann, P.G. (1974). An attribute frequency model for the abstraction of prototypes. *Memory and Cognition, 2,* 241-248.

Neves, D M. (1981). *Learning procedures from examples.* Unpublished doctoral dissertation, Carnegie-Mellon University, Pittsburgh.

Newell, A., Shaw, J.C., & Simon, H.A. (1957). Preliminary description of the general problem solving program I (GPS I). CIP Working Paper 7. Carnegie-Mellon University, Pittsburgh.

Newell, A. & Simon, H.A. (1972). *Human problem solving.* Englewood Cliffs, NJ: Prentice-Hall.

Newell, A. & Simon, H.A. (1976). Computer science as empirical inquiry: Symbols and search. *Communications of the Association for Computing Machinery, 19,* 11-126.

Nisbett, R.E. & Wilson, T.D. (1977). Telling more than we know: Verbal reports on mental processes. *Psychological Review, 84,* 231-259.

Norman, D.A. & Rumelhart, D.E. (1975). *Explorations in cognition.* San Francisco: Freeman.

Novak, G.S. (1976). Computer understanding of physics problems stated in natural language. *American Journal of Computational Linguistics,* Microfiche 53.

Osherson, D.N. (1975). *Logical abilities in children:* Vol. 3. Hillsdale, NJ: Erlbaum.

Paige, J.M. & Simon, H.A. (1966). Cognitive processes in solving algebra word problems. In B. Kleinmuntz (Ed.), *Problem solving.* New York: Wiley.

Pellegrino, J.W. & Glaser, R. (1982). Analyzing aptitudes for learning: Inductive reasoning. In R. Glaser (Ed.), *Advances in Instructional Psychology:* Vol. 2. (pp. 269-345). Hillsdale, NJ: Erlbaum.

Piaget, J. (1952). *The child's conception of number.* New York: Norton. (Original work published, 1941, in French.)

Polson, M.C., Restle, F., & Polson, P.G. (1965). Association and discrimination in paired-associates learning. *Journal of Experimental Psychology, 69,* 47-55.

Polson, P.G. (1972). A quantitative analysis of the conceptual processes in the Hull paradigm. *Journal of Mathematical Psychology, 9,* 141-167.

Polson. P G., Atwood. M E., Jeffries. R., & Turner. A. (1981). The processes involved in designing software. In J.R. Anderson (Ed.), *Cognitive skills and their acquisition.* Hillsdale. NJ Erlbaum.

Pople. H.E. (1982) Heuristic methods for imposing structure on ill-structured problems: The structuring of medical diagnostics. In P. Szolovits (Ed.). *Artificial intelligence in medicine* (pp. 119-185). AAAS Symposium Series. Boulder, CO: Westview Press.

Posner, M.I., Goldsmith, R., & Welton, K.E., Jr. (1967). Perceived distance and the classification of distorted patterns. *Journal of Experimental Psychology, 73,* 28-38.

Post, E.L. (1943). Formal reductions of the general combinatorial decision problem. *American Journal of Mathematics, 65,* 197-268.

Potts, G.R. (1974). Storing and retrieving information about ordered relationships. *Journal of Experimental Psychology, 103,* 431-439.

Quillian, M.R. (1968). Semantic memory. In M. Minsky, (Ed.), *Semantic information processing* (pp. 216-270). Cambridge, MA: MIT Press.

Reber, A.S. (1967). Implicit learning of artificial grammars. *Journal of Verbal Learning and Verbal Behavior, 6,* 855-863.

Reddy, D.R., Erman, L.D., Fennell, R.D., & Neely, R.B. (1973). The HEARSAY speech understanding system: An example of the recognition process. *3rd International Joint Conference on Artificial Intelligence.* Stanford, CA.

Reed, S K. (1972). Pattern recognition and categorization. *Cognitive Psychology, 3,* 383-407.

Reed, S.K., Ernst, G.W., & Banerji, R. (1974). The role of analogy in transfer between similar problem states. *Cognitive Psychology, 6,* 435-450.

Rees, H.J. & Israel, H.C. (1935). An investigation of the establishment and operation of mental sets. *Psychological Monographs, 46,* Whole No. 210.

Reif, F. & Heller, J.I. (1981). *Knowledge structure and problem solving in physics.* (Educational Science Paper 12.) University of California, Berkeley, Physics Department and Group in Science and Mathematics Education.

Reitman, J.S. (1976). Skilled perception in GO: Deducing memory structures from inter-response times. *Cognitive Psychology, 8,* 336-356.

Reitman. J.S. & Bower, G.H. (1973). Storage and later recognition of exemplars of concepts. *Cognitive Psychology, 4,* 194-206.

Reitman, W.R. (1965). *Cognition and thought.* New York: Wiley.

Reitman, W.R. & Wilcox, B. (1978). Pattern recognition and pattern-directed inference in a program. In

D.A. Waterman & F Hayes-Roth (Eds.), *Pattern-directed inference systems.* New York. Academic Press.

Resnick. L.B. (1983) A developmental theory of number understanding. In H.P Ginsburg (Ed.), *The development of mathematical thinking.* New York. Academic Press.

Restle. F (1962) The selection of strategies in cue learning. *Psychological Review, 69,* 329-343.

Restle. F. (1970). Theory of serial pattern learning: Structural trees. *Psychological Review, 77,* 481-495.

Riley. M.S., Greeno. J.G., & Heller. J.I. (1983). Development of children's problem-solving ability in arithmetic. In H.P. Ginsburg (Ed.), *The development of mathematical thinking.* New York: Academic Press.

Rips. L.J. (1983). Cognitive processes in propositional reasoning. *Psychological Review, 90,* 38-71.

Rips. L.J. & Marcus, S.L. (1977). Suppositions and the analysis of conditional sentences. In M.A. Just & P.A. Carpenter (Eds.), *Cognitive processes in comprehension.* Hillsdale, NJ: Erlbaum.

Rips. L.J., Shoben, E.J., & Smith, E.E. (1973). Semantic distance and the verification of semantic relations. *Journal of Verbal Learning and Verbal Behavior, 12,* 1-20.

Rosch. E.H. (1973). On the internal structure of perceptual and semantic categories. In T.E. Moore (Ed.), *Cognitive development and acquisition of language.* New York: Academic Press.

Rosch, E.H. (1975). Cognitive representations of semantic categories. *Journal of Experimental Psychology: General, 104,* 192-233.

Rosch. E.H. (1978). Principles of categorization. In E. Rosch & B.B. Lloyd (Eds.), *Cognition and categorization.* Hillsdale, NJ: Erlbaum.

Rosch, E.H., Mervis, C.B., Gray, W.D., Johnson, D.J., & Boyes-Braem. P. (1976). Basic objects in natural categories. *Cognitive Psychology, 8,* 382-439.

Rumelhart, D.E. & Abrahamson, A.A. (1973). A model for analogical reasoning. *Cognitive Psychology, 5,* 1-28.

Sacerdoti. E.D. (1977). *A structure for plans and behavior.* New York: Elsevier-North Holland.

Schank, R.C. (1972). Conceptual dependency: A theory of natural language understanding. *Cognitive Psychology, 3,* 552-631.

Schank, R.C. & Abelson. R.P. (1978). *Scripts, plans, goals, and understanding: An inquiry into human knowledge structures.* Hillsdale, NJ: Erlbaum.

Schoenfeld, A.H. (1979). Explicit heuristic training as a variable in problem-solving performance. *Journal of Research in Mathematics Education, 10,* 173-187.

Shortliffe. E.H. (1976). *Computer-based medical consultations: Mycin.* New York: American Elsevier.

Silverman. H. (1975) A digitalis therapy advisor (Rep. No MACTR-143.) Cambridge. MA MIT Computer Science Dept

Simon. D P & Simon. H A (1978) Individual differences in solving physics problems. In R Siegler (Ed.), *Children's thinking· What develops?* Hillsdale. NJ Erlbaum.

Simon, H.A. (1973). The structure of ill structured problems. *Artificial Intelligence. 4,* 181 201.

Simon, H.A. & Barenfeld. ' 9'' ·formation processing analysis of ' ʾ ʓ ʚses in problem solving. *Psychologi Review. 76,* 473-483.

Simon. H.A. & Feigenbaum. E A (1964) An information processing theory of some effects of similarity, familiarization, and meaningfulness in verbal learning. *Journal of Verbal Learning and Verbal Behavior, 3,* 385-396.

Simon, H.A. & Gilmartin. K. (1973). A simulation of memory for chess positions. *Cognitive Psychology, 2,* 29-46.

Simon, H.A. & Kotovsky. K. (1963). Human acquisition of concepts for sequential patterns. *Psychological Review, 70,* 534-546.

Simon, H.A. & Lea. G. (1974). Problem solving and rule induction: A unified view. In L.W. Gregg (Ed.), *Knowledge and cognition.* Potomac, MD: Erlbaum.

Skinner, B.F. (1938). *The behavior of organisms: An experimental analysis.* New York: Appleton-Century-Crofts.

Smith, E.E., Shoben, E.J., & Rips, L.J. (1974). Structure and process in semantic memory: A featured model for semantic decisions. *Psychological Review. 81,* 214-241.

Snow, R.E. (1980) Aptitude processes. In R.E. Snow, P.A. Federico, & W.E. Montague (Eds.), *Aptitude, learning, and instruction: Cognitive process analysis analyses of aptitude:* Vol. 1. Hillsdale, NJ: Erlbaum.

Soloway, E., Ehrlich, K., Bonar, J., & Greenspan. J. (1982). What do novices know about programming? In B. Schneiderman & A. Bodre (Eds.), *Directions in human-computer interactions.* Norwood, NJ: Ablex.

Staats, A.W. (1966). An integrated-functional learning approach to complex human behavior. In B. Kleinmuntz (Ed.), *Problem solving: Research, method, and theory.* New York: Wiley.

Stefik. M. (1981). Planning with constraints (MOLGEN: Part 1). *Artificial Intelligence, 16,* 1.:-140.

Sternberg, R.J. (1977). Component processes in analogical reasoning. *Psychological Review, 84,* 353-378

Sternberg, R.J. (1980). Representation and process in linear syllogistic reasoning. *Journal of Experimental Psychology: General, 109,* 119-159.

Sternberg, R.J. & Weil. E.M. (1980). An aptitude-