DOCUMENT RESUME

ED 154 775 IR 005 753

AUTHOR	Brown, John Seely; And Others
TITLE	Artificial Intelligence and Learning Strategies.
INSTITUTION	Bolt, Beranek and Newman, Inc., Cambridge, Mass.
SPONS AGENCY	Advanced Research Projects Agency (DOD), Washington, D.C.; Navy Personnel Research and Development Center, San Diego, Calif.
REPORT NO	BBN-3642
PUB DATE	Jun 77
CONTRACT	MDA903-76-C-0108; N00014-76-C-0083
NOTE	51p.
EDRS PRICE	MF-$0.83 HC-$3.50 Plus Postage.
DESCRIPTORS	*Artificial Intelligence; Cognitive Processes; *Comprehension; Educational Research; Electric Circuits; *Learning Processes; Learning Theories; *Problem Solving; *Research; Sequential Learning; Short Stories; Structural Analysis

ABSTRACT

Different kinds of basic knowledge and strategies necessary for comprehension are examined in three radically different domains: (1) the stories, (2) mathematical problems, and (3) electronic circuits. From analyzing the comprehension processes in these different domains, similarities have emerged in the role of planning knowledge and the strategies governing the application of that knowledge for synthesizing a deep structure analysis of a story, a math solution, or a circuit. Insights gained from these similarities can be applied to the problems of teaching learning strategies to students and developing an expanded theoretical basis for further research in learning strategies. (Author/CMV)

****
* Reproductions supplied by EDRS are the best that can be made *
* from the original document. *
****

To appear in:

Harry O'Neil (Ed.) Learning Strategies
N.Y.: Academic Press, 1978

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

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

ED154775

ARTIFICIAL INTELLIGENCE AND LEARNING STRATEGIES

John Seely Brown, Allan Collins, Gregory Harris

Bolt Beranek and Newman, Inc.
50 Moulton St., Cambridge, Mass.

June 1977

Acknowledgements:

We are indebted to Dr. Harry O'Neil whose suggestions on our firts draft led us to such a major revision that it now bears little resemblance to it. We would also like to thank Marilyn Jager Adams, Bob Donaghey, and Ned Benhaim for their suggestions and assistance on this paper.

This research was supported by the Advanced Research Projects Agency, Air Force Human Resources Laboratory, Army Research Institute for the Behavioral and Social Sciences, and Navy Personnel Research and Development Center, under Contract No. MDA903-76-C-0108 and by the Personnel and Training Research Programs, Psychological Sciences Division, Office of Naval Research and the Advanced Research Projects Agency, under Contract No. NQ0014-76-C-0083, Contract Authority Identification Number, NR154-379.

IR006753

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER BBN Report #3642	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) ARTIFICIAL INTELLIGENCE AND LEARNING STRATEGIES		5. TYPE OF REPORT & PERIOD COVERED Technical Report
		6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) John Seely Brown, Allan Collins, Gregory Harris		8. CONTRACT OR GRANT NUMBER(s) MDA903-76-C-0108 N00014-76-C-0083
9. PERFORMING ORGANIZATION NAME AND ADDRESS Bolt Beranek and Newman, Inc. 50 Moulton Street Cambridge, Massachusetts 02138		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS
11. CONTROLLING OFFICE NAME AND ADDRESS Defense Advanced Research Projects Agency 1400 Wilson Boulevard Arlington, Virginia 22209		12. REPORT DATE
		13. NUMBER OF PAGES
14. MONITORING AGENCY NAME & ADDRESS(if different from Controlling Office) Navy Personnel Research and Development Center, San Diego, California 92136		15. SECURITY CLASS. (of this report) Unclassified
		15a. DECLASSIFICATION DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)

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

Approved for public release; distribution unlimited

18. SUPPLEMENTARY NOTES

This research was supported in part by the Defense Advanced Research Projects Agency, Air Force Human Resources Laboratory, Army Research Institute for the Behavioral and Social Sciences, Navy Personnel Research and Development Center, and the Office of Naval Research.

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

Learning Strategies, Artificial Intelligence, Planning Knowledge, Deep Structure Trace, Top-down and bottom-up processing, learning, comprehension.

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

In this paper we examine the different kinds of knowledge and strategies necessary for "understanding" in three radically different domains -- namely stories, solutions to mathematical problems, and electronic circuits. From analyzing the understanding process in these different domains, some surprising similarities have emerged concerning the role of planning knowledge and the strategies governing the application of that knowledge for synthesizing a deep structure analysis of a story, a math solution, or a

(OVER)

DD FORM 1 JAN 73 1473 EDITION OF 1 NOV 65 IS OBSOLETE

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered) circuit. Insights gained from these similarities are applied to the problem of teaching learning strategies to students and of developing an expanded theoretical basis for further research in learning strategies.

ABSTRACT

In this paper we examine the different kinds of knowledge and strategies necessary for "understanding" in three radically different domains -- namely stories, solutions to mathematical problems, and electronic circuits. From analyzing the understanding process in these different domains, some surprising similarities have emerged concerning the role of planning knowledge and the strategies governing the application of that knowledge for synthesizing a deep structure analysis of a story, a math solution, or a circuit. Insights gained from these similarities are applied to the problem of teaching learning strategies to students and of developing an expanded theoretical basis for further research in learning strategies.

TABLE OF CONTENTS

Artificial Intelligence and Learning Strategies

John Seely Brown, Allan Collins, Gregory Harris

INTRODUCTION

The field of artificial intelligence grew out of the attempt in the late 1950's to build computer programs that could carry out tasks requiring human intelligence. The goal was to build machines that could understand language, recognize objects in scenes, act as intelligent robots, solve problems, play games such as chess, teach students about different subjects, etc. These problems have not been completely solved, but there has been a steady accumulation of tools and techniques in artificial intelligence, such that the programs designed to carry out these tasks have become more and more sophisticated (Bobrow and Collins, 1975; Schank and Abelson, 1977; Winston, 1977).

In order to build these programs, artificial intelligence has developed a variety of formalisms that in turn provide a new basis for analyzing cognitive processes. These formalisms are used to express structural and procedural mechanisms and theories about human problem-solving, planning, representing knowledge, and understanding text by computers. Our belief is that the cognitive and artificial intelligence theories expressible in these formalisms can begin to provide a domain-independent, theoretical foundation for research in learning strategies.

With the development of these formalisms, there has been renewed interest in what it means exactly to "understand" a piece of text, a set of instructions, a problem solution, a complex system, etc. This has repeatedly led to the realization that "understanding" requires different kinds of knowledge not explicitly referred to in the text or problem solution, as well as strategies for governing how this implicit knowledge should be used in synthesizing a structural model of the meaning of the text or problem solution. This model, which we call a deep structure trace, is a complex hypothesis about the plans and goals of the characters in the text or the person who solved the problem.

The intent of this paper is to explore the role of the different kinds of knowledge needed in the understanding process and to examine those insights we have gained about learning strategies through the recognition of the tremendous amount of tacit knowledge that must be exploited by students as they try to understand something. This is especially relevant for learning strategies because in analyzing comprehension tasks in a variety of _divergent_ knowledge domains, we have begun to see some surprising similarities in the kinds of strategies and knowledge used in these different domains. This suggests that there may be general learning strategies that will enhance a student's comprehension abilities over a wide range of content areas. Rigney (1976) has claimed that "The approach to teaching students cognitive strategies has been through content-based instruction and maybe that is wrong and should be reversed; i.e., content independent instruction." Rarely has anyone tried to make explicit or formalize the different kinds of strategies and knowledge needed for "understanding" something in even one content area--let alone in different ones. Perhaps that is why we have not seen powerful generality along the learning strategies dimension from content-based instruction.

One goal of a learning strategies curriculum might justifiably be to first teach the student all of the abstract, tacit knowledge and strategies that underlie problem solving and "understanding" for a particular content area, and then later to show him the generality of these strategies across content areas. Alternatively, a curriculum might teach the knowledge and strategies in a content-independent form, and then show how they apply to different content areas. Either approach would help the student to more readily acquire an understanding of that particular domain of knowledge. By transferring these skills, it would also have a significant effect in his ability to acquire other quite separate domains of knowledge.

It has been sometimes suspected that presenting a topic to students in the clearest way can be counterproductive in the long run, since they do not have to struggle with understanding the concept, and walk away expecting that real world situations will always be crisp, clear, and easy to grasp. At first glance this would seem to argue against articulating the tacit knowledge involved in understanding a concept or performing a task, since it might make "understanding" too easy and compartmental. However, it is precisely this lack of attention to tacit knowledge that often causes "optimal" presentations of a concept to have this effect. If a concept is explained without explicit reference to the complex processes necessary for understanding it, then the student will not be able to reconstruct the process himself. However if a concept is presented by showing how to successively refine one's understanding of the concept (or, more metaphorically, how to experiment with the concept in order to "debug" one's own understanding of it), such a presentation will not be counterproductive.

Before proceeding, let us restate our premises from a slightly different point of view. We believe that (1) by explicating the underlying domain-independent cognitive processes, strategies, and knowledge that a student must use to "understand" a new situation, text, set of instructions, solutions to a problem, etc. and (2) by finding ways to teach him a general awareness of these processes along with some learning strategies based on those processes, we can provide him with a foundation for acquiring new knowledge in the future and perhaps, more importantly, diminish his fear of being confronted with new conceptual material that he cannot instantly understand. How detailed these learning strategies must be in order to be effective is, of course, an open question. But simply making the student aware of the existence of some very simple strategies that are in concert with the cognitive processes involved in his synthesizing an "understanding" can be surprisingly useful. For example, the act of "understanding," in itself, can become less mysterious with the realization that comprehension is an active process requiring the <u>formation and revision of hypotheses</u> about the meaning of a given event or situation.

In this regard, we are reminded of an apocryphal story of a teacher who gave a young student a problem to work out. After several minutes of attempting (and failing) to solve it, the student asked for help, and was told to return to his chair and to THINK about it some more. At this point the student broke into tears, exclaiming that everybody tells him to "think", but he doesn't have the slightest idea of what that means! Naturally, he felt terribly frustrated.

"Thinking" was something that he could not see, feel, or touch. It seemed to him that everyone assumed he knew the secret to this magical process. When he was told to think about something, all he could do was stare blindly at the problem and panic. He kept wondering why no one would tell him the secret. Today, schools are flooded with experimental programs to teach students to "think" (a la problem-solving) but where are these students being taught how to understand something new on their own, let alone what it means to "understand"?

In the next three sections we will analyze the knowledge and strategies underlying three radically different domains: understanding stories, problem solving in mathematics, and understanding electronic circuits. We will proceed by describing the cognitive processing of a person performing these three tasks in terms of artificial intelligence concepts. This analysis is not meant to be definitive. Rather it is suggestive of a kind of analysis and concern that might be beneficial to learning strategists. We will conclude the paper with a discussion of the central ideas that have emerged from studying the invariances over these disparate domains. We will also specify the implications this analysis has for a learning strategies curriculum, and suggest some techniques that might be useful in teaching these strategies.

ERIC

UNDERSTANDING A STORY

We will begin with story understanding, since this is the domain that has been analyzed most thoroughly and because it is easier to understand the artificial intelligence terminology in a familiar context. At the same time, we think the reader may find it surprising how much problem solving knowledge is involved in the comprehension of a story. We have chosen an Aesop fable called Stone Soup that requires a fair amount of problem solving to interpret both the characters' actions and the author's intentions.

Stone Soup

A poor man came to a large house during a storm to beg for food. He was sent away with angry words, but he went back and asked, "May I at least dry my clothes by the fire, as I am wet from the rain?" The maid thought this would not cost anything, so she let him come in.

Inside he told the cook that if she would give him a pan, and let him fill it with water, he would make some stone soup. Since this was a new dish to the cook, she agreed to let him make it. The man then got a stone from the road and put it in the pan. The cook gave him some salt, peas, mint, and all the scraps of meat that she could spare to throw in. Thus the poor man made a delicious stone soup and the cook said, "Well done! You have made a silk purse from a sow's ear."

Surface Structure and Deep Structure Traces

The story recounts a set of events that occurred as the poor man solved the problem of obtaining food. This set of events is the surface structure trace of the story. They are the result of the man's problem-solving activity.

To understand this story in any deep sense, the reader must construct an interpretation of these events of the following type (Adams and Collins, in press):

1. The poor man is prevented from obtaining his initial goal.
2. He uses a clever means to get part way to the initial goal.
3. He then uses an even cleverer means to reach the initial goal.

This understanding of the story is not a simple trace of how the events in the story are linked up, but rather a deep structure trace; that is, it is not at all obvious from the surface form of the story. The reader must reconstruct from the surface events how the poor man solved the problems he faced in the story.

The reader must utilize many different types of knowledge in order to construct such an interpretation of Stone Soup. We will try to illustrate them and then we will briefly try to recount how a skilled reader uses these different kinds of knowledge to understand the story on several different levels.

Basic World Knowledge and Schema Theory

A large amount of basic knowledge about the world is necessary to understand Stone Soup: (1) that servants work in large houses of wealthy people, are paid with room and board and small amounts of money, and want to please their employers; (2) that maids clean and take care of the residence and play the role of butler if there is none; (3) that to make soup a cook slowly heats a base of some meat, bones, or vegetables plus other ingredients in water over a low fire; (4) that fables are short stories with a moral; designed to explain people's motivations or actions; and (5) that a moral is the summary of a story structure, usually in terms of what a person should do in a given situation, and usually in the form of a proverb or maxim. These are English descriptions of a small part of the basic knowledge readers have about these concepts.

Schema theory (Bartlett, 1932; Minsky, 1975; Rumelhart and Ortony, 1977; Winograd, 1975) provides a very general formalism for representing different types of knowledge, including the basic knowledge described above and the planning knowledge described below. One of the fundamental notions associated with schemas is that they have various slots for variables (Minsky, 1975) that can be filled with different values. For example, the slot of the master can be potentially filled by any adult and the place where the cook heats meals can be filled by a stove, oven, fireplace, etc. Associated with each slot are default values, which will be assumed if no value is specified. For example, the default value for a master is the owner

ERIC

of the house," and the default cooking place is a stove. Thus associated with any slot is information about the range of values that can fill that slot plus the most likely values to fill it for particular predictable contexts.

Means-Ends Analysis.

Means-ends analysis is a procedural formalism developed by Newell and Simon (1963) in the General Problem Solver (GPS) that was designed to simulate human problem-solving. Means-ends analysis operates as follows: If there is a method to reach a goal directly, then that method is applied. If there is none, then a subgoal is generated that reduces the difference between the present state and the goal. If there is a method to reach the subgoal directly, then this method is applied; otherwise a sub-subgoal is generated, etc. Often a potentially useful method cannot be directly applied because the prerequisites for that method have not been met. In this case, a new subgoal is generated that tries to alter the given state of affairs so as to enable the application of this method.

In Stone Soup the man is hungry and is trying to achieve the goal of obtaining food. First he applies the method of begging for food but fails. However he does achieve the prerequisite of attracting the attention of someone in the house. He then pursues the subgoal of getting into the kitchen near the fire, which reduces the difference between his current state and his goal. Then he pursues the second subgoal of getting the cook to help him make soup, further reducing the difference. As Newell and Simon argue, people solve problems in everyday life and understand the actions of other people in terms of the means-ends strategies described here.

Applying means-ends analysis to a problem-solving situation produces a tree structure of goals and subgoals. The tree structure for Stone Soup is illustrated in Figure 1. The events in the story are the terminal nodes in the tree structure--the begging, the asking permission, the going inside to the fire, etc. The deep structure trace is the structure of goals and subgoals above the terminal nodes.

Recently several researchers (Mandler and Johnson, 1977; Rumelhart, 1975, 1977) have developed story grammars to specify the structure of well-formed stories. These story grammars are formalisms

Figure 1

Nested Tree Structure of Goals and Subgoals for Stone Soup

```
Goal: Become fed
  Subgoal: Obtain food
    Method: Beg --> Fails
      Subgoal: Get inside
        Subgoal: Con maid for permissi[illegible] inside
          Method: Ask permission to d[illegible]mself --> Succeeds
        Subgoal: Move inside
          Method: Walk (Default value) --> Succeeds
    Subgoal: Obtain soup
      Subgoal: Con cook into giving him soup
        Subgoal: Con cook into helping him make soup
          Subgoal: Con cook for permission to make stone soup
            Method: Bargain recipe for permission --> Succeeds
          Subgoal: Get pan
            Method: Ask cook --> Succeeds
          Subgoal: Get stone
            Method: Go out to road --> Succeeds
          Subgoal: Cook soup
            Method: Heat stone in water filled pan --> succeeds
          Subgoal: Con cook into adding scraps of food
            Method: Bargain his contributions for hers --> Succeeds
    Subgoal: Ingest soup
          Method: Drink (Default value) --> succeeds
```

for specifying the possible _target structures_ that a story's deep structure trace must fit. In fact, they define the set of tree structures that mean-ends analysis would produce. Thus, they are compact representations for the target structure that a reader must construct in order to understand a story. Story grammars are specific to the domain of stories, but there are similar target structures that guide understanding in other domains.

Planning Knowledge

The problem solving strategies based on means-ends analysis provide a domain-independent method for constructing plans. Abelson and Schank (Abelson, 1975; Schank and Abelson, 1977) have developed a _deltact theory_ to account for the way that people construct social plans. In particular they are trying to specify in formal terms the goals and methods that apply in a social context using means-ends analysis.

A deltact does two things: it permits the factorization of the differences between a situation and an arbitrary social goal into a few familiar categories and it gathers together all the methods that might make each difference category reducible in an actual situation. A deltact is a class of acts that reduce a certain difference to achieve a social goal. That class is broken into methods, which are ordered to suggest which to try first. A method of a deltact is something that is done (a segment of a plan) plus the preconditions under which it may be expected to reduce the deltact difference as promised. The plan segment may contain anything: it can be as vague as a goal to satisfy or as specific as a goal + deltact + method already instantiated or a specific action to take "plugging in" the actual participants. The claim is that with this _planning knowledge_, elaborate plans can be constructed (such as those of the man in _Stone Soup_), and complicated sequences of actions can be understood.

Some aspects of deltact theory can be illustrated in terms of _Stone Soup_. Two deltacts serve high level goals in the story: (1) the man's goal of obtaining food is accomplished by a ΔHAVE (a change of possession) and (2) the goal of getting into the house is accomplished by a ΔPROX (a change of proximity). Each of the deltacts has a formal definition: a ΔHAVE has five variables, an _actor_ causes an object to change possession from the possessor to the receiver by some means. In Stone Soup the poor man causes soup to change possession from the cook to himself by some unspecified means. There is an ordered set of methods for obtaining a ΔHAVE. The methods are: ASK; INFORM REASON; BARGAIN OBJECT; BARGAIN FAVOR; THREATEN; OVERPOWER; STEAL. In Stone Soup the poor man first tries an ASK to obtain food from the maid, but ends up using a BARGAIN FAVOR with the cook; buying is a special case of BARGAIN OBJECT, where the object bargained is money. The methods are ordered according to the priority in which they should be used in constructing a plan, but the order changes in different contexts.

Each of these methods is formally defined as an act with various prerequisites and results (Charniak, 1975; Schank and Abelson, 1977). In Schank and Abelson's theory an ASK has the following prerequisites: (1) the asker is near the askee, (2) the askee knows the information, and (3) the askee wants to transmit the information to the asker. The result is that the askee transmits the information to the asker, which in turn causes the asker to have the information. When any precondition for applying a method is not satisfied, then a deltact can be used to propose a plan for obtaining the required precondition. This is how subgoals are generated in the theory.

Essential to Stone Soup is the notion of a CON. A CON has the same structure of prerequisites and results as does a method, but the actual act involved in a CON may be any of the methods that Schank and Abelson (1977) describe. In conning the maid, the man used an ASK, whereas in the case of conning the cook he used a BARGAIN FAVOR. What the reader knows about a CON is its result--X gets nearer his goal (G1) and its prerequisites: (1) X must have a goal G1, (2) Y must have a goal (G2) to prevent X from obtaining G1 and a plan for G2 that X and Y believe will work, and (3) X must perform some act that Y thinks is directed toward a different goal (G3) and that helps X obtain G1 without Y giving up either the goal G2 or the plan for it. To identify a CON in reading a story, the reader must match the preconditions of any act in the story against the prerequisites of a CON, and find (or guess) all of the participants by name.

Deltacts illustrate the notion of a difference. They facilitate the reduction of differences by suggesting methods for the means-ends analysis that will reduce some or all of the differences that have been noticed in a pair of actual and desired situations.

Means-ends analysis operates by searching among the known means or methods for those that will attain the ends that are sought, expressed, in terms of such differences. This analysis depends on two pieces of planning knowledge: an index of the known methods in terms of descriptions of the differences they can be expected to reduce, and a technique for computing the differences from a given situation plus goal so that a method indexed by that difference can be applied.

Strategic Knowledge For Understanding

Strategic knowledge refers to knowledge that the reader uses to drive the process of trying to make sense out of the story. It is the most elusive kind of knowledge because it is not at all apparent in any trace the reader may leave of his understanding (such as a summary of the story). Perhaps for this reason there are no explicit theories of what comprises this knowledge.

We will list a few general principles that skilled readers must use in understanding stories, as a first attempt to specifying what some of this knowledge must look like:

1. The deep structure trace constructed by the reader should make a well-formed story. (Things such as episodes should begin and end.)
2. The deep structure trace should somehow accommodate every event in the story.
3. Every slot in the schemas used to understand the story should be filled, preferably by values specified in the story, or by default values that do not contradict anything in the story.
4. Authors write stories for particular purposes and the reader should construct an interpretation of the author's intentions as well as an interpretation of the events in the story.
5. The reader should reread to synthesize a new interpretation of the story, if any of the strategic conditions (such as the four above) are not satisfied.

Principles such as these must be operating as a skilled reader tries to make sense of the story, but there may be many more such principles.

ERIC

Constructing and Revising Hypotheses about Deep Structure Traces

With this glimpse of the various kinds of knowledge needed to understand a story, we will briefly describe how a skilled reader synthesizes this knowledge to understand Stone Soup. Comprehension involves a notion of variable binding, where elements in the story are bound to slots in different knowledge structures. Where a value is not specified in the story, it must be assigned a default value. For example, when the man begs for food, this is bound to the method ASK, which in turn is bound to the goal of pursuing a ΔHAVE to obtain food. The reader makes the default assumption that the man's ultimate goal is eating to alleviate his hunger, rather than giving the food to his dog, for instance. The way objects and events in the story invoke different pieces of knowledge to suggest hypotheses is called bottom-up processing. For example, the poor man's begging for food suggests he wants to eat. In contrast, the way that knowledge schemas compete to provide the best hypothetical account for the input data is called top-down processing. For example, the goal of getting inside the house competes with that of getting dry as an explanation of why the man asked permission to dry himself by the fire. Taken together these two processes allow the reader to piece together the large amount of structure necessary to integrate the structural fragments in the text (Adams and Collins, in press; Rumelhart and Ortony, 1977).

The skilled reader uses bottom-up and top-down processing to formulate hypotheses about the deep structure underlying the various events as he encounters them in the story. As suggested earlier, the reader makes the inference that the man wants food because he is hungry. The default method of obtaining food is buying it--BARGAIN OBJECT, but this is not possible since the man is poor. The ASK is thwarted by the maid's refusal. There are a number of alternative methods but apparently the man goes off to beg elsewhere. He then returns to ask if he can dry himself by the fire. Apparently he has changed his top-level goal in the means-ends analysis structure. This change agrees with the fact that it is raining, and drying himself could be a reasonable lower-order goal in the man's goal structure. Thus an intelligent reader at this point may be led into constructing an incorrect hypothesis as to why the poor man asks to dry himself by the fire.

ERIC

When the man suggests making stone soup, the reader may construct a second incorrect hypothesis. The apparent goal is that he wants to teach the cook a new recipe. There are several clues, however, that allow the reader to formulate a different hypothesis about the man's goal: (1) Stones are not in the variable range of things from which one normally makes soup; (2) Stones have no food content; and (3) Because the man helped make the soup, the reader infers that he gets to eat it. This satisfies the original top-level goal that the reader constructed for the man's coming to the house in the first place. However, these three facts should lead the reader to construct a new goal structure in which making the stone soup is a subgoal, beneath the higher level goal of eating. In this new structure, making stone soup is a BARGAIN FAVOR for the subgoal ∆HAVE to obtain food. This nesting of the goal structure eliminates an unmotivated change of goals. Until this revision is made, the third strategic principle we named above--no important slots left unfilled--is violated because there is no purpose for the change of goals.

Given this restructuring, the reader should also be able to revise his earlier incorrect hypothesis as to why the man wanted to dry himself by the fire. To do this, the reader must notice that the man needed a ∆PROX to get into the kitchen, where he could bargain with the cook to obtain food. Thus asking to dry himself by the fire can be subsequently interpreted as a CON: it moved the man closer to the initial goal of eating that the maid had prevented; however once she thought the action was directed toward the goal of getting dry she allowed him to achieve his subgoal. By restructuring these two acts under the one goal of obtaining food, the reader has produced a tree structure that fits the constraints on a well-formed story according to Rumelhart's (1975) story grammar.

Knowing that this story is a fable, the skilled reader should infer that the story has a moral, a prescription for how to behave in a given situation. The reader can convert the deep structure trace he constructed to a moral something like the following: If one method fails, you can often reach your goal by a more circuitous one that is clever but not immoral. There are several underlying aspects to this moral: initial failure, persistence or repeated trying when you fail, changing methods when you fail, devising multi-step plans, using clever means such as a CON, not using immoral means such as THREATEN, OVERPOWER, and STEAL (see the section on Planning Knowledge above), and finally succeeding. The avoidance of immorality can only be realized from the reader's knowledge of the alternative methods the man didn't use. This ability to evaluate another person's plan derives from the reader's own ability to plan.

Furthermore, if the reader knows that the points of fables are often proverbs, he might be able to select the correct one for Stone Soup. This is done by matching the various aspects of the deep structure trace of the story against the deep structure trace of any candidate maxims. For example, "If at first you don't succeed, try, try again" matches the two aspects of failure and repeated trying. "Where there's a will there's a way" matches four aspects fairly well: persistence, changing methods, using a circuitous (multi-step) plan, and ultimate success. Neither of these proverbs matches perfectly or includes the cleverness aspect, but that's why we have fables.

By tracing the process of understanding through different stages, we have tried to show: (1) The problem solving processing necessary to forming hypotheses about the underlying structure; (2) The way the reader must construct revised hypotheses from the incorrect ones; and (3) How notions of means-ends analysis, goals, and methods for achieving those goals are integral to the understanding and evaluation of social events in the world. In particular, we have hinted at the quality between problem-solving and understanding where we, in part, achieve an understanding of this fable by recapitulating a hypothetical trace of how the beggar was achieving his goals, what his methods and intentions were at each step, and so on.

We, as readers, must actively invoke our own problem solving strategies in synthesizing a deep structure model or understanding of this story so as to be able to bridge the gaps between each line in the story. Hence, we see that even in simple stories, the reader cannot expect to be given or told everything. Indeed, he must participate, so to speak, in the event that he is trying to understand. This often happens almost unconsciously since the planning knowledge and problem solving strategies needed to participate are thoroughly ingrained in our heads. However, understanding less common events (instructions, systems, etc.)

requires an active invocation of this knowledge, as we shall see in considering the less natural domains of mathematics and electronics.

ERIC

UNDERSTANDING ELEMENTARY MATHEMATICS

In the previous section we discussed how higher-order knowledge in the form of plans, methods, and hypothesis construction strategies in the area of social interactions must often be used in order to understand even simple stories. In this section, we will sketch out an analysis of the understanding of a solution to an exercise in elementary mathematics that directly corresponds to our preceding analysis of story comprehension. The correspondence is between the strategies and processes used to conjecture and fill in the unmentioned plans in a story and those used to fill in the motivations for the steps in a solution to a math problem. In both cases, the lines comprising the surface structure of the story or solution must be <u>augmented</u> by the understander before a deep structure trace constituting an understanding can be generated.

While studying mathematics, probably everyone has experienced at one time or another the phenomenon of the almost magical nature of mathematic proofs or solution paths--the steps leading to a solution--that are encountered in studying most mathematics textbooks. Somehow the critical lines of a proof or critical steps in a solution seem to be pulled out of thin air, leaving one in awe about how these steps were ever conceived of or selected. Although each step of the proof, or solution, seems plausibly true as it is read, the proof as a whole is hard to remember; one could not summarize it except by reciting it verbatim from memory--much like what one does for a story which makes no sense, or a magic act in which the trick remains unknown. Worse, the proof as a whole does not seem to bear more than a coincidental resemblance to other proofs that are presented "on the same subject." For a student to develop the skills to understand, as opposed to rotely memorizing, a new solution--let alone skills to create his own solutions--the sense of what makes one proof or solution "like" the others is needed. In short, the answer is "there," but a student who does not know what to look for cannot really see it.

For the rare student who <u>has</u> seen how it all fits together, a newly "worked solution" seems well-planned, a deliberate sequence of steps culminating in the desired result. The steps are often so directly justified and self-evident that after a while the student begins to speak of steps "falling right out" and "moving toward the solution"--spatial metaphors. These metaphors are the inarticulate allusions to habits of thought in which the same knowledge is used in solution after solution. This planning and strategic knowledge is identical in structure and function to that used in story understanding.

To demonstrate this thesis, let us start with a concrete example drawn from Bundy (1975). Consider the task of solving the following equation:

$$\log (x+1) + \log (x-1) = 3$$

We seek all the expressions for x that make this equation true. (These logarithms are in base 2.) How are we to proceed? Observe that knowing all the basic mathematical transformations (e.g. commutativity, associativity) gives us no information as to what <u>direction</u> we should move or what transformation we should apply. Indeed, this basic knowledge tells us nothing more than the <u>legal</u> transformations that <u>can be</u> made on the expression. We know we can rewrite this equation in at least a dozen different ways, but which ones will move us closer to achieving the goal of solving the equation? For example, we could use the commutativity transformation:

$$A+B = B+A$$

which generates a host of new expressions such as:

$$\log (1+x) + \log (x-1) = 3$$

$$\text{or } \log (x+1) + \log (-1+x) = 3$$

$$\text{or } \log (x-1) + \log (x+1) = 3$$

Or we could use a transformation applying to logarithms such as:

$$\log A + \log B = \log (A B) \quad , \quad A > 0 \quad B > 0$$

$$\log (-A) + \log (-B) = \log (A B) \quad , \quad A < 0 \quad B < 0$$

which generates:

$$\log (x+1)(x-1) = 3 \quad , \quad x > 1$$

and so on.

Before proceeding, we encourage you, the reader, to generate your own solution. As you do so, try to keep track of <u>why</u> you applied a particular transformation, what kind of difficulties you experienced, and how you decided when to abandon an unsuccessful approach toward solving the equation. Now let us flip the coin from the typical problem solving process to the almost totally overlooked (in mathematics) understanding process. What follows is one of the many possible solution paths to this problem. Read through this solution and then step back and think about what it means to understand or summarize it in a way that might help someone else generate a solution to another problem.

Example 1

1. $\log(x+1) + \log(x-1) = 3$
2. $\log(x+1).(x-1) = 3$, $x > 1$
3. $\log(x^2 -1) = 3$
4. $x^2 -1 = 2^3 = 8$
5. $x^2 = 8+1 = 9$
6. $x = \sqrt{9} = \pm 3$
7. but $x > 1$, so $x = 3$ only

As we skim over the above solution, each step by itself seems to be almost obvious, but what about its overall structure? Can we scrutinize it as easily as we can the Stone Soup fable? Can we fill in the underlying motives or plans that directed the unfolding of this solution? To see that each step of this solution path indicates to the initiate a separable, distinct decision and has its own motivating piece of an overall plan for the solution, (that is, to see that there must exist some deep structure trace for solutions to math problems and that it plays a determining role in what steps were taken), compare the surface structure trace given in Example 1 with that given in Example 2 for a slightly different problem statement.

Example 2

1. $\log(x) + \log(x-2) = 3$
2. $\log(x).(x-2) = 3$, $x > 2$
3. $(x).(x-2) = 2^3 = 8$
4. $x^2 -2x = 8$
5. $x^2 -2x -8 = 0$
6. $(x+2).(x-4) = 0$
7. $x+2 = 0$ or $x-4 = 0$
8. $x=-2$ or $x=4$

9. but $x > 2$, so $x=4$ only

Although this is the same problem with $(x-1)$ substituted for the "x" of Example 1, the steps to solve the equation were different in each case. Why?

Planning Knowledge and Means-Ends Analysis

Alan Bundy (1975) has constructed an initial taxonomy and theory of the planning knowledge involved in solving a wide class of elementary equations such as this one. His theory involves two types of knowledge: first, there are planning rules for associating transformations that are applied with situations that arise in means-ends analysis; and second, there is strategic knowledge that selects the order of the application of these rules. Integrating these two types of knowledge results in a problem-solving procedure which Bundy calls the Basic Method; that is, a schema for an instance of a means-ends analysis strategy for seeking a solution. Below we give examples of such planning knowledge, cast as Bundy rules:

Isolation: Given a single occurrence of the unknown in the equation, apply a set of mathematical transformations that removes whatever functions surround this occurrence, so that it stands in isolation.

This covers any set of steps that selects the outermost function dominating the occurrence, selects an axiom which eliminates it by introducing its inverse on the right-hand side; and so on until the unknown sits by itself on the left-hand side of the equation.

Simplification: Place expressions in canonical form.

This covers adding and multiplying by zero, multiplying by one, logarithm of one, zero or one as an exponent, evaluation of terms with no unknowns, cancellation of factors across a quotient sign, etc. It is often enabled by the isolation strategy.

Collection: Given more than one occurrence of the unknown, select a transformation that reduces the number of occurrences of the unknown, thereby making the isolation strategy applicable.

This covers such steps as summing terms, adding constant exponents of products of powers of the same expression, etc.

Attraction: Given more than one occurrence of the unknown, apply a transformation that simply moves two occurrences of the unknown closer, to enable some transformation for the Collection strategy.

This covers such steps as finding common denominators for the sum of fractions, non-elementary applications of legal transformations, etc.

ERIC

Splitting: Given a complicated expression or subexpression, split it into a functional composition of some less complicated expressions, to enable the composed expressions to be treated separately.

This covers factorization, completing the square, cancellation of terms across an equals sign, etc.

Check: Given any additional relationships that must obtain, check whether they do.

This covers substitution of answers or expressions into a previous step, the extra case analysis for division by zero; indeed it includes almost any deliberately redundant processing, such as multiplying out the square that has been completed.

Basic Mathematical Knowledge

The above planning rules for mathematical problem-solving help specify which basic mathematical knowledge--the kind taught laboriously in most elementary mathematics curricula--should be applied at each step in the solution. This is the knowledge of what one may and may not do; (i.e., performing the same operation on both sides of an equation, multiplying by an expression equal to 1, adding an expression equal to zero, transposing commutative operands, distributive law for multiplication, adding exponents in multiplication... with the basic mathematical skills of algebra that make the difference between a sloppy victim of careless mistakes and a loyal upholder of the deductive laws of mathematics. Basic knowledge is not sufficient for flexible, independent mathematics--the kind of flexibility and independence derived from the ability and confidence to plan on one's own.(1)

Deep Structure Traces

The above sections have discussed some of the basic mathematical and planning knowledge needed to tie together the individual steps of a solution into a coherent deep structure, revealing the motivations and plans that lie beneath the surface of the solution path. But how

(1) In fact, perhaps one of the causes for a math student "bending the law" when he gets lost is that i) he is told he has to get from here to there, but he doesn't see how; ii) the paradigmatic math proof is usually given with unjustified leaps (i.e. referring only to the basic knowledge), but, since he didn't follow its thread when it was presented, he assumes nobody expects there to be one; iii) by not using planning knowledge, he views the process of constructing a proof as one of jumping forward from the premises and backward from the conclusions; and iv) he may as well jump from one such sequence of jumps to another whenever the expressions look sufficiently similar.

ERIC

is one to synthesize, for himself, this structure? Before exploring this issue, let us first show both a top-level summary of planning knowledge that might have been used for solving Example 1, and a more detailed example of the deep structure trace for Example 2.

For Example 1, we may briefly note that steps 4-6 reflect the successful application of the planning rule for Isolation. This was made possible by the prior application of the Collection planning rule, which in turn was enabled by the correct application of the Attraction rule.

We can see in Example 2 that steps 4-7 are motivated by the desire to split the quadratic into cases corresponding to its two roots; that this becomes possible if we could express it as a product of expressions of the form (ax+b) set equal to zero (as was done to get from steps 6 to 7); and that this would become possible if we were able to express our quadratic as $ax^2 + bx + c = 0$ and then complete the square. So the deep structure trace underlying 4-7 is shown in Figure 2.

Figure 2
Deep Structure Trace for 4-7 of Example 2.

```
Solve the quadratic equations x^2-2x=8 (line 4) by:
SPLIT (4) - into one equation for each root
      by: express as f(x).g(x) = 0
          by: ATTRACT (4) - towards desired form above
              by: SIMPLIFY (4) - into ax^2 + bx + c = 0
                      by: SPLIT: 8 --> 0 + 8
                      and: ATTRACT: move the 8 over,
                                         switching sign
                  yielding (5): x^2 -2x -8 = 0
              and: SPLIT (5) - into (ax+b).(cx+d) = 0
                      by: factoring (5)
                          yielding (6): (x+2).(x-4) = 0
      and: use the product-zero rule
              yielding (7): (x+2) = 0  or  (x-4) = 0
and solve the linear equations.
```

But what kind of reasoning/strategies did we use to synthesize this deep structure trace from steps 4 through 7 of the solution?

Constructing and Revising Hypotheses about Deep Structure Traces

Note the three patterns cited in the above deep structure trace: $f(x).g(x) = 0$, ax^2+bx+c and $(ax+b).(cx+d)$. Each of these patterns is related to a small chunk of basic algebraic knowledge that specifies what basic operations and truths can be linked to this pattern, such as the product-zero rule: $f(x).g(x) = 0 \rightarrow f(x) = 0$ or $g(x) = 0$. These patterns play a pivotal role in that they link fragments of the deep structure trace to elements in the surface structure trace via hypotheses that the understander forms. For instance, when one sees [by comparing the right-hand sides of (4) and (5)] that one side of the equation is being made zero, one might be able to apply the Product-zero rule which splits the equation into two simpler equations. This rule requires some instantiation of the pattern $f(x).g(x) = 0$ which is found when it is searched for further down the line of the proof, identifying (6) as an important clue in our reconstruction of the deep structure trace, since it confirms our hypothesis that this particular kind of split was attempted.

Our next hypothesis formation subtask is to determine how (6) is different from (4) and how and why it got that way. (We indent one level in the deep structure trace--Figure 2--since this is a subtask.) (4) is a quadratic which means it is a close relative of a three-termed polynomial in x--ax^2+bx+c. This mini-hypothesis is easily confirmed (all the terms of (4) are of the form ax2 or bx or c), but it is crucial to making sense of (6) because one of the many ways to express a quadratic is the next pivotal pattern--$(ax+b).(cx+d)$. This pattern matches (6), so we now know that the jump from (4) to (6) included placing a quadratic in the form $(ax+b).(cx+d)$. And that is another basic knowledge schema, for it tells the understander that one converts ax2 +bx +c into $(ax+b).(cx+d)$ by factoring (provided the roots exist).

Now there is a reason to look for ax2 +bx +c in the surface structure trace, which we find at excerpt element (5). Note that not every proposed surface structure element must be visible in forming a hypothesis; often a little basic knowledge is needed to see that what one is looking for is implied in what was there explicitly (e.g., a few steps were skipped, or the representation given is not in canonical form, etc.) This is what the understander will have to do if and when he explores just how the expression was factored and whether the jump from (5) to (6) was valid. But in this case he was fortunate to find $ax^2 + bx + c$ directly. The understander has only to note that the 8 came from the right hand side to get a zero, thereby splitting 8 into 0+8.

The overall process at work here for handling the understander's hypotheses is controlled by a group of strategies for taking information from the solution path that dictates possible new hypotheses, and for connecting existing hypotheses together in ways that are consistent with basic mathematical knowledge and planning knowledge for algebra. These are bottom-up and top-down activities, respectively. From this hypothesis growing and merging process, we can construct a model of how the problem was solved, what motivation lay behind selecting each step, and what the overall plan was behind the problem's solution.

Of course, the purpose of this planning knowledge is usually to help solve a problem rather than enhance one's understanding of a particular solution; however, without recourse to it, understanding the solution path is nearly impossible. This top-down/bottom-up hypothesis formation process is indeed a complex one--one that may seem more difficult than solving the problem in the first place. That this is so indicates how little experience we have in "reading" and understanding novel mathematical solutions (or proofs).

How Does This Relate to the Stone Soup Fable?

If a reader has trouble understanding the "point" of the *Stone Soup* fable, it is apt to be because he fails to perceive the existence or structure of the planning knowledge used by the global problem-solving strategies being invoked by the beggar. Less likely, but still possible, the troubled reader might never have learned the planning knowledge comprising the "social" plans and methods underlying each particular isolated action of the beggar. In story understanding, the individual schemas of planning knowledge, the plans and methods, are more apt to be recognized in a piecemeal fashion than the global hypothesis-handling strategy that weaves these schemas into a coherent model of what is really happening in the story.

However, in understanding (or generating) the solution path for solving an equation, the troubled student is apt to be completely unaware of either the existence, content, or use of the planning knowledge (in the form of, say, Bundy rules) that lies between the lines of the solution path and that provides the rationale for tying the individual steps together into a coherent plan. By being both unaware of this higher-order knowledge and of the hypothesis formation/revision strategies that use it, the student is deprived of the basic apparatus to make sense out of a solution or proof (a necessary but not a sufficient condition). He is therefore likely to believe that understanding math is a difficult and mysterious process, even after he has mastered all the basic knowledge of math, (i.e., when the transformations are applicable and how they are done). The end result is that whatever mathematical knowledge he does manage to absorb is in the form of heavily encoded procedures that he rotely memorizes and mechanically applies in order to solve special classes of problems. As he is likely to experience it, learning mathematics consists of categorizing or linking problem characteristics to rotely memorized procedures that "solve" them. Lacking the insight to see how each individual procedure naturally follows from applying some simple higher-order knowledge, he has little basis for understanding the semantics of the procedure, and therefore cannot reliably generalize or apply the procedure to slightly different problems.

To summarize the point of this section, we have seen that the cognitive process of understanding elementary mathematics does conform to our regularities: (1) There is a surface structure trace that results from the sequence of applications of mathematical transformations and laws; (2) there is a deep structure trace that recapitulates the understander's best guess as to what motivated the steps that were made; (3) there is a well-known literature of basic knowledge, consisting of axioms, notations, transformations, and the like, that serves to define the composition of the surface structure trace by saying what may follow directly from what, and to suggest all the possible ways that the surface structure trace may be extended to intermediate steps when a jump has been made; and (4) there is a body of planning knowledge, consisting in part of Bundy-style rules, concepts for types of equations given as patterns that may be matched.

ERIC

These serve to define the interrelations among trace elements (particularly those involving trace elements that don't appear in the surface structure trace, e.g., the schema for completing the square), and to suggest further bases for constructing hypotheses that would connect up with those the understander has found so far.

The fundamental coin of math understanding is the body of hypotheses about the deep structure trace. Teachers don't ever talk about them, but they reflect the mental steps that every student takes as he reads the lines of a proof and tries to piece it together. The strategies the student uses, (which threads to pursue before others, which logically-based predictions to make from the ones he accepts), as well as the different tools for handling hypotheses--how to confirm one, how to extend one, how to suggest one--are something teachable, like all "study habits," and are surely something most people could import wholesale from their deep familiarity with social attribution and planning. We think it might even be the case that some people do just that, once they grasp the planning knowledge underlying mathematics.

We think that people haven't thought of math this way before; that if they had, a teaching methodology that cites the planning knowledge explicitly and gives practice in its application would have evolved and ameliorated the mathematical illiteracy that presently offers such a stark contrast to people's familiarity with the analogously structured knowledge for social goals and attribution. Although most people find mathematics hard to understand (as compared with fables and other stories), its formal nature enables us to be substantially more precise about the planning knowledge, hypotheses formation strategies, etc., underlying the act of understanding than in the domain of general text understanding.

UNDERSTANDING ELECTRONIC CIRCUITS

In the last two sections we illustrated some of the important theoretical constructs and processes involved in understanding some event, story, or mathematical solution, etc., while stressing the surprisingly invariant nature of these concepts and processes over two radically diverse domains. It might seem to be belaboring this point by delving into yet a third knowledge area--understanding electronic circuits. However, it was from witnessing student technicians struggling and failing to "understand" a novel circuit that we first began to wonder what higher-order knowledge--knowledge besides basic electronic laws and concepts -- were actually needed to enable a technician to understand a new circuit well enough to troubleshoot it on his own. As we began to explore this issue by explicitly representing the tacit knowledge that a skilled troubleshooter uses in "comprehending" a new circuit schematic and then analyzing the protocols of both expert and student technicians using this knowledge, we discovered the strong similarity between this activity and that of story comprehension. In fact, comprehending a circuit schematic is a slow and conscious effort, with eye fixations complementing verbal protocols. Thus we had an unparalleled experimental setting for probing the understanding process. After discovering the strong correspondence between these two diverse domains of knowledge--story understanding and circuit schematic understanding--we questioned: (1) if other domains, equally diverse, would support this correspondence (and hence we began to examine the process of understanding mathematical solutions) and (2) if these "comprehension" skills were sufficiently domain-independent to enable us to find ways to teach them to technicians using the more intuitively understandable domain of, say, stories, and then transferring them to the domain of electronic troubleshooting (admittedly a bizarre idea).

Just as with student technicians, most of us will find the jargon and technical underpinnings of electronics rather unnatural. Therefore, in the remainder of this section we shall lapse into technical details only when absolutely necessary, and focus our attention primarily on the relationships between this domain and story understanding.

Surface Structure and Deep Structure Trace

In story understanding the basic elements of the surface structure were easy to identify since an element of the surface structure was basically a line or group of lines in the text. Identifying the basic elements in the circuit schematic involves segmenting the two-dimensional diagram into its primitive functional constituents (e.g., a transistor with its biasing network). Sometimes this segmentation is explicitly indicated with functional block diagrams superimposed on top of the schematic.

A circuits' deep structure trace, which is the result of the understanding process, captures the underlying teleology or causal mechanisms of the circuit. It should contain the information necessary to explain how the circuit works and why it works as it does, with each component of the schematic (or constituent of the surface structure) playing some role in the purposeful design of the circuit. Initially, one would expect the deep structure trace of a fable, for example, to have little in common with that of an electronic device. However, such is not the case. One of the key conceptual processes used in "reading between the lines" of a story consists of the skilful application of social attribution theory--a theory of social plans, motives, intentions--for providing the grist for filling in the plot of the story.

We have begun to appreciate that schematic understanding has its own <u>attribution</u> theory. The mental glue used for cementing the constituents of a circuit schematic are the designer's <u>plans</u>. Constructing an understanding of a circuit schematic requires one to realize a sequence of plans and sub-plans where fulfilling each piece of a higher-order plan generates a sub-plan. Therefore, understanding a novel schematic involves recapitulating, to a limited degree, the problem-solving activity that hypothetically went into designing it. Each function block or component becomes associated with a piece of a plan which, in turn, is a piece of a higher-order plan, continuing up the planning tree until a top-level plan is reached. This plan accounts for all of the components in the circuit--much like the moral explicates the fable. Understanding schematics, therefore, requires access to both the planning knowledge and the problem-solving strategies that expand and refine these plans, just as understanding stories requires access to, for example, what is involved in a CON.

Planning Knowledge

In the last several years there has been a flurry of activity in the discovery, the representation and use of Plans in circuit design, circuit understanding, and teaching (A. Brown (76), Sussman (73), DeKleer (77), Rich and Shrobe (76), Goldstein (74,76), J.S. Brown (76)). A detailed discussion of this knowledge is beyond this paper, but to give the reader some idea of its scope, we will illustrate some of the planning knowledge underlying one class of circuits -- regulated power supplies. For our purpose here, this planning knowledge is meant only to facilitate understanding circuits, as opposed to designing them from scratch, and therefore there is little need for extensive mathematical detail. What is more important here are those aspects of the planning knowledge that provide guidance in uncovering which particular plan underlies a given circuit (such as the knowledge about a CON that helps us <u>recognize</u> a variant of a CON as opposed to performing a CON).

An active regulated power supply is most likely to be constructed from one of three Top-Level Plans Types:

1. Series-Regulated-Plan
2. Shunt Regulated Plan
3. Switching Regulated Plan

Each one specifies a connected set of circuit plan "elements"; recursively each element can be constructed from one of a set of sub-plan types.

In Figure 3a we present a diagram of the set of connected elements in the Series-Regulated Plan. The top-level plan is, by definition, abstract. It specifies the top-level functional elements, their interrelationships, and the various constraints that each element must meet relative to the design goals of the top-level plan. Since there are many ways to realize each of these elements, the plan at this level of abstraction covers a large variety of series-regulated power supplies. An actual circuit appears only when each of the top-level functional elements is expanded according to a repertoire of lower level plans for realizing that element (see Figure 3b).

Plans at any level of abstraction are multi-faceted specifications embodying several other kinds of knowledge. These can

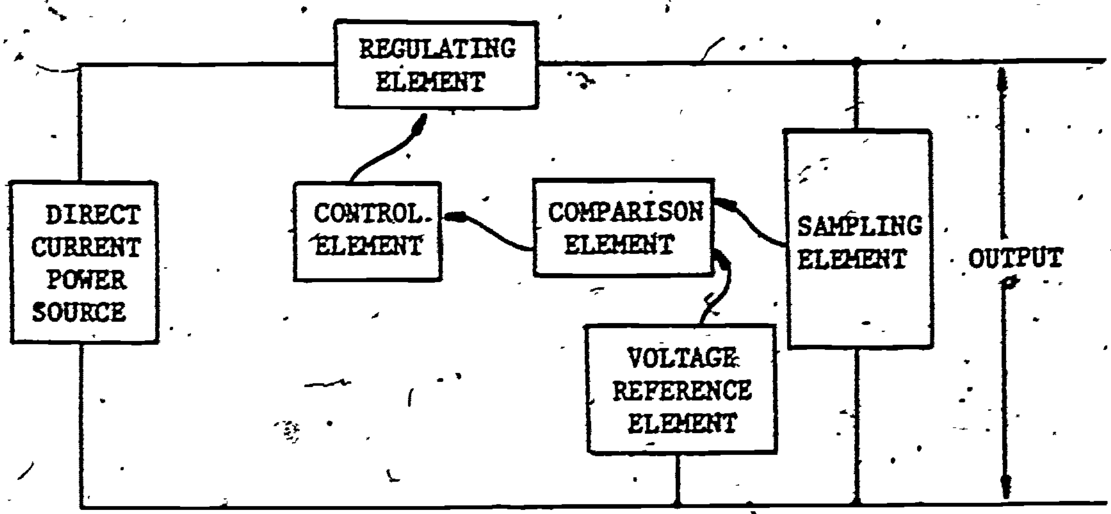

Figure 3a. Top Level Series-Regulated Plan

ERIC

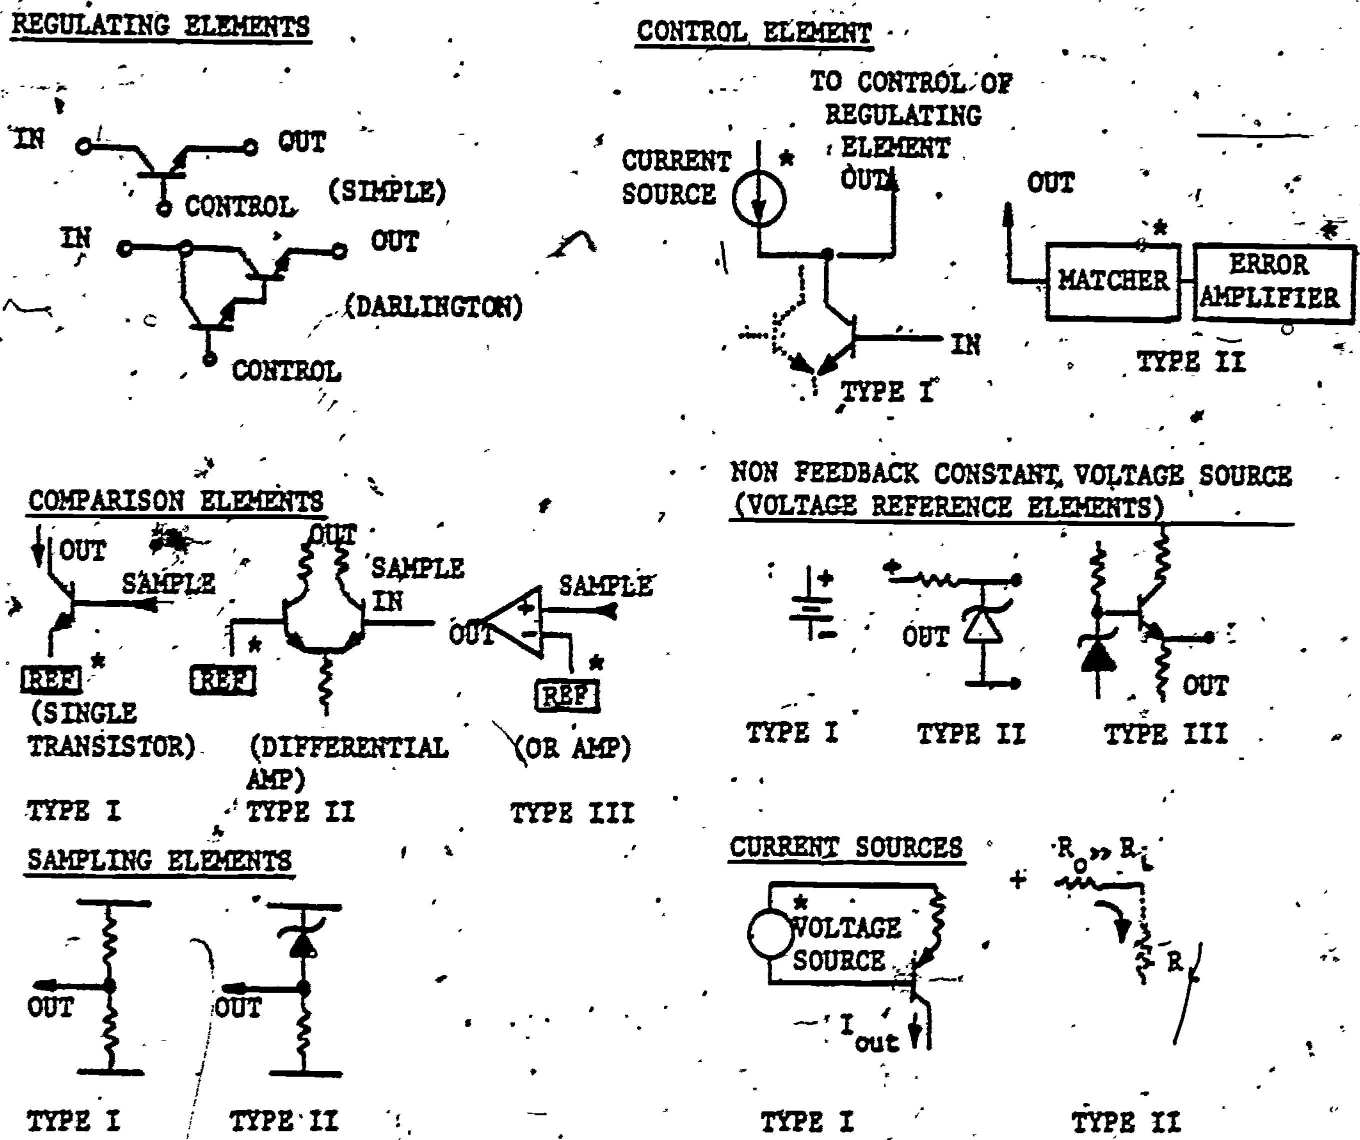

*These, in turn, are instantiated by still lower level plans.

Figure 3b. Lower Level Subplans

Various possible expansions for each of the functional elements of the top level plan. Only the circuit form is shown here; the annotations are omitted for simplicity.

Stereotype Form:

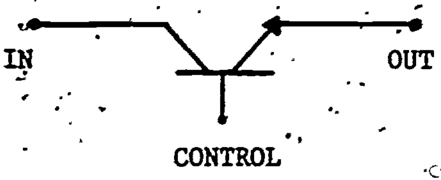

Input/Output: The current through the in-out line changes as a function of the control input.

Viewpoints:

1. May be seen as electronically controlled variable resistance forming, together with the load, a voltage divider across the power source.
2. By including the load resistance, may be seen as an emitter follower as shown below:

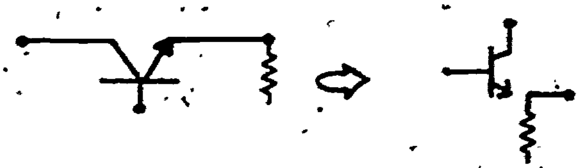

Recognition Features (for parsing a schematic): Transistor is in series with the load in a closed path across the power source. There are no other significant impedances in this path.

Commentary: Transistor must be operating in its active region.

Typical Faults and their Manifestations: Example: Control terminal open would cause the current referenced in "the I/O behavior" slot to be independent of the Control. (Note that the global symptoms of the fault are then determined by "lifting" the altered I/O behavior up through the teleology of the higher-order plans.)

Knowledge and metaphors for Understanding the Teleology of the Plan: The regulating element has an input driven by the power source and an output delivering current to the load. The control input mediates power flow similar to how a valve mediates flow in a pipe.

Boundary Conditions: Surrounding circuits must provide sufficient current to control input and maintain emitter base junction forward-biased and collector base junction reverse-biased. There is a lower bound on output current below which the element ceases to operate.

Teleology: [Basically none since this plan has only one component -- unless we discuss the junctions in the transistor.] Ordinarily, teleology would describe how the elements of the plan function together so as to achieve the "I/O behavior". For example, in the plan scheme of the Series-Regulated Plan the teleology would specify <u>how</u> the elements function together as a feedback control system to achieve the goals of the "I/O" slot whereas the "Knowledge to Understand..." slot contains the <u>conceptual</u> knowledge about feedback.

Fig. 4. A simplified example of the kinds of knowledge in the Regulating (sub) Plan Schema of the <u>Regulating Element</u> contained in the Series-Regulated Plan.

be brought together to form a plan schema. The kinds of knowledge in a plan schema are illustrated in a sub-plan schema for the "regulating element" of the Series-Regulated plan shown in Figure 4. If other, alternative realizations of this element existed, then each would also have a corresponding plan schema. Of course, these plans may consist of functional descriptions that require a still lower level expansion before an actual series regulated power supply becomes fully specified.

According to our theory, understanding a circuit schematic involves using this planning knowledge to propose a sequence of design (problem solving) steps that will eventually culminate in the given schematic. This planning knowledge, which is so tightly structured that it could even be viewed as a planning grammar(2) (much like a story grammar), captures the set of abstract plans and methods that could be used to construct (up to some level of detail) any one of a potentially infinite number of circuits pertaining to some generic class of electronic devices. The challenge of understanding a particular circuit schematic involves discovering a sequence of plans (and sub-plans, ad infinitum) that will eventually account for the way that each surface structure fragment becomes an integral part of the overall plan.

Without knowing this planning grammar for the generic device being examined, the process of understanding a schematic is as difficult as understanding a fable from a foreign culture. By knowing this planning grammar, the understanding process becomes one of examining the schematic in a bottom-up way, isolating fragments of the schematic and guessing what part of a lower level plan it might match. This bottom-up process constantly interacts with the top-down process for conjecturing the nature of the high level plan. The process is complete and the circuit understood when the two "meet," accounting for all the components in the schematic.

Hypothesis Formation and Revision

Strategies for facilitating this comprehension process not only concern how to apply the higher-order knowledge in the form of plans

(2) A concept originally used by Goldstein (1976) to formalize basic problem-solving methods as augmented transition networks.

but also how to coordinate and allocate processing resources between top-down hypothesizing about a possible global plan and bottom-up processing of the data contained in the schematic. Understanding how to coordinate these two approaches is critical, since it is often difficult to know how to interpret a fragment of the schematic without the advantage of using a conjecture about how to view it which finally stems from some top-level plan. The person trying to understand a circuit must often be willing to make educated guesses about how some fragment of the circuit might be functioning in terms of some high order plan, and then attempt to either verify or reject that guess.

An "Understanding" Scenario

Rather than provide a theoretical description of the hypothesis formation and revision process, we have included below an annotated protocol of a subject, having access to the planning knowledge, describing his process of understanding a particular voltage regulated power supply. The protocol has been described in a way that (hopefully) the casual reader can skim, gleaning the flavor of the process to sufficient depth so as to be able to perceive its relationship to the understanding process for fables, etc.

Event 1. An initial scan is made of the schematic (see Figure 5) and immediately the pair of transistors Q3,Q4 leaps out as an instance of the Darlington plan. (The Darlington transistor pair is such a common device that it's not unreasonable for an electronics technician to be able to pick it out nearly instantly.) This leads to the conjecture that this pair of transistors is an instance of a Darlington schema which functions as the Regulating Element in the Series-Regulated plan for Feedback/Regulated power supplies.

> This conjecture follows from two facts: the first is that we know this circuit is some kind of regulated power supply and the second is that the only top-level plan of the three (i.e. SRP SP SWP) which naturally uses a Darlington sub-plan as an element is the Series-Regulated Plan (SRP). Additional support for this conjecture comes from the fact that the Darlington pair lies along a path in series with the load -- a clue sought for in the recognition knowledge part of the plan schema; as well as satisfying the topological constraints imposed by the Series-Regulated plan.

Event 2. Continuing to scan the schematic, zener CR4 is detected in series with the resistor R10. This grouping satisfies one of the

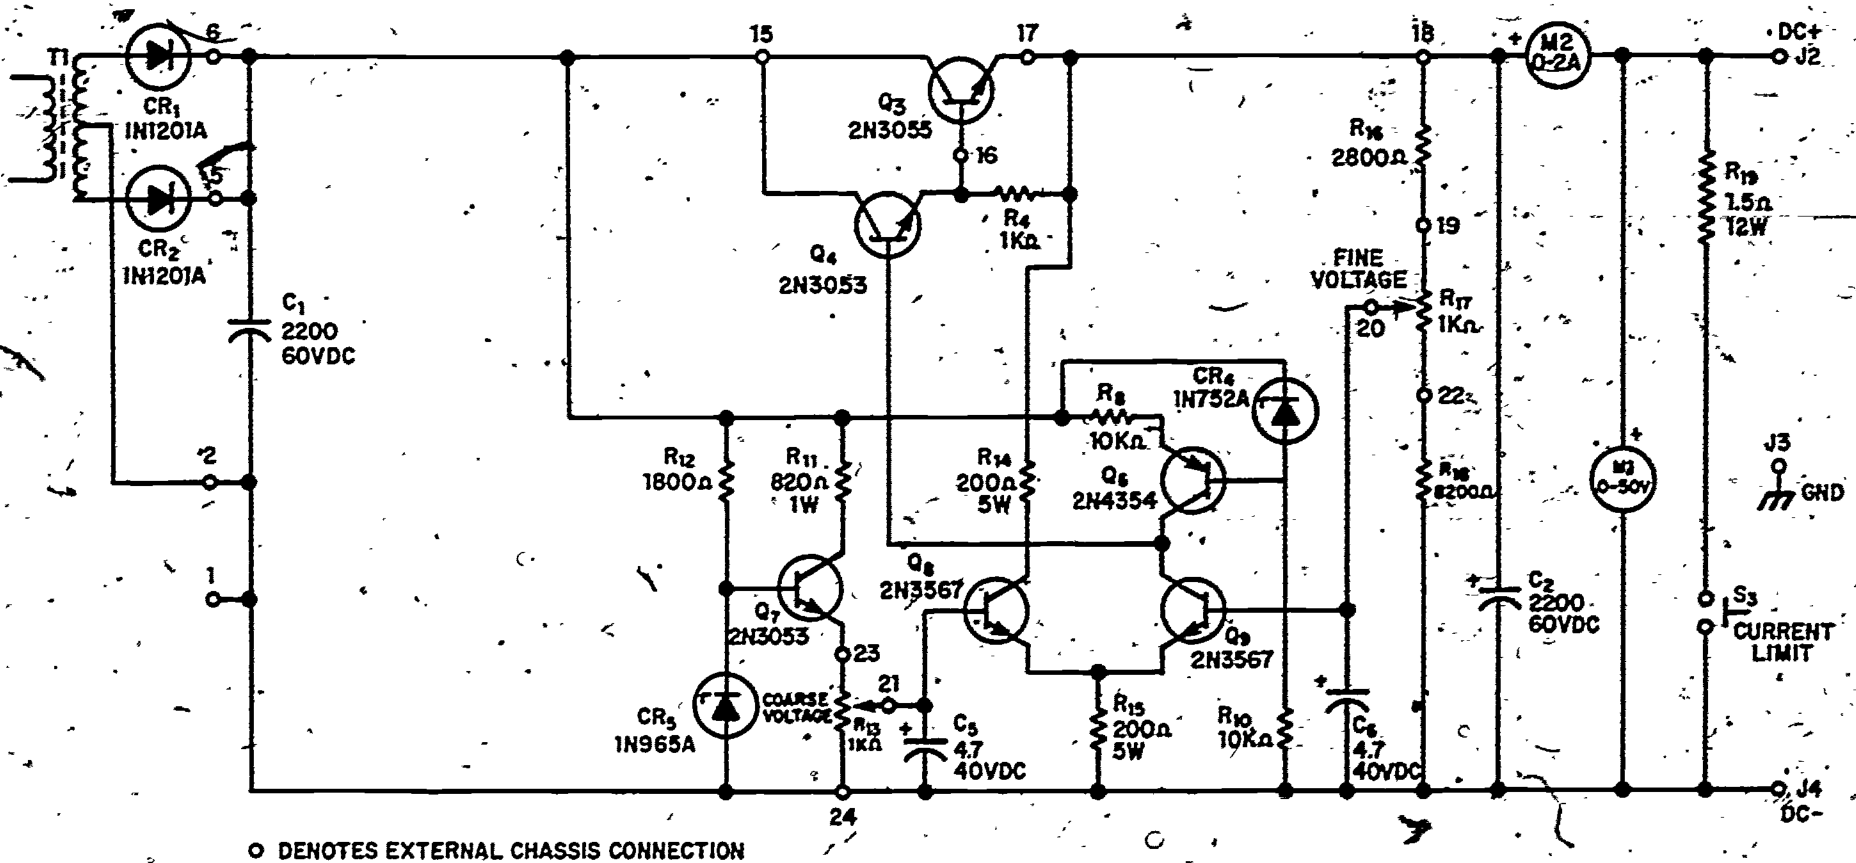

Fig. 5. Circuit Schematic for a regulated power supply

intermediate level plans for a non-feedback Constant Voltage Source and is therefore conjectured to be the Voltage Reference Element under the hypothesized Series-Regulated plan.

> Note how the initial hypothesis about the top-level plan is beginning to affect how a low level element is interpreted.

<u>Event 3</u>. Next, the pair of transistors Q8;Q9 are superficially examined and guessed to be the kernel of the plan for a differential amplifier.

> This low level conjecture seems reasonable since the Series-Regulated plan calls for a Comparing Element which can be realized by a Differential Amplifier plan.

<u>Event 4</u>. Believing this, the bank of resistors R16,R17,R18 is guessed to be an instance of a Voltage Divider plan which serves as the Sampling-Element in the Series-Regulated plan.

> This again seems reasonable except for the fine-voltage control which is not expected as a component in the Sampling-Element. But this <u>objection</u> to a piece of contradictory evidence is temporarily ignored, perhaps because there is a coarse voltage adjusting element which is not connected to the fine control in an obvious way (i.e., no known plan-schemas account for this).

<u>Event 5</u>. At this point all active components (e.g., transistors, diodes) have been accounted for in the schematic except for zener CR5 and transistors Q6,Q7. Hence, there could be something amiss. There is only one element of the Series-Regulator plan that is still unfulfilled, namely the Control Element, and since transistors Q6 and Q7 don't appear to be topologically close, it seems doubtful that they can be made to instantiate any of the potential Control Element plans.

> Note the use of heuristic knowledge about topology to accrue more evidence that something might be wrong with the current deep structure trace-hypothesis.

<u>Event 6</u>. This causes a re-examination of what has been accounted for thus far by the current hypothesis (which is a prelude to a hypothesis revision step). It seems that interpreting CR4,R10 as the Voltage Reference Element cannot possibly be correct since it doesn't feed into the Comparing Element as dictated by the top-level Series-Regulated plan. Further examination reveals an even more important clash: under the above interpretation, one side of the Differential Amplifier plan has no input and the other side has two contradicting inputs.

Enough evidence has certainly been accrued to call for a revision of the current hypothesis but should the whole hypothesis be abandoned and if not, what parts of it can be saved and the remainder intelligently revised?

Event 7. Feeling confident that the conjecture about the role of Q8, (he feels he can save this part), Q9 is correct, a decision is made to reconsider the two inputs of the Comparing Element. (Note that he determines where the two inputs should be from the Differential Amplifier plan.) There is little doubt that the low level conjecture about instantiating the Sampling Element with R16,R17,R18 is correct since this string of resistors is such a usual realization of that element.

Event 8. A match is attempted of the unrecognized active devices topologically connected to Q8. In fact, CR5,Q7 trivially matches a low-level plan for constant voltage sources which only leaves Q6,CR4 unexplained.

This process combines both a bottom-up data-driven grouping with a local top-down hypothesis expectation.

Event 9. Hopefully, these remaining devices will satisfy one of the Control Element plans. Since the hypothesized output of the Differential Amplifier is directly connected to the input of the Regulating Element, that rules out viewing the Control Element as a Matcher (one of the possible plans for the Control Element). Hence this leads to viewing it as a Constant Current Source (i.e. the other known plan for the Control Element).

If this doesn't work then another major revision is called for. But now after he concentrated all his processing resources on this goal, it becomes clear how to match these remaining components to one of several possible plans for a Constant Current Source.

Event 10. Now all the active components have been grouped together and consistently interpreted as elements in sub-plans within the context of the overall Series-Regulator plan.

The resulting deep structure trace/hypothesis can now tie together all the knowledge associated with each Plan schema yielding a teleological model (structured by the top-level plan) of how the circuit works and how to troubleshoot it. For example now that devices (CR5,Q7) have been successfully accounted for as instantiating one of the Constant Current Source plans, the role or purpose of CR5 can be determined from additional knowledge in the given plan.

This scenario captures the essence of how one person made sense of a novel schematic. It does not describe very much of the problem-solving effort that went into fulfilling each plan in terms of satisfying any constraints required by the laws of electronics. Rather it focussed on satisfying topological constraints dictated by the plans themselves. In part, this was to show how this higher-order planning knowledge can, in fact, be useful to technicians (who don't have the electronic theory needed by circuit designers) and, in part, it was to show that the problem solving required to handle these issues goes beyond our means-ends analysis scheme and involves a collection of more sophisticated problem solving strategies.

In concluding this section, it might be of interest to note that the above-mentioned understanding process involving a hypothetical recapitulation of a sequence of design/plan steps has been used as the primary explanatory methodology for teaching student technicians why a given piece of equipment works, i.e., what its underlying teleological model is (Brown et al., 1976). In this scenario, we first present a simplified model/design of the circuit, examine why this simplified circuit fails to perform satisfactorily, and then examine how that failure might be patched or modified and so on until this hypothetical sequence of design patches finally yields the given circuit.(3) In this way, the student understands what each component's role is, either in terms of its role in the simplified circuit model, or as a patch around some understood shortcoming of that model.

(3) A pedagogical idea inspired by Sussman's research in electronics (A. Brown & Sussman, 1974).

SUMMARY OF THEORETICAL CONCEPTS IN COMMON OVER THE THREE DOMAINS

In the last three sections we have examined what kinds of knowledge and strategies are used in the understanding process over three diverse domains. Although each domain has its own idiosyncratic and domain specific knowledge, there is a fair amount of invariance over the domains. In this section we summarize the underlying concepts in this theory of "understanding."

The surface structure trace is an information structure that spells out what actually happens in the story, solution path, etc. It is a sequence of the reported or described elements of the behavior to be understood. Since one can never describe everything about a behavior, there will always be gaps in the information that the surface structure trace provides; one of the measures of how thoroughly the behavior has been understood will be the ability to fill in these gaps.

The deep structure trace is an information structure that spells out the decisions that were made and that resulted in the particular behavior. It is composed recursively out of:

Goals--desired situations, usually described in the same terms as the behavior. A goal always occurs in a deep structure trace in contrast with another actual situation. This contrast is factored into differences that the decision maker hopes to reduce. Hence, the deep structure trace also contains:

Deltacts--reducible difference categories. The reason a difference is abstracted as a category and given a deltact name is that the understander knows some:

Methods--how various things may be achieved; the means to an end. Methods are attached to one or more Deltacts, ("To get less hungry [deltact] try eating [method].") Methods are where the recursion comes in. A method may consist of reducing certain differences or adopting certain other goals, as well as some fully specified behavior. By having methods that use goals and deltacts, a wide range of possible behaviors may be regarded as pursuing a particular method.

The deep structure trace is an explanation of the surface structure trace; one of the measures of how thoroughly it has been understood will be the plausibility and completeness of this explanation. It is an information structure, and the processes that produce it may be quite different in form from the problem solving process it traces. We believe that understanding proceeds by assembling hypotheses about the two traces of behavior and of problem solving.

We use the term means-ends analysis target structure for a familiar pattern of deep structure trace elements used in building up hypotheses. "Means-ends analysis" describes the purposive aspect of the deep structure trace. The target structure guides the construction of the deep structure trace, filling in for certain parts of the pattern with known behavior from the surface structure trace or with other (presumably confirmed) hypotheses. (The part that is filled in is called a "slot.") Thus, the means-ends analysis target structure yields a basic hypothesis about the deep structure trace that may be revised in light of other details.

When revising a hypothesized deep structure trace, there are constraints on what may be changed, and on what must be changed, in addition to deciding which kind of change should be made. We will use the term planning knowledge for the rules that specify how deep structure trace elements may be combined. The planning knowledge defines the target structure for: (1) understanding the behavior as a unified whole; (2) the plausibility of the deep structure trace; and (3) the habits as to which combinations to try and which basic hypotheses to suggest.

The understander must have a feel for the problem solving processes that will explain or generate the behavior; that is, he must know those processes involved in finding possible solutions in order to solve the problem by himself. In effect the understander is being asked to "catch up" with the plans and motives underlying some behavior that has already happened. Thus he must have a mastery of the hypothesis formation processes that are available. These are his tools for getting from behavior to explanation. They enable him to know the range of other possible behaviors and plans; to fill in choices in the deep structure trace that appear to have been glossed over; to understand and profit from conventions or restrictions in the planful behavior for a given domain; to select major deep structure elements to dominate the hypothesis; and to incorporate details into a hypothesis (either by substitution or by composition), etc.

Since these hypothesis formation tools are incomplete and imperfect (to say nothing of the information they may be given to work on), it is equally important that the understander have a mastery of the hypothesis selection process, since there will always be inconsistent, alternative hypotheses to choose from. Some of the devices and criteria for making such a choice are: (1) the integrity, wholeness, or appropriateness of a plan; (2) its formal plausibility; (3) its consistency with the situations contrasted in the various goal elements; (4) the ease with which the planning knowledge (grammar) can splice it into larger, accepted structures; (5) the presence of confirming behavior for a plan; (6) predictions and consequences for further behavior; and (7) whether or not two hypothesis elements can be interchanged or combined.

An incorrect hypothesis can be salvaged: It may have been almost right, or the detailed understanding of most of its evidence may have been correct. For this reason, we have introduced the notion of revising the hypothesis to conform to the evidence. Revising a hypothesis consists of focussing criticism and responding with proposed improvements.

The hypothesis formation and elaboration process is dominated in a "top-down" way by a target structure or model, such as is given by means-ends analysis, the planning knowledge, the hypothesis manipulation procedures, and the basic deductive strategies taken together. The hypothesis proposal, confirmation, and revision processes are necessarily driven onward by the actual evidence available--the elements of the two kinds of trace--in a "bottom-up" way, for the fundamental direction of information flow is from behavior narrative to complete explanation. These two, one pulling, one pushing, strike a balance in the understanding of purposive behavior that we call top-down/bottom-up processing.

Herein lies the beauty of the model notion and its potential relevance to learning strategies: Without guessing in advance what the eventual explanation is going to be, the understander can use each new piece of evidence to drive his deduction forward while using the largely intensional target structure provided by the model to focus his efforts toward the most viable explanation. Given the right distribution of knowledge between the representational framework and the principal target structure representation, the basic strategies can be expressed in a domain-independent way.

FORMULATIONS AND DELIVERY OF SOME NEW LEARNING STRATEGIES

Thus far we have examined the understanding process over three diverse domains, describing the processes, strategies and conceptual structures for each, and the invariances over these domains. Within each domain we have focussed on the correspondence between the knowledge needed for problem solving and the knowledge needed for understanding. Without explicit awareness of the largely tacit planning and strategic knowledge inherent in each domain, it is difficult for a person to 'make sense of' many sequences of behavior as described by a story, a set of instructions, a problem solution, a complex system, etc. Our premise was that before one could begin to formulate new learning strategies for enhancing a student's abilities to acquire an understanding of some new piece of knowledge (as opposed to just rotely memorizing it), these processes and tacit knowledge had to be made more explicit. Having partially accomplished this, the question naturally arises as to what impact this has on the formulation and teaching of learning strategies. We suggest that the above theory be used to make as explicit as possible how "understanding" is an active process requiring the understander to synthesize, verify, and refine a deep structure trace or hypothesis about the underlying motives, plans, and intentions that fit each separate piece of the "puzzle" into a coherent structure. Teaching this process can probably best be accomplished by focussing on the domain of knowledge the student is to specialize in. The teacher should articulate for that domain the higher-order planning knowledge and the strategic knowledge for formulating and revising hypotheses about what something means. By carefully choosing a set of situations for the student to understand, each strategic rule can be instantiated, providing him with practice in the coordination of the top-down, bottom-up hypothesis formation and revision process. Some situations might be devised to be inherently "garden path" where the student's most likely first guess of the underlying meaning is apt to be wrong, requiring him to focus on how he detects that his guess is wrong and how he then intelligently goes about revising it.

Since this hypothesis formation/revision process is so complex, it might be useful to construct a hypothetical understander in a film animation, who shows the process in an expert's head (in slow motion) as he goes about understanding some novel situation. At the very least, this will suggest to the student that understanding is not a simple process, but rather a complex and very active one.

After a student has begun to master strategies for constructing and revising deep structure traces over the given knowledge domain, his attention could be drawn to this same process as it applies to story comprehension. In this way, he could begin to witness the generality of what he has been taught, especially since the planning knowledge needed in story comprehension is usually well understood (albeit tacitly), as are the rudimentary strategies and processes of weaving together the lines of a story into a coherent explanatory structure.

There is one new kind of instructional technology we are developing that might provide a unique capability for exposing students to the underlying problem-solving strategies and knowledge for a domain in a way that is apt to be enticing and meaningful. Recently we have been designing an "Articulate Expert" instructional computer system that explicitly contains all of the planning knowledge, basic knowledge, means-ends problem-solving strategies, as well as a limited class of hypothesis revision (debugging) strategies necessary for solving on its own a wide class of student-generated problems. The expert's articulateness is especially significant: not only can it solve a problem, but it can also explain (at various levels of detail) <u>why</u> it performed each step. It can explain its overall plan of attack, how it formulated that plan, and why it did not do it some other way.

In other words, the student can pose a problem to this system and witness all the inner thinking, mistakes and false attempts that an expert makes, thereby exposing the student to strategies and knowledge sources that are hidden by looking only at the final solution to a problem. We believe that by letting the student pose his own problems to the Articulate Expert and having him witness the unfolding of the plans of a problem solver, he is on his way to appreciating what he must fill in when he tries to make sense of a problem solution.

ERIC

Rumelhart, D.E. Understanding and summarizing brief stories. In D. LaBerge and S.J. Samuels (Eds.), Basic processes in reading: Perception and comprehension. Hillsdale, N.J.: Lawrence Erlbaum Associates, 1977.

Rumelhart, D. E. and Ortony, A. Representation of knowledge. In R. C. Anderson, R. J. Spiro, and W. E. Montague (Eds.), Schooling and the acquisition of knowledge. Hillsdale, N. J.: Lawrence Erlbaum Associates, 1977.

Schank, R. and Abelson, R. Scripts, Plans, Goals, and Understanding. Hillsdale, N.J.: Lawrence Erlbaum Associates, 1977.

Sussman, G.J. A computational model of skill acquisition. Artificial Intelligence Laboratory, TR-297, Cambridge: M.I.T., 1973.

Winograd, T. Frame representations and the declarative-procedural controversy. In D. G. Bobrow and A. M. Collins (Eds.), Representation & understanding: Studies in cognitive science. New York: Academic Press, 1975.

Winston, P.H. Artificial Intelligence. Reading, Mass.: Addison-Wesley, 1977.

Rumelhart, D.E. Understanding and summarizing brief stories. In D. LaBerge and S.J. Samuels (Eds.), Basic processes in reading: Perception and comprehension. Hillsdale, N.J.: Lawrence Erlbaum Associates, 1977.

Rumelhart, D. E. and Ortony, A. Representation of knowledge. In R. C. Anderson, R. J. Spiro, and W. E. Montague (Eds.), Schooling and the acquisition of knowledge. Hillsdale, N. J.: Lawrence Erlbaum Associates, 1977.

Schank, R. and Abelson, R. Scripts, Plans, Goals, and Understanding. Hillsdale, N.J.: Lawrence Erlbaum Associates, 1977.

Sussman, G.J. A computational model of skill acquisition. Artificial Intelligence Laboratory, TR-297, Cambridge: M.I.T., 1973.

Winograd, T. Frame representations and the declarative-procedural controversy. In D. G. Bobrow and A. M. Collins (Eds.), Representation & understanding: Studies in cognitive science. New York: Academic Press, 1975.

Winston, P.H. Artificial Intelligence. Reading, Mass.: Addison-Wesley, 1977.

51