ARI Research Note 87-71

PROBLEM SOLVING IN A NATURAL TASK AS A FUNCTION OF EXPERIENCE

Juliana S. Lancaster and Janet L. Kolodner
Georgia Institute of Technology

for

Contracting Officer's Representative
Judith Orasanu

BASIC RESEARCH LABORATORY
Michael Kaplan, Director

U. S. Army
Research Institute for the Behavioral and Social Sciences

December 1987

U. S. ARMY RESEARCH INSTITUTE

FOR THE BEHAVIORAL AND SOCIAL SCIENCES

A Field Operating Agency under the Jurisdiction of the

Deputy Chief of Staff for Personnel

EDGAR M. JOHNSON
Technical Director

WM. DARRYL HENDERSON
COL, IN
Commanding

Research accomplished under contract
for the Department of the Army

Georgia Institute of Technology

Technical review by

Dan Ragland

A-1

COPY INSPECTED 6

REPORT DOCUMENTATION PAGE

READ INSTRUCTIONS BEFORE COMPLETING FORM

1. REPORT NUMBER: ARI Research Note 87-71	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle): Problem Solving in a Natural Task as a Function of Experience		5. TYPE OF REPORT & PERIOD COVERED: Interim Report July 86 - July 87
		6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s): Juliana S. Lancaster and Janet L. Kolodner		8. CONTRACT OR GRANT NUMBER(s): MDA903-86-C-0173
9. PERFORMING ORGANIZATION NAME AND ADDRESS: School of Information and Computer Science, Georgia Institute of Technology, Atlanta, GA 30332		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS: 2Q161102B74F
11. CONTROLLING OFFICE NAME AND ADDRESS: U.S. Army Research Institute for the Behavioral and Social Sciences, 5001 Eisenhower Avenue, Alexandria, VA 22333-5600		12. REPORT DATE: December 1987
		13. NUMBER OF PAGES: 12
14. MONITORING AGENCY NAME & ADDRESS(if different from Controlling Office): - -		15. SECURITY CLASS. (of this report): Unclassified
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE: n/a

16. DISTRIBUTION STATEMENT (of this Report)

Approved for public release; distribution unlimited.

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

- -

18. SUPPLEMENTARY NOTES

Judith Orasanu, contracting officer's representative

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

Cognitive Psychology
Problem Solving
Expertise
Experience

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

This research note investigates the effects of experience on problem solving behavior and the knowledge base of workers in an applied setting -- that of automobile mechanics.

The automobile is a highly complex system with many interconnected subsystems. Problem descriptions presented to a mechanic who needs to diagnose a car are usually quite sketchy, however. Novices are less able than experts to diagnose any but the most obvious problems. This research note concerns itself with

(OVER)

ARI RESEARCH NOTE 87-71 (Continued)

20. Abstract (Continued)

identifying the qualitative differences between mechanics with different levels of expertise. In this research note, three student mechanics are observed in a post-secondary technical school, each at a different level of expertise, diagnosing six problems introduced into cars in the school. Collected protocols are then analyzed to find the knowledge and strategies used in solving each problem. Series of protocols for each student were also analyzed to find the changes in knowledge and strategies used in solving later problems as compared to earlier problems. Differences were seen in both the knowledge used by the subjects and in their general approach to diagnosis. As a result of experience, the student mechanics seemed to improve in three areas: 1) their knowledge of the relationships between symptoms and possible failures was augmented, 2) their causal models of the car's systems were augmented, and 3) their general troubleshooting procedures and decision rules were much improved.

Problem Solving in a Natural Task as a Function of Experience*

Juliana S. Lancaster
Janet L. Kolodner

School of Information and Computer Science
Georgia Institute of Technology
Atlanta, GA 30332

Abstract

Problem solving is known to vary in some predictable ways as a function of experience. In this study, we have investigated the effects of experience on the problem solving behavior and knowledge base of workers in an applied setting: automobile mechanics. The automobile itself is a highly complex system with many interconnected subsystems. Problem descriptions (i.e., symptoms) presented to a mechanic who needs to diagnose a car, however, are usually quite sketchy, requiring the collection of more information before solution. Novices are less able than experts to diagnose any but the obvious problems, and we are interested in identifying the qualitative differences between mechanics at different levels of expertise. In the study reported, we observed three student mechanics in a post-secondary technical school, each at a different level of expertise, diagnose six problems introduced into cars in the school. We then analyzed the protocols we collected to find the knowledge and strategies used in solving each problem. We also analyzed the series of protocols for each student to find the changes in knowledge and strategies used in solving later problems as compared to earlier problems. Differences were seen in both the knowledge used by the subjects and in their general approach to diagnosis. As a result of experience, the student mechanics seemed to improve in three areas: (1) their knowledge of the relationships between symptoms and possible failures was augmented, (2) their causal models of the car's systems were augmented, and (3) their general troubleshooting procedures and decision rules were much improved.

1. Introduction

Problem solving is known to vary in some predictable ways as a function of expertise. When the process of problem solving first came under scrutiny by psychology and computer science researchers, the problems studied were in knowledge-lean domains in which well-defined situations have known solutions (Reed, Ernst, & Banerji, 1974; Reed & Johnson, 1977; Reitman, 1976; Simon, 1975). In that work, the behavior of interest was generally a variable such as number of steps to completion or number of correct solutions. Recently however, interest in problem solving has leaned more toward problems in knowledge-rich domains such as physics (Chi, Glaser,& Rees, 1982; Simon & Simon, 1978), thermodynamics (Bhaskar & Simon, 1977), architecture (Akin, 1980), and political science (Voss, Greene, Post, & Penner, 1983; Voss & Tyler, 1981). Within these domains, researchers have continued to look at the steps and plans generated in coming to a solution, but they have also developed a further interest in the nature or organization of the knowledge used in the process of problem solving. A major question regarding the nature or organization of knowledge has been how that knowledge and its changes influence performance.

Our knowledge of the differences between novices and experts has reached the point where several general statements can be made. First, experts in any field are more able to recognize and remember typical conditions within their area of expertise. Second, experts generally organize their knowledge by functional characteristics of problems while novices are more likely to use surface features to characterize problems.** There have not been a lot of explicit conclusions, however, about the particular knowledge structures used by experts and novices. Nor has there been work describing the particular changes in knowledge and processing behavior that happen as a result of a single experience.

---

* This research is supported in part by the Army Research Institute for the Behavioral and Social Sciences under Contract No. MDA-903-86-C-173. Thanks to Ken Allison and Gita Rangarajan, who provided representations for the paper and ideas about analyzing the protocols.

***See Chi, et al (1982) and Glaser (1985) for more discussion of novice/expert differences.*

Our primary goal is to discover the changes that individual experiences have on a problem solver. In order to achieve that goal, we first have to find out what knowledge the problem solver starts with before solving any problem and what knowledge he has later to solve a similar problem. While earlier work has indicated that "good" diagnostic ability is a function more of knowledge about the problem area being diagnosed than of general diagnostic skills (Miller, 1975), we find that the diagnostic skills of novices and experts also differ, and therefore also observe initial strategies of problem solvers and those used after a particular experience.

In the particular experiment to be discussed, we had two goals. Our first was to find out what knowledge subjects at different levels of expertise had and to be able to state the problem solving strategies used by subjects at varying levels of expertise. This, we felt, would give us a good idea of what things experience teaches. Based on our previous work on memory and problem solving (Kolodner, 1985; Kolodner & Simpson, 1984; Kolodner & Kolodner, 1987), we expected that differences would be in both the amount known and accessibility (or organization) of known knowledge. Our second goal was to identify particular changes over time in each individual's handling of specific problems and types of problems. The sequence of problems presented to the subjects was derived such that this would be possible.

The task domain we have chosen to look at, *diagnosis of automotive problems*, is interesting for several reasons. The automobile engine is a highly complex entity. It consists of a number of interacting systems acting to produce the car's motion. Failures in any component or system of the engine usually produce noticable symptoms or changes in the car's performance, but the failures themselves are seldom obvious to the amateur. In addition, a given symptom can indicate numerous possible failures within the engine. The person who comes to the shop with a problem describes a symptom or set of symptoms to the mechanic, and it is the mechanic's job to further investigate the car to find out which of the many possible problems that could cause the reported symtom(s) is in fact responsible for it. Experts are much better than novices at determining the causes of automotive problems. (As the old story goes; it's ten cents for the screw and twenty dollars for knowing which one to replace.)

The domain is knowledge-rich, and the depth of knowledge and ability to use it are both important in making a good diagnosis. Schools teach about cars in general, but since there are so many different kinds of cars, each of which have their own peculiarities, textbooks and schools can't teach everything. Diagnosing a car with a given set of symptoms may depend as much on the age and type of engine as on the symptoms presented. A given failure can be a common cause of a particular symptom in one engine and not possible in another. Experience with different types of cars and different types of problems is thus essential in gaining expertise. Furthermore, there are too many types of cars (most models change at least a little every year) and too much in the sets of manuals for individual cars for a mechanic to know everything about every car. Thus, it is essential for the expert mechanic to draw his own generalizations about cars that allow him to organize and access knowledge appropriate to any particular car and problem he is looking at.

In the work reported here, three student mechanics were observed while diagnosing car failures. Six problems were presented at weekly intervals and think-aloud protocols were collected while the students worked and were transcribed and coded for later analysis. Each week the instructor demonstrated the correct or optimum troubleshooting sequence for diagnosis of the failure after all subjects were finished. Thus, each student had an opportunity for feedback and an explanation of the car's problem whether or not he had diagnosed it correctly. Each failure was introduced into the car deliberately and each problem was caused by only one failed part. Analysis of the data focussed on the knowledge and strategies used by students at different levels of training, how their knowledge was organized, and how their knowledge and strategies changed with experience.

We expected that the more experienced student would solve more problems and would give evidence of having a more organized knowledge base than the less experienced students. In addition, we expected that individuals would show evidence over the series of problems of acquiring new diagnostic skills and new knowledge and connections within their knowledge.

2. Method

2.1. Subjects

Three students at a post secondary technical school volunteered to participate in the project. The technical program is a two-year, eight-quarter program. During much of the second year, the students work in a shop setting within the school. Cars belonging to school personnel and friends of the students and instructors are diagnosed and repaired by students. In addition, the school owns several cars that can be used in teaching students to teach about specific problems.

Each of the three student volunteers was at a different point in the program. The novice student was in his first quarter of the program and had no prior training or experience. The intermediate student was at the beginning of his second year in the program. The advanced student was near the end of the second year and held a part-time job as a mechanic outside of school. Each student worked on at least four of six problems.

2.2. Procedure

Subjects were observed once a week while diagnosing an actual problem in a car. The problems used were selected by an instructor in the program in consultation with the experimenter. The problems and the information given as the customer's complaint are described in Table 1. Each fault was introduced into a car by the instructor or by a student not in the study under the direction of the instructor. The cars used were all owned by the school with one exception: a new car brought in by a school official that had symptoms we had been presenting to the students in previous weeks. In every case, a single complaint was given and a single fault could be traced to account for the complaint. Students were told to track down the fault, but not to fix it unless repair was necessary to confirm the diagnosis.

In each session, the student was led to the car and, with the experimenter posing as a customer, told that the car was exhibiting a particular symptom. The student was then allowed to perform any tests desired on the car and its engine, with the exception of a driving road test, prohibited primarily by the symptoms presented by the car. The student was instructed to think aloud as he worked to find the failed component in the car. His comments were tape recorded by the experimenter, who also served as an assistant to the student when necessary.

Table 1
Faults and their complaints as presented to subjects

Problem	Complaint (Symptom)	Fault
1	cranks but will not start	sediment or other blockage in gas line
2	cranks slowly when starting	bad cell in battery-will not hold charge
3	cranks but will not start	bad connection behind fuse panel and fuel pump fuse
4	cranks but will not start	loose ground wires from Electronic Control Module (computer)
5	cranks but will not start	open tach circuit
6	detonation on acceleration	poorly adjusted timing

2.3. Coding

After all protocols were transcribed, each statement was coded into one of six categories, shown in Table 2 with examples. Statements coded as *hypotheses* were those in which a specific system or component was first named as a possible source of the failure or in which the system or component was accepted or rejected as the

Table 2
Coding Categories for Protocols

Category	Additional Specifications	Examples
Hypotheses	Number and Status	Could be starved for gas (N-P1) It could be, could be the starter (N-P2)
Rules	Topic(Failure, normal functioning, or troubleshooting)	Fuel Pump should come on for 3 seconds (I-P4) First of all, I have to locate the connector to the back of the fuel pump (A-P3)
Information Gathering	Source of information obtained	Before I look in the book, I'm going to check the fuse (A-P3)
Observation	Topic (hypothesis(number) or complaint)	What we don't have is fuel to the throttle body (A-P3) I don't believe I hear it running (A-P3)
Restatements	Topic (complaint or summary of observations)	to rephrase that-the throttle body is not injecting fuel (A-P3)

source of the failure. Hypotheses were numbered in order of appearance and, each time one was mentioned, its status was noted. Its status could be *open, accepted, confirmed, or rejected*. *Rules* were statements giving known, constant information about an engine or about the process of diagnosis. Statements coded as *information gathering* were generally descriptions of the actions being taken by the subject at the time. Such actions could elicit or obtain information from the customer, from a book, or via a procedure or test applied to the engine. *Observations* were statements giving the information obtained from the action taken. *Restatements* were repetitions of previously stated or collected information rather than new information. Each statement falling into one of the last three categories was identified with a specific hypothesis by its number if possible. All other statements were uncodable and were marked as such.

3. Results and Discussion

As expected, the ability of the students to correctly diagnose the problems changed substantially between the novice level and the intermediate and advanced levels. The diagnoses given by each subject and the number of hypotheses considered are shown in Table 3. The novice correctly diagnosed only one of four problems attempted, while the intermediate student correctly diagnosed three of six and the advanced student three of four. In addition, the number of hypotheses considered increased with expertise. The novice generated a mean of 3.0 hypotheses per problem and the intermediate and advanced students generated 6.8 and 5.0 hypotheses per problem respectively.

Table 3
Final Diagnoses and Number of Hypotheses Considered by Each Subject

Problem	Novice	Intermediate	Advanced
1	not getting fuel(4)	clogged fuel line(3)	-------
2	dead battery cell(5)	starter(4)	dead battery cell(6)
3	-------	fuel pump relay(5)	fuel pump fuse(5)
4	fuel pump(3)	no diagnosis(6)	injector solenoid(9)
5	no diagnosis(0)	open tach circuit(9)	-------
6	-------	bad timing(10)	bad timing(2)

3.1. Knowledge Structures and Knowledge Organization

In general, the diagnostic behavior we saw was similar to that reported by other researchers (Hunt, 1981; Rasmussen, 1978; 1979; Rasmussen & Jensen, 1974). Students generated one or more possible hypotheses for the failure immediately after observing the symptom(s). These hypotheses were then tested in a fairly systematic (albeit sometimes idiosyncratic) way either by observation of the inputs to and outputs from specific components and systems or by performance of specific diagnostic tests. In successful cases, a single diagnosis ultimately was given, accompanied by an explanation of how or why that failure would generate the observed symptom(s).

We interpret this process as being indicative of an interaction between two types of knowledge structures. The first, a *causal model* of the car's engine, contains knowledge about individual components and their inputs, outputs, and normal behavior; relates components within a system to one another; and describes the relationships and connections between systems. It is used to evaluate hypotheses in light of the evidence obtained from the failed engine and to lead the mechanic through the engine to the source of the problem in a systematic way. The causal model is generally quite large, and the second type of knowledge structure, *symptom-fault sets*, is used to index into the causal model at appropriate places. Symptom-fault sets represent the relationships between particular symptoms or sets of symptoms and failures. For example, given the symptom "the car cranks but will not start", the symptom-fault sets will identify three systems as possible locations for the failure: the fuel system, the air intake system, and the ignition system. Within each of these systems, additional symptom-fault sets will identify individual components that may cause the symptom(s). For the fuel system, these would be a failed fuel pump, an empty gas tank, or a blocked fuel line. For the ignition system, these would be a bad distributor, bad spark plug wires, or bad spark plugs. These symptom-fault sets are used to derive initial hypotheses, directing the mechanic to look at only appropriate places in the causal model.

If, in fact, mechanics are using these two types of knowledge structures during troubleshooting, then we can predict several changes we should expect to see in these structures as a result of experience, and from those, we can predict the processing differences that would result from these changes. First, we predict that through experience, a mechanic's set of symptom-fault sets increases and that the sets he already knows become more accurate. As a result of these changes, the mechanic should have better ways to index into the causal model, leading to more efficient searches for the correct failure. Second, the causal model should become more filled out with experience, both through addition of components and/or systems that were previously unknown and through addition of relationships and dependencies between the known components. The causal model, like symptom-fault sets, should also become more accurate. As a result of having a better causal model, a mechanic should be better able to systematically reason about the way the car works, allowing him to find engine failures more systematically and in more cases.

We did, in fact, see clear differences between students at different levels of experience reflecting exactly these changes in their knowledge structures. First, we saw evidence that both the organization and number of symptom-fault sets increased with experience. The advanced student seemed to know more symptom-fault sets than the novice, as evidenced by the larger number of hypotheses he was able to generate for each problem. In addition, the advanced student seemed to organize his symptom-fault sets differently than the novice, evidenced by the more systematic procedure he used for generating and testing hypotheses. The advanced student's procedure was to zero in on one of the engine's subsystems and then to consider which component of that system was faulty, while the novice did not differentiate between systems and components of systems in diagnosis. While for the novice, all faults are equal and an hypothesis at the component level was as likely to be selected as the first to investigate as an hypothesis at the system level, the more advanced troubleshooter seemed to organize his symptom-fault sets into two categories, each used for different purposes. One set pointed to faulty *subsystems* within the car (e.g., fuel system, electrical systems) and was used early in diagnosis to zero in on the faulty subsystem, while the second set pointed to faulty *components of these systems* (e.g., the fuel pump, the battery) and was used to diagnose the problem within that system. Such a change requires that the mechanic also reorganize his knowledge about the car's engine in a more hierarchical way that differentiates between systems and components of systems. Figure 1 shows a portion of the novice and advanced student's organizations of the causal model of the engine.

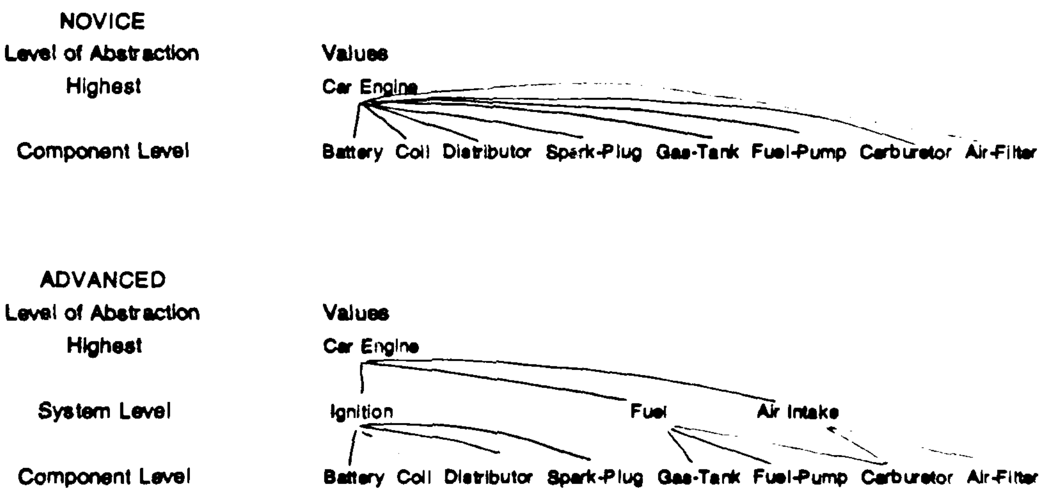

Figure 1
Novice and Advanced Student Representations of the Car's Engine

We also saw evidence that content of the causal model changed with experience. The causal model of the more advanced students contained not only more knowledge about individual components, but also more knowledge about the interconnected nature of the engine's systems. The behavior of the students during troubleshooting illustrates these findings. Consider, for example, the behavior of the advanced student in Problem 4. His reasoning went as follows:*

> The first thing you want to do, which is the easiest thing to do, is look and see if we have any fuel, because you gotta have fuel, air, and heat... Dont have fuel..The first thing I want to do is check the fuse...they're OK... hook this jumper lead to the bypass to the fuel pump...the fuel pump is running... check and see our connection up here to the energizer...going from the ECM up to the injector is OK...try to energize this solenoid by hand...check to see if we got any gas...all the lines are alright...got gas to the throttle body... my diagnosis is the solenoid is bad because everything else checks out.

The hypotheses generated by this student are in an order that reflects the multi-level and highly integrated

* For a full protocol of the session, write to the first author.

organization of both his causal model and his symptom-fault sets. He first determined which of three possible systems of the engine was affected and then investigated its components and others that could impinge on the behavior of the system under focus. In fact, his primary focus was on the electronic (or computer controlled) influences on the behavior of the fuel pump and fuel injectors. This reasoning showed an awareness (reflected in the student's causal model) of the interdependencies between subsystems. His reasoning shows that he knows that systems (such as the fuel and electronic systems) may intersect at several points and that an apparently or possibly failed component in one system may reflect an action, or lack of action in another system.

In contrast, the novice generated relatively few hypotheses for any given problem. His protocols indicate that this is because he has little knowledge about the relationships between given symptoms and their causes and also because his causal model is inadequate. In solving the same problem the advanced student was working on above, the novice reasoned:

> This problem could be in the fuel system, ignition system...we know it's not in the starting system because the car will crank over...One small drop of fuel...in that bowl...so it's in the fuel system...the fuel pump's... supposed to turn for 10 to 15 seconds...I can't hear it...it might just be a bad fuel pump.

We can see little evidence of an integrated hierarchy of levels in his organization of symptom-fault sets. While his hypotheses were sometimes at the system level (i.e., fuel system) and sometimes at the component level (i.e., fuel pump is bad), in only one problem (this one) did he clearly consider first a system and then a component within that system. More commonly, he generated hypotheses at both levels and then investigated only specific components. Furthermore, he showed a similar lack of integration in his causal model. Specifically, he never considered the possibility that one system could affect the behavior of another. His knowledge appeared to stop at the individual component's behavior and did not include the possibility that the actions of another system (the electronic system) could be affecting the behavior of the component he was considering (the fuel pump).

While the novice knew about many of the components of the car's engine and about what their connections were within a single system, he did not know how the systems and the components in different systems were interrelated. The advanced student, on the other hand, knew both the connections between components and the connections between systems. Thus the advanced student had a more integrated and complete understanding of the car's engine, while the novice's understanding seemed to be highly disjoint. Figure 2 shows our interpretation of what the novice and advanced students knew about the fuel pump, for example.

PUMP	Source: a container
	Substance: a substance in the container
	Conduit: a pipe
	Destination: a container
	Energy-Source: an energy device

NOVICE		ADVANCED	
FUEL PUMP	ISA PUMP	FUEL PUMP	ISA PUMP
	Source: gas tank		Source: gas tank
	Substance: gasoline		Substance: gasoline
	Conduit: hose		Conduit: hose
	Destination: carburetor		Destination: carburetor
			Energy-Source: electrical system

Figure 2

Novice and Advanced Student Representations of a Fuel Pump

Note that the general information about pumps is available to both the novice and the advanced student. However, the information that the fuel pump requires an energy source which is the electrical system of the car is not part of the novice's representation of the fuel pump. If asked "What makes the fuel pump run?", the novice is able to construct the appropriate answer by using the more general information about pumps, but he does not use this knowledge during problem solving. The same pattern is probably true of knowledge about systems and components. The novice can undoubtedly tell an inquirer what system of the engine a particular component resides in, but he does not maintain this information where it is readily usable during problem solving.

We also saw within-subject changes in these knowledge structures over the course of the experiment. These changes were most evident in the intermediate student. Two examples will serve to demonstrate changes across problems. In working on problem three, the intermediate student made a long and protracted search for the fuel pump relay using both written reference materials and extended visual examination of the engine. While working on problem four, he was able to immediately locate and check the same part. This component, and its physical relationship to others, had been incorporated into the causal model during or following problem three. Similarly, the symptom-fault sets changed as new information was acquired. For example, the first hypothesis the intermediate student checked at the component level for problem four was the fuel pump fuse, which was the correct diagnosis for problem three. He made the point as he worked that he was checking this possibility out first because of the previous case. ("I'm gonna check the fuel pump fuse *first* [this time].")

3.2. Diagnostic Strategies

In addition to the changes experience makes in knowledge structures and organization, we also saw differences in diagnostic style. Diagnostic strategies seemed to be used differently by subjects at different levels of expertise and evaluation criteria changed significantly with experience. Some of these changes are due to the development of better strategies for testing and confirming hypotheses with experience while others appear to result from the differences in the knowledge available for diagnosis as a mechanic gets more experienced.

The change in how the mechanics tested and confirmed hypotheses was striking. As the example above showed, the novice student was willing to accept an hypothesis when preliminary evidence could be interpreted as congruent with that hypothesis and not pursuing the task any further (i.e. "can't hear the fuel pump"). In contrast, the advanced student sought, for each hypothesis, specifically confirming or disconfirming evidence that was part of a causal explanation. While he was willing to select an hypothesis to pursue on the basis of preliminary evidence, he would not accept or reject it without causally based information (i.e. "the fuel pump's not running, now we have to find out why").

The changes in diagnostic strategies that resulted from changes in the knowledge structures were more apparent in the efficiency of diagnosis. As the causal model gets filled out, it should allow the mechanic to pursue a longer systematic search through the engine and also allow him to evaluate information in more detail and with more concern for the real effects of the behavior observed. At the same time, as the number and complexity of symptom-fault sets increases, long searches should become less necessary, because the mechanic is able to index into his model in more, and more effective, locations.

These two types of changes in the mechanic's diagnostic strategies work together to produce the results we saw. As the mechanic gains experience with making correct and incorrect diagnoses, he gains a sense of what kind and how much information is "enough" to be sure of his opinions. In addition, as his causal model and symptom-fault sets become more complete and accurate, he is more able to select hypotheses for investigation appropriately and to continue invgestigating a problem to the point that only one hypothesis remains as a possible diagnosis. Consequently, the conditions under which he will accept an hypothesis as a final diagnosis will become more accurate and the path by which he reaches his diagnosis will become more efficient.

This result is clearly evident in protocols of the novice and advanced students. When the novice's working hypothesis was a that a particular component was faulty, he either accepted it or rejected it as the cause of the symptom. He never investigated other effects on or inputs to that component. For example, in problem 2, the failure was a dead battery cell which caused the car to crank very slowly. The novice based his diagnosis on the following information:

> First of all, we'll have to check this battery...it could be the starter... it could be the alternator...it could be a voltage loss...could be a dead cell in the battery...we've only got 10 volts in the battery--each battery cell is 2 volts and there's 6 cells in the battery, so dead battery cell.

Here we see the novice generating both system and component level hypotheses but, because his knowledge is not hierarchically organized, not pursuing them in that order. Rather, he looks first at the battery charge. Because it is low, he accepts the hypothesis of a dead cell. His diagnostic strategy does not require that he consider any hypotheses relating to why the battery might be low, such as a malfunction in another system.

In contrast, the advanced student generally collected more information before giving a diagnosis. If possible, he confirmed his diagnosis by visually finding the condition that created the symptom (i.e., the disabled fuse panel connection in Problem 3). When that was not possible, he justified his diagnosis within his causal model. For example, in Problem 2, the failure could not be confirmed by visual evidence. Instead, the advanced student reaches his diagnosis with the following information:

> ...check the starter draw...it's pulling enough down to get the starter to go alright...We put the battery under load, you can see the amps rising and it's charging the battery...So the alternator's working OK...what I believe we have is the cell is dead in the battery...Try the test on the VAT...As you see on the indicator is also showing that it needs charging for the battery is bad...So what we have here is a battery with a couple of cells dead, and it's a sealed battery and you cannot check the specific gravity with a hydrometer to check and see which one's dead.

He reached and justified his diagnosis by eliminating all other possibilities from his symptom-fault sets and the causal model. In other words, he tested and verified normal functioning of both the starting system ("It's pulling enough down to get the starter to go alright") and the charging system ("So the alternator's working OK"). These are the only two systems, other than accessories such as headlights and radio, that affect the level of charge in the battery. Consequently, according to the student's causal model, if the battery's charge is low and the starting and charging systems are functioning correctly, the only remaining component in which the failure can be located is the battery itself. In some types of batteries, this conclusion can be tested directly, but in the car used in this problem, the battery is sealed. Therefore, the mechanic must stop with his explanation rather than attempt to verify the diagnosis any further. In comparison to the novice, he selected his hypotheses more efficiently, first eliminating competing systems from consideration. In addition, he based his acceptance of the diagnosis on a full causal explanation rather than on superficial evidence.

4. Conclusions

The results are as predicted by our interpretation of the diagnostic behavior as an interaction between several knowledge structures. Both the causal model and the symptom-fault sets change with experience, and we have seen some examples of exactly what changes occur. In the causal model, the most notable change is the increasing complexity of the model, reflected in the growing awareness of the interconnectedness of systems within the engine. The novice is clearly unaware of the possibility that electronic failures can affect things like fuel delivery, since he knows little about the dependencies between the fuel system and the electrical system, while the more advanced mechanic not only knows that such relationships exist, he considers them a highly common source of failures. Similarly, the number, organization, and accuracy of the symptom-fault sets changes with increasing experience. Ultimately, they are able to represent a complex, hierarchical system of relationships. The data suggest that components are organized hierarchically under their respective systems and are never directly considered unless their system is determined to house the failure, or at least to be the source of information crucial to locating the failure.

Building partly on these changes in the knowledge structures, and partly on independent effects of experience on decision processes, the mechanic's procedures and guidelines for accepting hypotheses as diagnoses also change. The processes or procedures used become increasingly focussed on information that will allow a causal interpretation of the behavior observed. At the same time, the developing knowledge structures allow the mechanic to search for and aquire more, and more accurate, information from his symptom-fault sets and his causal model. The interaction of these changes in both knowledge and process lead to the more accurate and efficient problem solving seen in experts.

Thus, we see that experience is providing the mechanic with three things. His overall level of knowledge is increasing; the organization and integration of his knowledge structures, both the symptom-fault sets and the causal model, are increasing; and his processes and criteria for reaching diagnoses are becoming more accurate, more efficient, and more focussed on causal information.

5. References

Akin, O. (1980) *Models of architectural knowledge* London: Pion.

Bhaskar, R. & Simon, H.A. (1977) Problem solving in semantically rich domains: An example from engineering thermodynamics. *Cognitive Science, 1*, 193-215.

Chi, M.T.H., Glaser, R., & Rees, E. (1982) Expertise in Problem Solving. In R.J. Sternberg (ed) *Advances in the Psychology of Human Intelligence* Hillsdale: Lawrence Erlbaum.

Glaser, R. (1985) Thoughts on expertise. Technical Report #8. Learning Research and Development Center, University of Pittsburgh, Pittsburgh, PA 15260.

Hunt, R.M. (1981) Human pattern recognition and information seeking in simualated thought diagnosis tasks. Report #T-110, Coordinated Science Laboratory, University of Illinois at Urbana-Champaign.

Kolodner,J.L. (1985) Experiential processes in natural problem solving. Technical Report #GIT-ICS-85/23. School of Information and Computer Science, Georgia Institute of Technology, Atlanta, GA 30332.

Kolodner,J.L. & Kolodner,R.M. (1987) Using experience in clinical problem solving: Introduction and framework. *IEEE Transactions on Systems, Man, and Cybernetics.*

Kolodner,J.L. & Simpson,R.L. (1984) Experience and Problem Solving: A Framework. *Proceedings of the Sixth Annual Conference of the Cognitive Science Society.*

Miller, P.B. (1975) Strategy selection in medical diagnosis, Project MAC. Report #TR-153, Massachusetts Institute of Technology, Cambridge, MA.

Rasmussen, J. (1978) Notes on diagnostic strategies in process plant environment Riso National Laboratory Report #RISO-M-1983, Roskilde, Denmark.

Rasmussen, J. (1979) On the structure of knowledge-A morphology of mental models in a man-machine context. Riso National Laboratory Report #RISO-M-2192, Roskilde, Denmark.

Rasmussen, J. and Jensen, A. (1974) Mental procedures in real life tasks: A case study of electronic troubleshooting. *Ergonomics, 17(3)*, 293-307.

Reed, S.K., Ernst, G.W., & Banerji, R. (1974) The role of analogy in transfer between similar problem states. *Cognitive Psychology, 6*, 436-450.

Reed, S.K. & Johnson, J.A. (1977) Memory for problem solutions. In G.H. Bower (ed) *The Psychology of Learning and Motivationk* New York: Academic Press.

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

Simon, H.A. (1975) The functional equivalence of problem solving skills. *Cognitive Psychology, 7*, 268-288.

Simon, D.P. & Simon, H.A. (1978) Individual differences in solving physics problems. In R. Siegler (Ed.) *Children's Thinking: What Develops?* Hillsdale, N.J.:Lawrence Erlbaum

Voss, J.F., Greene, T.R., Post, T.A., & Penner, B.C. (1983) Problem solving skill in the social sciences. In G. Bower (Ed.), *The psychology of learning and motivation: Advances in research theory.* New York: Academic Press.

Voss, J.F. & Tyler, S. (1981) *Problem solving in an social science domain.* Unpublished manuscript, University of Pittsburgh, Learning Research and Development Center.