Learning Through Growth of Skill in Mental Modeling[1]

Jill H. Larkin
Herbert A. Simon
Carnegie-Mellon University

The mental models of a skilled scientist are often different from those of an untrained person. For example, in thinking about the interaction of physical objects, the untrained person seems largely restricted to envisioning objects in a sequence of motions. The entities in these "naive" or *immediate* mental models correspond directly to objects in the real world. The inferential rules that control the running of these models correspond roughly to rules reflecting how events unfold in real time. While such mental models are perfectly adequate for getting around in everyday situations, they are sometimes dramatically wrong (Green, McCloskey, and Caramazza, 1980, Clement, 1981) and they certainly seem less effective in solving scientific problems than are the more extended representations of the scientist.

The mental models of a person with training in physics are not limited to entities and inferential rules based directly on experience. Instead, these models can and do include entities that have technical meanings defined only by the scientific discipline, and that are related by special inferential rules again defined in the discipline. For example, a person with training in physics is not restricted to considering perceivable objects like cats and coffee cups, but may also represent situations in terms of technical entities like forces or pressure drops. Similarly, capacity to make inferences about a situation need not parallel imagined development of the situation in time, but instead may reflect special constraint laws of physics, e.g., that the momentum of an isolated system must remain constant. These ideas are discussed more fully in (Larkin, 1981a) and in the discussion of physical intuition in (Simon and Simon, 1978).

In this paper we consider how an individual might develop the ability to re-represent situations in terms of scientific entities. Presumably this development is one goal of science instruction. We shall present preliminary results from an experimental and theoretical case study in such development. Subjects with backgrounds in physics studied sections taken from a physics textbook that described material (fluid statics) they had not previously encountered, and then used this material in efforts to solve problems. In a coordinated theoretical effort we are developing a computer-implemented model of learning from text that is capable of using declarative statements of facts (in this case, relations of physics) both to "understand" the derivation of new results and to apply these results in solving problems.

1. The ABLE system

The system, called ABLE, is a descendant of a system that learned through practice to apply principles of mechanics, and that accounted for strategy differences between skilled and less skilled individuals (Larkin, 1981b).

ABLE is a *production system* written in the current implementation of OPS, a LISP-based efficient production-system language (Forgy, 1980, Forgy, 1979). Thus ABLE has a *working memory* composed of passive elements of knowledge that are acted on by a large *production memory* composed of elements of procedural knowledge encoded as condition-action pairs. When the conditions of a particular production are found to match some of the contents of working memory, this match cues the execution of the corresponding actions which then act to modify working memory. Then the conditions of some new production are found to match, and the cycles continue. Production systems have a continued history of fruitfulness in psychological modeling (Newell and Simon, 1972, Newell, 1973, McDermott, 1978) but the major feature used in ABLE is the easy modeling of learning. Each production is an independent piece of knowledge, and the circumstances under which it applies are determined only by the contents of its own conditions. Thus the addition of knowledge (learning) is modeled simply by the addition of new productions.

To explain the working of ABLE we consider its application to solving part of the problem given in Table 1 and Figure 1a, and presented as a worked example in Halliday & Resnick (1970). The first paragraph of the example states the problem. We currently give ABLE a good understanding of this paragraph, i.e., a good immediate representation of the problem. It is coded as a set of related declarative elements in working memory, indicated by the graph structure in Figure 1b.

1.1. Encoding of Principles

After achieving this immediate representation of the situation, how does a solver make scientific inferences of the kind illustrated by the textbook solution given in Table 1? In other words how is a scientific mental model run?

Such inferences must be based on scientific principles that are in some sense "known" to the solver. "Known" might initially mean that the appropriate textbook page is available for inspection. We discuss later the growth of other kinds of knowing. Thus we provide ABLE with knowledge of relevant principles in the form illustrated in Figure 2. Like all principles and definitions in ABLE, it includes a symbolic *statement* of the principle $\Delta p \quad \rho g h$ together with a *setting* to which the principle applies and in terms of which of the symbols in the statement are defined. Here the setting includes a portion of liquid with density ρ, two points in that liquid separated by a height h. The "gravitational acceleration" $g = 9.8 m/s^2$ is not specified but assumed to be known outside the context of this principle.

Table 1: Worked example from a textbook (Halliday and Resnick, 1970) showing the application of relations of fluid statics to relate densities of liquids in a U-shaped tube.

A U-tube is partly filled with water. Another liquid, which does not mix with water, is poured into one side until it stands a distance d above the water level on the other side, which has meanwhile risen a distance l (Fig. 1a). Find the density of the liquid relative to that of water.

In Fig. 1 points C are at the same pressure[1] Hence, the pressure drop from C to each surface is the same[2], for each surface is at atmospheric pressure[3]

The pressure drop on the water side is $\rho_w 2l^4$, where the $2l^5$ comes from the fact that the water column has risen a distance l^6 on one side and fallen a distance l on the other side, from its initial position. The pressure drop on the other side is $\rho g(d + 2l)^7$, where ρ is the density of the unknown liquid. Hence,

$$\rho_w g2l = \rho g(d + 2l)^8$$

and

$$\rho/\rho_w \quad 2l/(2l + d)^9.$$

The ratio of the density of substance to the density of water is called the *relative density* (or the *specific gravity*) of that substance.

[1-9] These numbers label inferences for reference later in the text.

This knowledge of a principle is encoded as a passive link-node structure involving no knowledge of how or when to apply the principle. In this sense it is *declarative* knowledge, although clearly

[1] This work was supported by NIE-NSF grant number 1-55862, by NSF grant number 1 55035 and by the Defense Advanced Research Projects Agency (DOD), ARPA Order No 3597 monitored by the Air Force Avionics Laboratory under Contract F33615-78 C 1151 The authors acknowledge the important contributions of Susan Cotten in the collection, transcription and coding of the protocol data

Figure 1: (a) Diagram provided by the textbook for the example in Table 1. (b) Annotated "diagram" (immediate representation) provided for ABLE in starting to work the example.

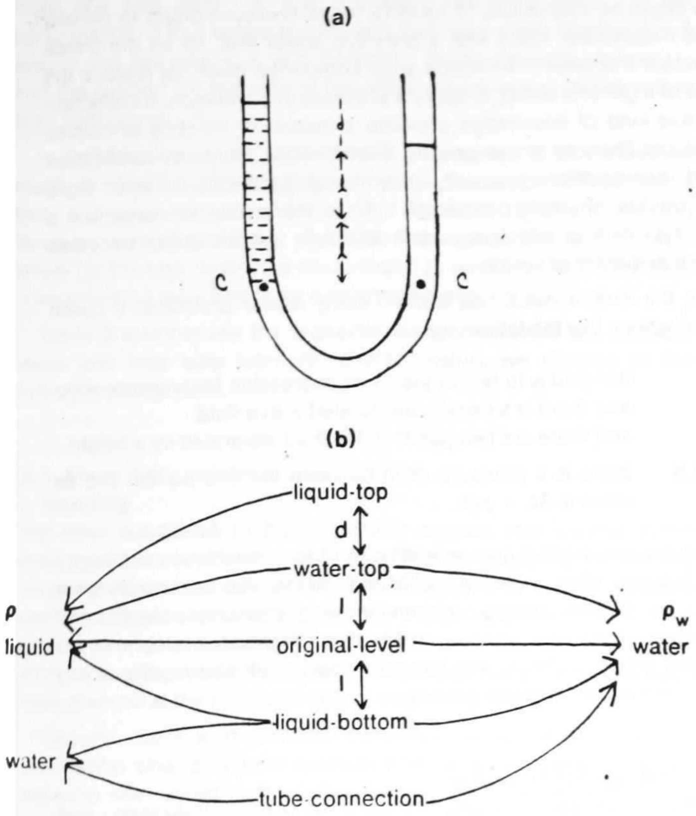

Figure 2: Graph structure representation of the principle $\Delta p = \rho g h$.

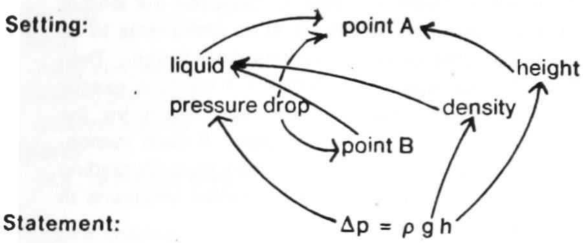

it goes beyond minimal propositional encoding of the phrases that may have been used to describe it in the textbook. To illustrate how ABLE uses this knowledge of principles, we consider its application to develop the inference labeled 4 in Table 1, that is to infer that the pressure drop from C to A on the water side of the tube is $\rho g 2l$.

1.2. Interpretive Use of Principles

ABLE applies declarative knowledge through general procedural knowledge that first matches the setting of the principle to the setting of the problem, and then uses the statement of the principle (interpreted in the setting of the problem) to make inferences.

The control structure of ABLE, shown in Figure 3, is based on that proposed by Neves and Anderson (1981) for the analogous task of supplying reasons for the statements in a geometry proof. The following paragraphs provide English statements of the productions that control the shift of goals, Figure 3 indicates these productions by labeled arrows linking goal statements.

ABLE starts by using the knowledge in production P1 that if the

Figure 3: Goals used in ABLE with labeled arrows indicating productions that change goals.

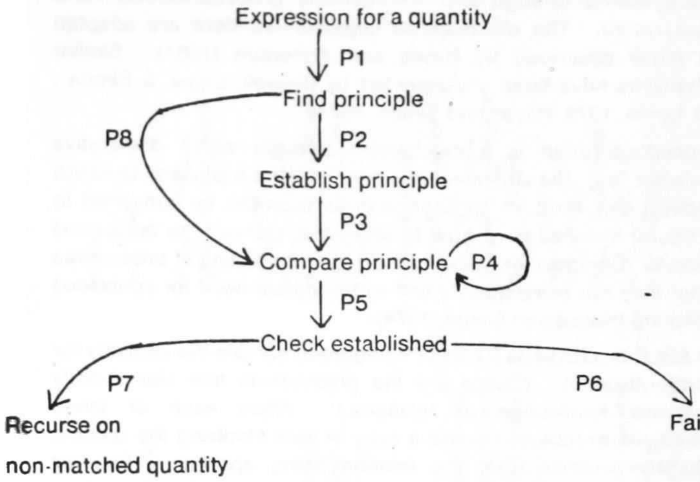

goal is to find an expression for a quantity in a problem setting, then one should search for a principle that can provide further information about that quantity. The initial search can be based on a variety of criteria (cf. Simon & Simon (1978), Larkin 1981b). ABLE currently uses a very rough process, embodied in P2, that to be considered a principle must involve a quantity of the type currently desired (here pressure drop).

ABLE then applies the knowledge in production P3 to set the goal of comparing elements of the problem situation with elements of the principle, including its statement and setting. Production P4 contains knowledge of how to trace and compare the graph structures representing the current problem situation (e.g., Figure 1b) and the principle situation (Figure 2). When all possible correspondences between the two graph structures have been matched, production P5 sets as a goal to check whether all parts of the principle have been matched. If they have, the goals are successively marked as succeeded. If this is not the case, production P6 recognizes that part of the situation crucial to the applicability of the principle has not been satisfied. The goals are then successively marked as failed, and ABLE ultimately must seek a different principle. If, however, the only part of the principle situation that does not have a correspondence in the problem situation is a particular theoretical entity (quantity) then production P7 recognizes that if this correspondence could be established then the principle would succeed. Thus in our example, if ABLE had not already established the correspondence $h = 2l$ (as the text solution has not -- see Table 1), then ABLE would set as a subgoal to establish an expression for the height h, beginning work again with production P1 in Figure 3.

Much of ABLE's work involves establishing a detailed match between a principle setting and a subset of the problem setting. In the single inference discussed above, of the total of 9 cycles of production execution, 5 were concerned with matching settings. This costly and compulsive matching seems, however, to be necessary for good problem solving. For example, in the current problem there are several densities, pressures, and heights. Without careful matching between settings, the solver can easily "infer" relations between quantities that in fact have no connection.

1.3. Reducing the Costs of Interpretive Matching

Because matching a principle to a setting is costly, it is crucial that a solver develop good search procedures for locating principles likely to produce useful inferences. The primitive search algorithm embodied in production P1 (pick a principle involving the kind of quantity you're trying to solve for) is certainly not good enough. The following paragraphs describe ABLE's ability to develop better search mechanisms.

ABLE learns through two mechanisms *proceduralization* and *generalization*. The mechanisms implemented here are adapted from those described by Neves and Anderson (1981). Similar mechanisms have been implemented by (Newell, Shaw, & Simon , 1960, Lewis, 1978, Hayes and Simon, 1974).

Proceduralization is a mechanism through which declarative knowledge (e.g., the statement of a principle and a situation to which it applies) can, through application in an example, be converted to procedural knowledge of how to apply that principle to analogous situations. Composition is the collapsing or compiling of procedures so that they run more quickly and with reduced need for conscious monitoring (Hayes and Simon, 1974).

In ABLE productions P1 and P4 (Figure 3) contain the capacity for proceduralization. (These are the productions that use directly declarative knowledge of relations.) When each of these productions executes, it builds a copy of itself involving the specific declarative entities from the relation being applied. Thus for example, when P4 applies to the principle Δp ρgh, it may apply in the form:

IF If the goal is to compare a principle to the current setting
and there are corresponding theoretical entities in the principle and problem settings
and these entities refer to two physical entities of the same type

THEN mark the two physical entities as corresponding.

This production builds a copy of itself of the form:

IF the goal is to compare the principle $\Delta p = \rho gh$ to the current setting
and the density of a liquid in the problem setting corresponds to the density in the principle setting

THEN mark this liquid in the problem as corresponding to the liquid in the principle $\Delta p = \rho gh$.

This new production is part of the procedural knowledge needed to apply the principle $\Delta p = \rho gh$.

As Neves and Anderson (1981) point out, these proceduralized productions are almost always shorter (contain fewer conditions) than the original general productions that built them. Thus they may immediately provide the advantage of reducing working-memory load. However, their main importance is that they are the ingredients for building efficient productions that recognize useful configurations in a problem and relate these configurations to potentially useful principles.

The second mechanism of learning is composition. If two proceduralized productions (any of those built by P1 and P4) execute in sequence, they are combined to form a single production that does the work of both. This is done first by collecting the condition and action elements from both and deleting repetitions. (Further details given by Neves and Anderson (1981).) For example, the proceduralized production above can be composed with a production built by P1 in Figure 3 to form the following:

IF the goal is to find an expression for Δp of a fluid
and there is a density for the fluid

THEN set the goal to compare to the current situation the principle $\Delta p = \rho gh$
with the correspondence between the pressure drops and the densities already established.

Productions of this form, indicated by P8 in Figure 3 short-circuit the primitive search algorithm of P1. They do not, however, short circuit all the interpretive matching between the principle and problem settings. In the current case, ABLE must still match lines and points and heights.. In human language the knowledge in a production like P8 might be expressed, "I have to relate pressure-drops to density and I remember there was a principle about that, so let me check whether it applies." Although such knowledge does not replace the use of a general ability to apply a principle in a situation, a collection of this kind of knowledge provides a means of locating principles that are likely to prove useful. Furthermore, as proceduralization and composition proceed, they build productions with more information in their conditions. Thus the ability to recognize a configuration in which a certain principle will be useful becomes more and more accurate.

In the limit, if ABLE has worked many similar problems, it builds productions like the following:

IF the goal is to find or justify an expression for pressure drop
and there is a density associated with a fluid
and there are two points in the fluid separated by a height

THEN there is a pressure drop between the two points, and its value is $\Delta p = \rho gh$.

At this point ABLE has at least one of the capabilities necessary for building what we have called a scientific representation for a problem. On encoding a problem involving appropriate heights and densities, ABLE immediately knows that the problem also involves related pressure drops, and ban use these entities as readily as any in the immediate representations.

1.4. Settings for Learning

When can the learning involved in proceduralization and composition occur? Clearly solving problems is one such setting. However, as we shall see in our discussion of human learners, it seems likely that this learning also takes place during study of text material. The following is one mechanism through which this might occur.

Consider the textbook example in Table 1. Suppose the learner considers the various inferences (labeled 1-9) as statements to be understood or verified on the basis of previous knowledge. Then understanding the sentence involving inference 4 involves exactly the process discussed above, justifying the expression for the pressure drops by using the principle Δp ρgh. If such reason-giving is part of active study of scientific text, then through reading the learner should acquire some ability to recognize situations to which principles will apply.

These comments are not limited to the study of text examples. New principles themselves are presented in much the same manner. A setting is described, and a sequence of inferences are stated, ultimately leading to a statement of a new principle.

2. Human Learners

In previous work (Simon and Simon, 1978, Larkin, McDermott, Simon, and Simon, 1980a, Larkin, 1981b), we have compared the problem solving of true experts, individuals with extensive professional experience in physics, with that of novices, individuals whose experience is limited to the equivalent of less than one college level course. The performance of the expert subjects would correspond here to the ABLE model after all proceduralization and composition had occurred. This very ABLE model is essentially equivalent in performance to the models of expert subjects described in earlier papers. Proceduralization and composition have produced a collection of sensitive productions that can recognize a configuration of knowledge in a problem situation, and make an immediate inference based on an appropriate principle.

The human solvers considered here are all novices with respect to the physics material, they knew varying amounts of general physics.

but none had previously studied fluid statics. The question is then to what extent do human learners, confronted with a novel section of physics text and associated problems, perform like the ABLE system? The answer is that in interesting ways human learners are more and less able. In this discussion we shall focus on the performance of four subjects, two who look very much like the ABLE system and two who look very different. These data are preliminary and the support for features of the ABLE model is suggestive rather than conclusive.

In individual sessions each subject was asked to talk aloud as much as possible while reading a six-page discussion of fluid statics from a physics textbook (Halliday and Resnick, 1970) and working three associated problems. Subjects were given one and one-half hours for the task, and were encouraged to work just as if they were completing an assignment for a science course.

Table 2 summarizes the characteristics of subjects we consider more and less able learners, characteristics we discuss in the following paragraphs.

2.1. More Able Learners

Reading

We have suggested in the preceding section that careful active reading of text may be an important setting for acquiring the partially proceduralized and composed productions that aid in locating useful principles. Indeed the two more able subjects used a great deal of effort in processing the text. First, both subjects began their work by reading the text completely from beginning to end, although both glanced at the problems before beginning to read.

Second, these learners show consistent evidence that they are processing what they read carefully and conscientiously. The two subjects interrupted their reading 53 and 70 times respectively to make comments on what they were reading. Most of these comments (75%, 81%) suggest that the subject is relating what is

Table 2: Characteristics of more and less able learners.

More Able	**Less Able**
Reading:	
Aloud, many comments	Often silently, usually few comments
Reading precedes problem solving	Problem solving before all text read
Slowly (e.g.,19 min, 20 min)	Rapidly (e.g.,7 min, 5 min)
Problem solving:	
Correct or factor of 2 error	Other errors
No search	Search common
Order of principles like ABLE	Order of principles means-ends

being read to previous knowledge ("Ok, intuition would tell you that"; "...which is analogous to just the weight of something in mechanics."), or expressions of understanding (e.g., "Ok, that's easy enough").

Problem Solving

We consider here performance on the following problem, which is very analogous to the worked example in the text (Table 1).

> A simple U-tube contains mercury. When 13.6 cm of water is poured into the right arm, how high does the mercury rise in the left arm from its initial position?

The text provides the densities of water ($1.0 \times 10^3 kg/m^3$) and mercury ($1.36 \times 10^4 kg/m^3$).

The two subjects solved this problem without any search through the text. They readily recalled the relevance of the U-tube example, found the right page in the text, and based their solutions on inferences made in that example. Thus before beginning the problem solution, these subjects had some internalized knowledge of useful principles to apply to a U-tube setting.

The order in which principles are applied is similar to that generated by the ABLE system, and differs slightly from the order presented in the original example in that information is generated in a forward working manner so that information is always available at the time it is needed. These data are presented in Table 3. Part (f) of Table 3 shows the nine inferences, labeled 1-9 in Table 1, used by the textbook in solving the analogous problem stated in Table 1. The remaining solutions are for the problem solved by the human subjects.

Table 3: Order of major steps in solution to the U-tube problem by (a) ABLE, (b) and (c) More able human solvers, (d) and (e) less able human solvers, (f) Text example.

(a) ABLE	(b) S1	(c) S2
$p_C(original) = p_C(added)$	$p_C(original) = p_C(added)$	$l_{up} = l_{down}$
$p_A(original) = p_A(added)$	$l_{up} = l_{down}$	$\Delta p_{original} = \Delta p_{added}$
$\Delta p_{original} = \Delta p_{added}$	$\Delta p_{original} = \rho_m g 2l$	$p_A(original) = p_A(added)$
$l = l_{up} = l_{down}$	$\Delta p_{added} = \rho_w(13.6\ cm)$	$h_{original} = 2l$
$h_{original} = 2l$	$\rho_m g2l = \rho_w g(13.6\ cm)$	$\Delta p_{original} = \rho_m g2l$
$\Delta p_{original} = \rho_m g2l$	$l = 0.486$ cm	$\Delta p_{added} = \rho_w g(13.6\ cm)$
$\Delta p_{added} = \rho_w g(13.6\ cm)$		$\rho_m g2l = \rho_w g(13.6\ cm)$
$\rho_m g2l = \rho_w g(13.6\ cm)$		$l = 0.486$ cm
$l = 0.486$ cm		

(d) S3	(e) S4	(f) Text Example
$p = p_0 + \rho gh$	$p = p_0 + \rho gh$	$p_C(added) = p_C(original)$
(decide U-tube example is more relevant)	h = height of mercury column	$\Delta p_{added} = \Delta p_{original}$
$\rho_w/\rho_m = 2l/(2l+d)$	$p = p_0$ = atmospheric pressure	$p_A(added) = p_A(original)$
$l = 13.6$ cm	(no place for desired quantity try another approach)	$\Delta p_{original} = \rho_w g2l$
[algebra and arithmetic]	$p = p_0 + \rho g(y_2 - y_1)$	$h_{original} = 2l$
$d = 353.6$ cm	$y_2 = 13.6$ cm	$l = l_{up} = l_{down}$
(too large try another approach)	y_1 = height of mercury column	$\Delta p_{added} = \rho g(d + 2l)$
$d = 13.6$ cm	ρ = density of mercury	$\rho_w g2l = \rho g(d + 2l)$
[algebra and arithmetic]	$p = p_0$ = atmospheric pressure	$\rho/\rho_w = 2l/(2l + d)$
$l = 0.523$ cm	$y_2 = 1.35 \times 10^{-1}$ meter	

The two subjects considered here both solved this problem correctly. All of the subjects we consider to be more able either solved this problem correctly or made the simple error of solving for the height of the original fluid (mercury) above the point C (Figure 1a), rather than solving for half that distance, the distance the mercury rose.

2.2. Less able Learners

Reading

Unlike ABLE and the more able solvers, the less able solvers skimmed through the text rapidly. The two considered here complained that reading aloud interfered with understanding and were permitted to read the text silently, which they did rapidly (see Table 2). These subjects also both stopped reading and began solving the first problem as soon as they encountered material relevant to it.

Problem Solving

Unlike the more able subjects, the less able subjects do search through the text for appropriate relations. As Table 3 shows, S3 and S4 each tried two different principles. In all cases the selection was preceded by an episode of searching the text material.

As shown in Table 3 the order in which principles are applied by the less able subjects is very different from that produced by the more able subjects, by the ABLE system, and in the analogous text example. The procedure of these subjects seems to be the following: (1) search through the text for an equation that involves distances (presumably because a distance is the quantity to be found). This equation may be the equation resulting from the U-tube example (subject S3 in Table 3) but it may not be (subject S4). (2) Substitute values for quantities appearing in this equation, using as a criterion for substitution merely left over that the value substituted must be of the same type as the symbol for which it is substituted (i.e., a height for a height, a density for a density). Indeed, even this simple constraint is sometimes violated when subjects fail to distinguish between pressure p and density ρ. (3) If after this substitution all values in the quantity have been used, and there is in the equation a quantity of the appropriate type (here distance) left over, then solve the equation if possible. (4) If it is impossible to fit all the information in the problem into the equation, then abandon it and get a new one. (5) If there remain in the equation symbols not assigned values, then search either for an expression involving this symbol, or for some "standard" value for this symbol (e.g., atmospheric pressure).

This procedure is very different from that executed by ABLE. However, it is not hard to see how such performance might be produced by ABLE through appropriate deletion of strategic knowledge. First, one would have to use the initial ABLE system, before it had built any procedural knowledge about applying principles. This absence corresponds to the lack of processing of the text observed in these human solvers. Second, one would have to remove from ABLE its strategic knowledge that a principle can be applied only if all aspects of the setting of that principle are matched against the setting of the problem (production P4 in Figure 3). This production would be replaced with one that would allow use of a relation if all symbols in it could be matched by quantities of the same type in the problem by quantities of the same type (i.e., a length for a length). This uncritical matching perhaps is associated with the less able subjects' poor abilities for selecting useful principles. They have to match a lot of principles, and so may do it in a less costly way, even though this economy has devastating effects on their problem solutions. With these changes the ABLE system could produce any of the incorrect solutions we have observed in the less able subjects.

The result is a weak means-ends procedure of searching for relevant principles observed elsewhere in novice solvers (Simon and Simon, 1978, Larkin, McDermott, Simon, and Simon, 1980b). A first principle is proposed because it contains a quantity of the type to be solved for. Subsequent principles are proposed because they can be used to replace in the original equation quantities without known values. Substitution is based on the weak criterion that the two quantities must be of the same type (e.g., two lengths, two pressures).

Because of their uncritical matching of principles to the problem situation, the errors made by the less able subjects are varied and exotic compared to the simple "sensible" error characteristic of the more able subjects. The most common error is illustrated by subject S3 in Table 3. The equation from the U-tube example is used, the distance 13.6 cm is substituted for one of the distances in the equation, l and d, and the equation is solved for the remaining symbol for distance. However, as illustrated by subject S4, the errors can be far more exotic. However, all errors are produced by copying some equation from the problem, substituting for the symbols in that equation values that correspond in type, and then solving for a symbol for distance.

3. Conclusion

As we have noted elsewhere (Larkin, 1981a, Simon and Simon, 1978) individuals trained in physics seem to work with mental models that are different than those used by less trained individuals. In particular, skilled individuals re-represent the problems in terms of technical entities (e.g., pressure drops) that have no special meaning outside the discipline of physics. Here we suggest that the general learning mechanisms of proceduralization and composition provide some explanation of how this ability to re-represent problems might be acquired.

Our prototype ABLE system acquires a principle in declarative form, as a student might by reading a chapter. This initial encoding does not itself include any information about how or where to apply the principle. Thus initial applications of the principle are *interpretive*, achieved through general procedural knowledge about how to apply any principle or definition. Through such application ABLE builds new specific procedural knowledge associated with the principle. Initially fragments of procedural knowledge aid in the search process. They contain patterns of information that have been used with that principle in the past, thus short-circuiting ABLE's original general and weak method of selecting principles. Ultimately a principle can be completely proceduralized (for a set of analogous contexts) so that application is completely automatic. This final automatic knowledge may well be an ingredient of what one would want to call an expert's mental model in which technical entities (e.g., pressure drops) are seen as readily as visible entities like heights.

References

Clement, J. Students' preconceptions in introductory mechanics. *American Journal of Physics*, 1981, , in press.

Forgy, C.L. *Implementing a Fast Syntactic Pattern Matcher.* Technical Report, Computer Science Department, Carnegie-Mellon University, August 1979.

Forgy, C.E. *Preliminary OPS5 Manual.* Technical Report, Computer Science Department, Carnegie-Mellon University, 1980.

Green, B., McCloskey, M., Caramazza, A. Curvilinear Motion in the Absence of External Forces: Naive Beliefs About the Motion of Objects. *Science*, December 1980, *210(5)*, 1139-1141.

Halliday, D. and Resnick, R. *Fundamentals of Physics.* New York: John Wiley & Sons 1970.

Hayes, J.R., and Simon, H.A. Understanding written task instructions. In Gregg, L.W. (Ed.). *Knowledge and Cognition,* Hillsdale, NJ: Lawrence Erlbaum Assoc., 1974.

Larkin, J.H., McDermott, J., Simon, D.P., Simon, H.A. Models of competence in solving physics problems. *Cognitive Science*, 1980, *4*, 317-345.

Larkin, J.H., McDermott, J., Simon, D.P., and Simon, H.A. Expert and novice performance in solving physics problems. *Science*, June 1980, *208*, 1335-1342.

Larkin, J.H. *The Role of Problem Representation in Physics.* C.I.P. 429. Carnegie-Mellon University, 1981. Presented at a Conference on Mental Models, La Jolla, CA. October, 1980.

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

Lewis, C.H. *Production system models of practice effects*. PhD thesis, University of Michigan, 1978.

McDermott, J. Some strengths of production system architectures. In *NATO A.S.I. proceedings: Structural/process theories of complex human behavior*, The Netherlands: A.W. Sijthoff, International Publishing Company, 1978.

Neves, D. and Anderson, J.R. Knowledge Compilation: Mechanisms for the Automatization of Cognitive Skills. In Anderson, J.R. (Ed.), *Cognitive Skills and Their Acquisition.*, Hillsdale, NJ: Lawrence Erlbaum Associates, Inc., 1981.

Newell, A. and Simon, H. A. *Human Problem Solving*. Englewood Cliffs, NJ: Prentice-Hall, Inc. 1972.

Newell, A., Shaw, J.C., & Simon, H.A. A variety of intelligent learning in a general problem solver. In Yovits, M.C. & Cameron, S (Eds.), *International Tracts in Computer Science and Technology and Their Application*, : Pergamon Press, 1960.

Newell, A. Production systems: Models of control structures. In Chase, W.G. (Ed.), *Visual information processing*, New York: Academic Press, 1973.

Simon, D.P. and Simon, H.A. Individual differences in solving physics problems. In Siegler, R. (Ed.), *Children's thinking: What develops?*, Hillsdale, NJ.: Lawrence Erlbaum Associates, 1978.