Title
Incremental Envisioning: The Flexible Use of Multiple Representations in Complex Problem Solving

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 12(0)

Authors
Bauer, Malcolm I.
Reiser, Brian J.

Publication Date
1990

Peer reviewed

Incremental Envisioning: The Flexible Use of Multiple Representations in Complex Problem Solving

Malcolm I. Bauer   Brian J. Reiser

Cognitive Science Laboratory
Princeton University

Abstract: In this paper we describe two properties of most psychological and AI models of scientific problem solving: they are one-pass, and feed-forward. We then discuss the results of an experiment which suggests that experts use problem solving representations more flexibly than these models suggest. We introduce the concept of *incremental envisioning* to account for this flexible behavior. Finally, we discuss the implications of this work for psychological models of scientific problem solving and for AI programs which solve problems in scientific domains.

1 Introduction

Complex problem solving often involves the use of several representations. For example, when problem solving in scientific domains, people often use formal mathematical equations, many types of diagrams, and informal conceptual intuitions in the course of reasoning. In this paper, we consider people how coordinate the use of several representations while solving a problem in a scientific domain.

2 The Traditional View

In his book, *How to Solve It*, Polya (1945) described what he believed were the four phases of problem solving: 1) understanding the problem 2) developing a plan to solve it 3) carrying out the plan, and 4) checking over the solution. Most psychological models of mathematics and science problem solving (e.g., Larkin, McDermott, Simon, & Simon, 1980; Chi, Feltovich, & Glaser, 1981; Riley, Greeno, & Heller, 1983) and programs in AI which solve problems in science (de Kleer, 1975; Skorstad & Forbus, 1989) may be considered instantiations of this view within specific domains. For example, Larkin et al. (1980) propose that physics experts solving mechanics problems will begin by sketching a picture of the described problem and selecting a set of principles. From these they construct a representation of the problem containing relevant physical entities (i.e. understanding the problem). This conceptual representation of the problem is then re-represented as a set of physics equations (constructing a plan). Finally the equations are solved algebraically (carrying out the plan). In AI, de Kleer's (1975) Newton program solves simple kinematics problems in a similar manner. Given a description of a physical scenario and a question about that scenario (e.g. "How fast will the block be traveling at point A?"), Newton first constructs an *envisionment* or representation of what will happen in the scenario. Second, it constructs a general plan of what must be done to solve the problem. Finally it accesses the relevant mathematical knowledge and solves associated equations until the desired quantity is found.

In both models, as in most subsequent models of scientific problem solving, problem solving is accomplished by proceeding sequentially through the phases described by Polya. First some sort of general conceptual model of the problem situation is constructed (in the case of Newton, this conceptual model is called an envisionment). Second, a general plan for solving the problem is constructed, embodied either as the general principles operating or as the crucial states that must be solved for. Finally, the plan cues the relevant equations, which are then solved. Polya's fourth phase, checking over the solution, is rarely included. To summarize, these models can be characterized as comprised of some or all of the following phases.

1. Construct a conceptual understanding of what is occurring in the problem.

2. Develop a plan to solve the problem. Typically this involves deciding what principles are relevant and which equations will be used.

3. Construct and solve the relevant equations.

4. Check the solution.

The models have two key properties. First, the models are *feed-forward* because information flows from earlier phases to later phases, but never in the other direction. For example, inferences made while solving equations are never passed back to earlier phases to enhance conceptual understanding.

Second, the models are *one-pass serial.* All processing within each phase is completed before the next phase is initiated. There is a block of processing using to understand the problem conceptually followed by a planning phase, which is followed by a solving phase. This property is distinct from the first one in that it is possible to have an iterative model rather than a one-pass model which would still be feed-forward. Such a model would loop through the phases several times but only passing information "forward," for example, never using phase 3 inferences in later phase 1 processing.

For fairly simple problems, there is psychological evidence that expert problem solving fits this one-pass, feed-forward model (Larkin et al., 1980). Similarly, in AI, this model has been used to solve several types of physics problems (Skorstad & Forbus, 1989; de Kleer, 1975). However, there are several reasons to believe this model does not completely characterize how experts solve all types of problems. Similarly, there are reasons to believe that this kind of AI architecture will not be able to solve many kinds of problems.

First, for complex problems, short-term memory restrictions may require people to cycle through the phases, solving pieces of the problem each time, to put together a coherent complete solution, rather than doing all the required reasoning in each phase before initiating the next phase.

Second, memory considerations aside, for difficult problems, experts may need to use several kinds of representations simultaneously to characterize a problem conceptually. This may include particular equations, theoretical models from physics, and commonsense intuitions. Roschelle and Greeno (1987) give anecdotal evidence to support this in protocols where expert physicists use both Newtonian physics models and commonsense intuitions about a physical situation to how objects will behave.

Third, de Kleer (1975) describes a class of problems he terms *indefinite* that his program is unable to solve. He claims it can't solve these problems because the program lacks flexibility. It needs to access information from different phases of its problem solving, but cannot because it is a one-pass feed-forward model. For example, certain problems may require some calculations be performed (phase 3) in order to complete conceptual understanding (phase 1).

Finally, recent work in qualitative reasoning (Sacks, 1988) has focused on interpreting formal symbolic solutions qualitatively. In many scientific disciplines, coming up with a formal symbolic solution to a problem (the result of phase 3) is not the final goal, as it is in the models above. Instead the goal is to understand what the solution means at a conceptual level (phase 1). The work in qualitative reasoning focuses on interpreting the results of phase 3 in terms of the conceptual representations utilized in phase 1. In one-pass feed-forward models, this is impossible as passing back results from phase 3 to phase 1 does not occur.

3 An Empirical Investigation of Multiple Representations

The present research is concerned with understanding expert performance in situations that experts find more challenging. The psychological models described above were typically derived from expert performance on problems requiring little effort for the experts. For the reasons above, we suggest that a one-pass feed-forward model will be inadequate to completely characterize expert performance on more complex problems. We examine expert performance on moderately difficult mechanics problems in physics. Expert performance on "easy" problems has been studied extensively in mechanics so this provides a good basis for comparison. We are interested in investigating whether one-pass feed-forward models are inadequate to explain expert behavior, and if not, what is it that experts do beyond these models in those situations.

3.1 Design and Materials

We selected four hard mechanics problems. Three of the four were taken from a review text (Wells & Slusher, 1983). The fourth was created by one of us. Simplified versions of each of the hard problems were constructed. These used the same principles necessary to solve the hard problems, but the physical scenarios in which those principles had to be used were greatly simplified. Examples are shown in Figure 1.

Subjects were graduate students drawn from the Mechanical Engineering and Physics departments at Princeton University. There were 16 subjects in total, although the analyses in this paper focus on the first 6 subjects. Each subject solved four problems, two easy and two hard. No subject was given a hard problem and its corresponding easy problem. Subjects were asked to "think out loud" while solving the problems. The sessions were videotaped. We transcribed all subjects' actions which included verbal statements, writing an equation, drawing a diagram, modifying an equation or diagram, and pointing to an equation or diagram. Because we wished to examine the transitions among the kinds of representations used, we coded protocol statements according to the kind of information used and the type of action being performed. Our analyses consider only the *information heeded* by the subject, rather than attempting to categorize the actual processes which are acting upon that information (Ericsson & Simon, 1984). Recognizing the transitions was also facilitated by the fact that, in addition to the verbal protocols, the transcripts also contained all cases where subjects modified or pointed to an equation or diagram. Protocol statements were classified into one of the eight basic categories described below:

Categorization: Subject states a category to which the problem belongs.

Rehearsal: Subject reads or re-reads problem, or restates a fact previously found.

Physical Reasoning: Subject identifies a particular physical quantity in problem, or states what occurs in the scenario, without the use of equations.

Diagram Use: Subject draws, labels, or points to a diagram.

Miscellaneous: This category includes explicit statements involving planning, and stating basic physics principles.

Mapping: Subject explicitly maps information from one representation to another.

Formal Symbolic Manipulation: Subject recalls, writes down, or performs any operation on an equation.

Qualitative Mathematical Reasoning: Subject considers an equation and reasons about it qualitatively.

Setting Goals, Hitting Impasses: Subject states a goal, or makes a statement that he or she has hit an impasse (e.g. "I'm stuck.")

Each general category was divided in several subcatgories to code the kind of action performed, if one was explicit. For example, Diagram Use had three subcategories: writing down part or all of a diagram, pointing to a diagram, and labeling part of a diagram with an equation or symbol.

3.2 Analyses

First, subjects indeed found the hard problems more difficult that the easy problems. On average it took the subjects 4.3 minutes to solve the easy problems and 18 minutes to solve the hard problems. In terms of the coding scheme, transcripts for the easy problems contained an average of 38.3 steps and for the hard problems, 122.3. All of the easy problems were solved correctly, but only 75% of the hard problems.

5.13. In the turntable arrangement shown in Fig. 5-13, block A has a mass of 0.9 kg, block B has a mass of 1.7 kg, and the blocks are 13 cm from the axis of rotation. The coefficient of static friction between the blocks, and between the blocks and the turntable, is $\mu_s = 0.1$. Consider the friction and the mass of the pulley in Fig. 5-13(a) as negligible. Find the angular speed of rotation of the turntable for which the blocks just begin to slide.

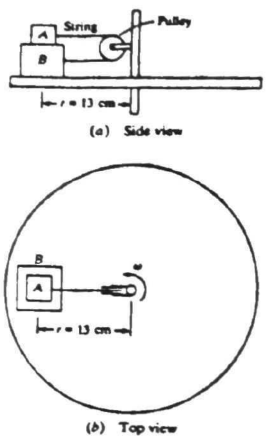

Fig. 5-13

A 4 kg block rests 0.6 meters from the center of a turntable. If the coefficient of static friction between the block and the turntable is .2, find the maximum angular velocity of the turntable for which the block will not slide.

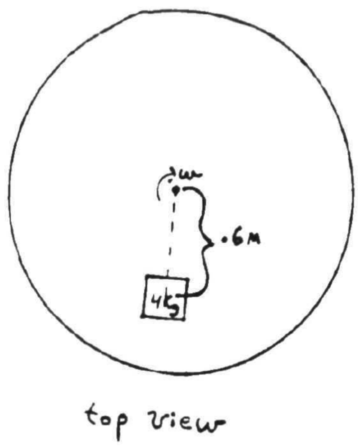

Figure 1: A hard problem (left) and its corresponding easy problem (right)

There are several analyses which can be used to evaluate the types of problem solving sequences. As a first step towards investigating these reasoning events, we divided each protocol into quarters according to the total number of codes in the protocol, and plotted the average percentage of each category of action within a quartile. The graph in Figure 2 is a quartile plot for the most prevalent codes.

The most dramatic effect is the rise in the percentage of formal symbolic manipulation, starting at 13% in the first quartile and rising fairly linearly to 60% in the fourth quartile. In addition, the physical, diagrammatic, and rehearsal codes start at around 20% in the first quartile, and slowly drop until they are all about 4% in the last quartile. This overall trend of the increase in formal manipulation and decrease in actions of with conceptual understanding is generally in keeping with the one-pass feed-forward model. However, the fact that there is even 15% formal manipulation in the first quartile and some actions of conceptual understanding in the last quartiles suggests that the one-pass feed-forward model of phases does not tell the whole story. In fact, some solving of equations (associated with phase 3) occurred before physical reasoning (associated with phase 1) thus the strict ordering of the phases is not being followed completely (the "one-pass" property).

Another way to evaluate the model is to examine the number of transitions between the formal and conceptual representations. A one-pass model would predict very few. Transitions would occur only when creating the formal equations from the conceptual representation. For the easy problems the average number of transitions was 4, while for the hard, 14. This again is evidence which supports the claim that subjects are not strictly "one-pass." Looking at the step/transition ratio, we find that this is approximately 1 transition for every 9 steps, which seems to be many more than would be expected if problem solving occurs as a long episode of conceptual work followed by an episode of planning, followed by an episode of formal symbol manipulation. Instead, it appears that subjects shift between phase 1 and phase 3 relatively frequently.

There is also evidence that subjects were not strictly feed-forward (property 1) while solving these problems as well. In many instances subjects

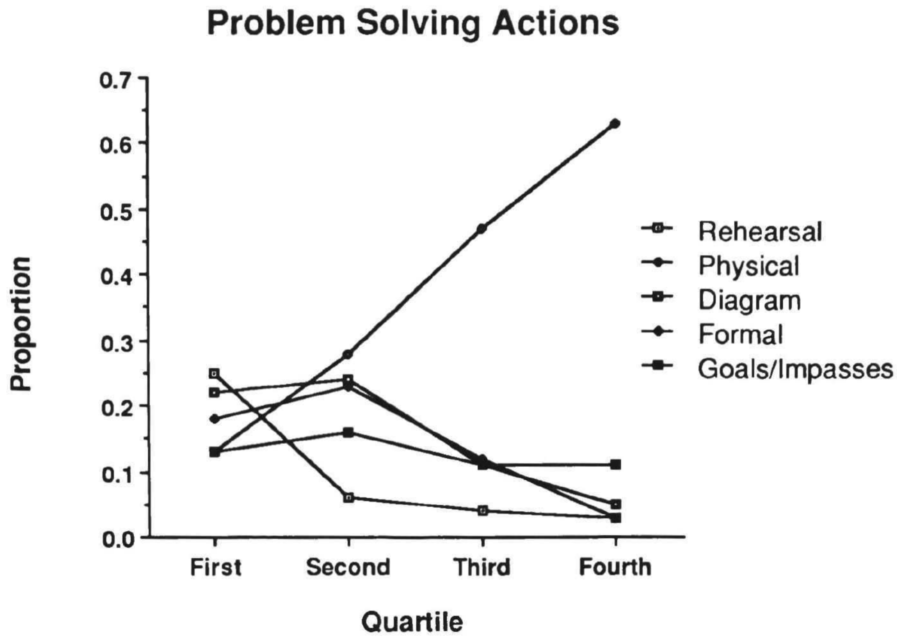

Figure 2: Quartile Plot of Problem Solving Actions

actually interpret derived equations to enrich their conceptual understanding in the course a problem. This *backward mapping* violates the feed-forward property described above. A typical example of this type of episode occurred when one subject was solving a problem about a falling rope:

> "Here's an old equation: $V^2 = {V_0}^2 + 2as$ [subject writes equation down] This is 0 [subject crosses out V_0 leaving the equation $V^2 = 2as$], so the velocity as it hits the table [points to rope in diagram] is gonna be a function of how far away from the table it was."

Here the subject has recalled an equation, applied some known quantitative information ($V_0 = 0$) and then interpreted the meaning of the expression conceptually, updating his conceptual understanding of the problem. The subject has done some work which would be classified as phase 3, but then applied it to work which would be classified as phase 1. For hard problems, this type of backward mapping occurred an average of 3.6 times per problem. This type of shift cannot occur in any model which is purely feed-forward. A more flexible model of reasoning is required.

4 Discussion

We have presented evidence that suggests that often people do not completely follow the one-pass feed-forward model. Their behavior is not strictly one-pass: they shift between understanding (phase 1) and solving (phase 3) many times in the course of reasoning about a problem. Similarly, their behavior is not strictly feed-forward: often they will use the results of solving aspects of a problem (phase 3) to enhance their conceptual understanding of the problem (phase 1) which will in turn enable them to solve other parts of the problem. Interestingly, there are different representations associated with each phase. Phase 1 is associated with building a conceptual understanding of a problem. In previous models, conceptual understanding is represented as a mental model (Roschelle & Greeno, 1987; Hegarty, Just, & Morrison, 1988), or an envisionment of what could happen (de Kleer, 1975; Roschelle & Greeno,1987). The representation used in the solution phase (phase 3) is a set of formal symbolic equations (Larkin et al, 1980). In this section, we discuss how experts coordinate these representations in problem-solving.

In most models of scientific problem solving, all of the understanding phase occurs at the beginning. Experts read the problem description, construct a representation of what's going on, and then set about to solve the problem. Instead, we propose that experts perform *incremental envisioning*: they successively refine their conceptual understanding of a problem as they work through it. There are two general ways in which subjects shift from one representation to the other. The first is a shift from envisioning to a representation associated with another phase. The second is a shift from working on the equations back to increasing the conceptual understanding of a problem.

We propose that the first kind of shift occurs largely to reduce the load of working memory. Here subjects will be thinking about what's going on in a problem and discover something relevant to one of the other phases – either an equation will be cued, an important subgoal discovered, or an important physical insight gained. The subject will then stop developing their conceptual understanding and shift to preserve that relevant piece of information, either by writing down the equation and doing some formal symbolic manipulation, by adding the physical insight to a diagram, or by stating clearly what that new goal is. In this type of shift, envisioning is momentarily halted and the important ramifications of it are propagated to other representations and preserved. Then envisioning is resumed. This is still feed-forward in that work in the other representations does not affect the envisioning, however, in contrast to the one-pass models described earlier, the process is incremental. By propagating new relevant information to other representations, the other representations are built up as the envisionment progresses. In this way, neither the entire envisionment, nor its ramifications for the other representations, have to be held in memory all at once.

To investigate this kind of incremental envisioning more precisely, we constructed the transition table containing the probability of each kind of action directly following a physical reasoning (envisioning) event (see Table 1). It is clear from the transition table that all the physical reasoning does not occur in one block, but instead shifts to other kinds of actions quite regularly. The probability of shifting to working on diagrams, shifting to working on equations, and shifting to setting goals, are all close to, if not greater than the probability of continuing with the physical reasoning.

Categorization	0.015
Rehearsal	0.073
Physical Reasoning	0.188
Diagram Use	0.272
Miscellaneous	0.019
Mapping	0.042
Formal Symbolic Manipulation	0.226
Qualitative Mathematical Reasoning	0.004
Setting Goals, Hitting Impasses	0.153

Table 1

The following examples demonstrate these feed-forward shifts:

Physical cuing Formal:

> we have the man..of mass M ..on the ladder which is a force Mg..

Here the subject began to describe the scenario (man on ladder) and shifted to writing the symbolic expression for the weight of the man (force Mg).

Physical cuing Diagram:

> that's the force needed to keep that thing going in the circle.. [subjects adds arrow to diagram]

In this example, the subject envisioned what a particular force is going to do, and then preserved that inference by adding an arrow (drawn in a circle) to a diagram.

physical cuing goals:

> ...we remain on the same circle, but I am moving on the circle and he's not..so the

> maximum time it's going to take is the time I need to make a complete circle, and that will be the worst case...so all I know have to know is the tangential speed...

Here the subject reasoned conceptually about what will happen, and then recognized a new goal to be acheived (phase 2).

In second kind of shift, the conceptual understanding phase is resumed because of an event that occurs in one of the other representations. Often, though not always, the subject returns to phase 1 to help resolve a difficulty arising in another phase. Envisioning might be resumed because of the realization that the problem cannot be solved from the current equations, for example the subject realizes that there are too few equations for the number of unknowns. This kind of shift also may be necessary to solve indefinite problems. Often it involves interpreting the results of symbolic manipulation (phase 3), as in the rope example given in the previous section. This may be done to check the validity of the derived equation, (if the interpreted equation makes sense conceptually it's more likely to be true). Occasionally, this kind of shift occurs simply to update the conceptual understanding of the problem, not to resolve any particular difficulty. However, this may cause the subject to gain new insight into the problem. Again, the rope protocol above is an example of this. Each of these cases is an example of iterative refinement as the understanding of the problem is updated with each return to envisioning. Some examples of this second type of shift are:

Too few equations:

> that's for that equation so that's two equations two unknowns.. no wait a minute two equations four unknowns<laughs> but we have two more equations.. which are the uhhh... wait a minute, [points to problem diagram] why does l interfere here?.... uhmmm.. well let's see physically what happens? If we start at this point....

Here the subject was working on the equations when he realized there are too few equations for the number of unknowns. In trying to come up with the other two equations, the subject shifted back to thinking about the problem conceptually.

Shift after finding impossible results:

> and we don't want a negative because we have to take a square root... and I screwed up the sign somewhere here.. A minus B... [points at minuses in equations] that's a minus, minus, minus, minus 2.... b....where did I lose my sign?.. [pointing at equations] T zero minus minus minus .. did I .. mess up...my forces while on this little thing?..... [points to diagram and checks equations against diagrams] where are my forces going here?..tension.... $muMg$...$muMg$ that's gotta be the opposite of that one... that one going that way...

Here the subject, in the course of solving the problem came up with the square root of a negative number. In trying to track down the error, he switched back to reasoning conceptually about what happens in the physical scenario.

To summarize, incremental envisioning occurs in two main ways. The first is done to preserve the results of envisioning and involves propagating each new result to other representations. In the second, an event in phase 3 causes envisioning to be resumed often to help resolve a difficulty arising in phase 3 processing. In this case, results from phase 3 are propagated back to phase 1.

5 Conclusions

In this paper we have described a class of scientific problem solving models called one-pass feed-forward models. We then described an experiment which suggested that expert problem solving behavior involved more than could be accounted by these models. Finally, we proposed that experts perform *incremental envisioning* as a way of describing the kinds of behaviors not characterized by one-pass feed-forward models. In this section, we briefly elaborate the implications of this work for AI and psychology.

At first glance, the first type of incremental envisioning, feed-forward shifts, may not seem useful for AI programs that solve scientific problems. We suggested this first type is used to overcome short-term memory constraints in humans. However, computers have no such limitations, so there is less need to preserve the representational ramifications of envisioning in the course of performing envisioning. Instead, the entire envisionment may be saved, and used as a whole to help in planning and solving the problem. Indeed, most scientific problem solvers work this way (de Kleer, 1975; Skorstad & Forbus, 1989). However, there are many scientific problems in which a complete envisionment need not be performed. Creating one, without reference to what must be solved for in the problem, is inefficient in these cases. Also, there are many problems for which it is impossible to construct a complete envisionment from the given information, yet the problems are solvable for the particular question being asked. In these cases, creating partial envisionments and propagating the results of envisioning during the course of envisioning is essential to deriving a solution.

The importance of the second type of incremental envisioning, backward mapping, is clear. In many cases, it may be necessary to solve part of a problem in order to complete an envisionment to solve the rest of the problem. De Kleer's indefinite problems fall into this class. Similarly, in other circumstances, it may be desirable to interpret equations conceptually. Work in AI along these lines is already being done (Sacks, 1988; see Forbus, 1988 for comprehensive review).

For psychological models, this work demonstrates that experts use representations more flexibly than has been thought. Models of expert problem solving must take into account this added flexibility. We are currently developing a computational model of incremental envisioning in problem solving. Roschelle and Greeno's (1988) relational model is a step in this direction as well.

Acknowledgments: The research reported here was supported in part by contract MDA903-87-K-0652 from the Army Research Institute, by a research grant from the James S. McDonnell Foundation, and by a grant-in-aid of research from the Sigma Xi, The Scientific Research Society. The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the U.S. Army or the McDonnell Foundation.

References

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

de Kleer, J. (1975). *Qualitative and quantitative knowledge in classical mechanics.* Master's thesis, Department of Electrical Engineering, Massachusetts Institute of Technology, Cambridge, MA.

Ericsson, K. A., & Simon, H. A. (1984). *Protocol Analysis.* Cambridge, MA: MIT Press.

Forbus, K. D. (1988). Qualitative physics: Past, present, and future. In H. E. Shrobe (Ed.), *Exploring Artificial Intelligence*, San Mateo, CA: Morgan Kaufmann.

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

Polya, G. (1945). *How to Solve It.* Princeton, NJ: Princeton University Press.

Riley, M. S., Greeon, J. G., & Heller, J. I. (1983). Development of children's problem solving ability in artihmetic. In H. P. Ginsberg (Ed.), *The development of mathematical thinking*, NY: Academic Press.

Roschelle, J., & Greeno, J. G. (1987). mental models in expert physic reasoning. *Report No. GK-2*, University of California, Berkeley.

Sacks, E. (1988). Qualitative analysis by piecewise linear approximation. *Artificial Intelligence in Engineering*, 3, 3.

Skorstad, G. & Forbus, K. (1989). Qualitative and Quantitative Reasoning about Thermodynamics. *Proceedings of the 11th Annual Conference of the Cognitive Science Society*, Ann Arbor, Michigan: Lawrence Erlbaum.

Wells, D. A., & Slusher, H. S. (1983) *Physics for Engineering and Science.* U.S.A: McGraw-Hill.