RC 15861 (#69354) 3/22/90
Cognition Science 8 pages

Research Report

Promoting Abstract Strategies in Algebra Word Problem Solving

Mark K. Singley

IBM Research Division
T. J. Watson Research Center
Yorktown Heights, NY 10598

John R. Anderson and Jody S. Gevins

Department of Psychology
Carnegie-Mellon University
Pittsburgh, PA 15213

IBM Research Division
Almaden • T J Watson • Tokyo • Zurich

Promoting Abstract Strategies in Algebra Word Problem Solving

Mark K. Singley

User Interface Institute
IBM T.J. Watson Research Center
Box 704
Yorktown Heights, NY 10598

John R. Anderson

Department of Psychology
Carnegie-Mellon University
Pittsburgh, PA 15213

Jody S. Gevins

Department of Psychology
Carnegie-Mellon University
Pittsburgh, PA 15213

We describe an experiment comparing three abstract strategies for algebra word problem solving in terms of both learning and transfer. The skeletal strategy is a working-forwards, expert strategy that most closely approximates standard textbook instruction. The means-ends strategy is a variant of means-ends analysis that begins with the goal and works backwards. The diagram strategy makes heavy use of a diagram to control the generation and combination of constraints. We found that the diagram strategy was superior to the skeletal strategy in terms of learning, and that the means-ends strategy was superior in terms of transfer. Our approach stresses the importance of explicit task analysis as a prerequisite to promoting transfer.

Promoting Abstract Strategies in Algebra Word Problem Solving

Mark K. Singley
User Interface Institute
IBM T.J. Watson Research Center
Yorktown Heights, NY 10598

John R. Anderson, Jody S. Gevins
Department of Psychology
Carnegie-Mellon University
Pittsburgh, PA 15213

Abstract

We describe an experiment comparing three abstract strategies for algebra word problem solving in terms of both learning and transfer. The skeletal strategy is a working-forwards, expert strategy that most closely approximates standard textbook instruction. The means-ends strategy is a variant of means-ends analysis that begins with the goal and works backwards. The diagram strategy makes heavy use of a diagram to control the generation and combination of constraints. We found that the diagram strategy was superior to the skeletal strategy in terms of learning, and that the means-ends strategy was superior in terms of transfer. Our approach stresses the importance of explicit task analysis as a prerequisite to promoting transfer.

Introduction

Lave (1988) has argued that our formal educational system is predicated on a mistaken belief in the transfer of abstract knowledge from one situation (the classroom) to another (e.g. the home, the job, the marketplace). Indeed, nearly a century of research in psychology has generated a depressing lack of evidence for the notion of general transfer. Thorndike, the first psychologist to study transfer systematically, showed that training in formal subjects like Latin and geometry had no effect on performance in reasoning tasks relative to training in more mundane subjects like bookkeeping and shopwork (Thorndike, 1924). More modern attempts at showing transfer have also been discouraging. First, there are the classic studies of Hayes and Simon (1977) showing little transfer between isomorphs of the tower of Hanoi puzzle. Continuing in this vein, Jeffries (1978) found little transfer between missionaires and cannibals and waterjug problems, although according to her simulation model, both involved use of a means-ends analysis strategy. Gick and Holyoak (1983) observed little transfer between isomorphs of Duncker's radiation problem. Aside from these laboratory studies of puzzle problems, studies of transfer in more realistic school settings have also been largely unsuccessful. Most relevant to our current interests is a study by Post and Brennan (1976), who trained students for several weeks on a general heuristic procedure for solving word problems. Their instructions included such things as "determine what is given" and "check your result." On a problem solving posttest, the performance of the trained subjects was no better than that of a control group. Other attempts at teaching general skills for mathematical problem solving have been largely negative. Finally, despite early claims that computer programming would emerge as the mental discipline that would revolutionize children's thinking, empirical studies have shown little benefit of learning to program on general problem solving abilities (for a useful review, see Mayer et al. (1986)).

Against this backdrop of failed experiments, and in response to the apparent failure of our educational system, we see the recent emergence of the situated learning movement. The situated learners deny the existence of abstract cognitive tools which can be applied in diverse situations (Lave, 1988) and are thus seemingly denying the possibility of transfer entirely. Although the movement is still defining itself, it seems clear that this is a radical departure from the cognitive tradition. Others, not willing to abandon the notion of transfer entirely, have remained squarely in the cognitive tradition but have adopted a more realistic view of what kinds of transfer are possible. There is a new realism in transfer research which is informed by recent advances in the understanding of task structure and the mechanisms of transfer (Singley and Anderson, 1989).

The New Realism in Transfer Research

In the last decade, some exciting progress has been made in codifying and teaching moderately general skills which represents a turnaround of sorts of the historical trend. First, in the domain of mathematics, Schoenfeld (1985), after numerous failed attempts by other researchers, has been moderately successful in teaching abstract heuristic strategies to college students. His basic strategy has been to take the heuristics of Polya (1957) as a starting point and to elaborate them with more specific rules and guidelines. Reflecting on his success, Schoenfeld concludes that "Heuristics are comquantities onto elements of the problem situation. If a diagram is being used, this amounts to labelling elements of the diagram with the givens, and finally, labelling the goal of the problem with a variable. Needless to say, the correct construction of the diagram does not logically imply the correct labelling of the diagram.

Step 3: Generate constraints. One important feature of the more difficult algebra word problems is that much of the information required for their solution is not stated explicitly in the problem text. Important quantitative relationships are missing from the problem statement and must either be derived from the problem situation or retrieved from the problem schema. For example, in the picture frame problem, the text simply provides a series of variable assignments. Not stated explicitly are the three primitive constraints that are critical for solution:

- The area of the picture equals the length of the picture times the width of the picture.
- The length of the picture equals the longer dimension of the frame minus two times the width of the frame.
- The width of the picture equals the shorter dimension of the frame minus two times the width of the frame.

When properly combined, these three primitive constraints are sufficient for the solution of this problem. However, many more constraints can be derived from the problem situation, such as the area of the picture equals the total area minus the area of the frame, and the total area equals the longer dimension of the frame times the shorter dimension of the frame. Before solving the problem, it is impossible to know which of these constraints will be useful.

Step 4: Combine constraints. A very important component of word problem solving that has been largely overlooked is search. Singley (1986) found that, in the domain of calculus related rates word problems, the representation phase was not the most difficult to master. Rather, the search through the system of equations took the most time and was the locus of most errors.

According to our task analysis, algebra word problems involve the solution of systems of equations, too. For example, in our picture frame problem, the three constraints presented above must be combined in such a way that the goal of the problem is expressed in terms of known quantities. The search problem is that there are more than just these three constraints that are potentially relevant and there are untold ways of combining the constraints. Indeed, it is rather impressive that beginning algebra students can solve these problems at all, given that this search process has been largely hidden from them in standard instructional texts. Given that students have no systematic method for dealing with search, it is no wonder that performance in word problem solving is often brittle and shows little transfer to new problem types. Given the lack of strategic instruction, students must learn particular solution patterns for particular problems.

Step 5: Solve the final equation. Once an equation is written that incorporates all the necessary constraints, the problem is finished except for algebraic solution. The details of this process have been explored in other work in our laboratory. We have concluded that algebraic solution is largely decoupled from the other components of word problem solving skill and can be profitably ignored in a word problem curriculum (Singley, 1986).

In sum, our task analysis claims that word problem solving is difficult because it is a knowledge intensive task that involves a number of difficult components, not the least of which is the search through the space of constraints (steps 3 and 4). Indeed, data from an initial study we conducted showed that students make ten times as many errors on word problems that involve systems of equations, and that the difficulty of a problem is directly related to the number of primitive equations that make up the solution. This is consistent with the results of the National Assessment of Educational Progress, a national survey of mathematical problem solving abilities (Carpenter et al., 1980). The survey showed that performance on one-step word problems was quite good, but that performance on multi-step problems was quite bad.

These results all suggest that students need help managing their search processes, and this is what we chose to focus on. Through our task analysis, we defined three distinct strategies for the generation and combination of constraints:

- *Skeletal strategy.* The skeletal strategy is a working forwards strategy that most closely approximates the kinds of example solutions presented in textbooks (e.g. Foerster (1984)). In the skeletal strategy, one first writes the primitive equation that is in some sense key to solving the problem. This primitive equation in most cases does not contain the goal variable, and its presentation is entirely unmotivated except for the fact that it "works." For example, to solve our example problem, one would first recognize that the area of the picture is equal to the length of the picture times the width of the picture. One would then realize that the length of the picture equals $20 - 2x$, and the width of the picture equals $14 - 2x$, and suddenly one has an equation that is solveable for x:

  $$(20 - 2x)(14 - 2x) = 160$$

  The skeletal strategy gets its name from the fact that the first equation written is the easily-identified skeleton of the final equation. It is easily identified because the skeletal equation (which again typically does not contain the goal variable) provides for an equation tree that has minimal depth and maximal breadth. Thus, the skeletal strategy generates final equations whose structures are easily understood. This may be why it is popular in terms of exposition: it provides the most easily understood post-mortem of how a problem was solved. However, the strategy is not generative in that particular skeletal equations and subsequent solution steps are largely unmotivated. The search problem is totally ignored; learning is reduced to the more-or-less rote acquisition of specific problem-solution patterns or schemas.
- *Means-ends strategy.* Means-ends analysis is a working backwards strategy that has been well-documented as a novice strategy for solving systems of equations (Larkin

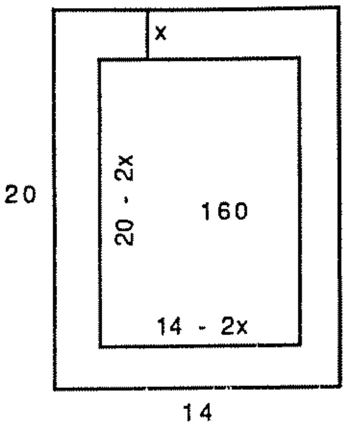

Figure 2. Labelling the picture frame using the diagram strategy.

	pennies	nickels	coins
number	x	18-x	18
value per coin	.01	.05	
total value	.01x	.05 (18 -x)	.38

Figure 3. Labelling a table using the diagram strategy.

on a keypad. The system is very flexible in terms of the labels it accepts. For example, the following are just a few of the valid labels for the leftmost vertical line in Figure 2: 1) 20, 2) y, 3) $2x + l$ (where l is the length of the picture), 4) $(160 + f)/14$ (where f is the area of the frame), and 5) $(l * (14 - 2x) + f)/(w + 2x)$ (where w is the width of the picture). If a student asks for help and all known quantities are not posted on the diagram, the tutor will help the student post the missing quantities through a succession of hints. The tutor does no work directly for the student.

Step 3: Generate constraints. The tutor allows students to write primitive (or composed) constraints at any time during the solution of the problem. There are two ways to write constraints: either by writing a full-fledged equation (in support of the means-ends or skeletal strategies) or by labelling the diagram with an expression (in support of the diagram strategy). In Figure 4, we see that the student has not labelled the diagram with expressions, but has instead written three primitive equations: $l = 20 - 2x$, $l = 160/w$, and $w = 14 - 2x$. If a student asks for help and not all the necessary constraints have been generated, the tutor recommends different things depending on the strategy being used:

- *Skeletal strategy.* If no constraints at all have been generated, the tutor recommends writing the so-called skeletal equation. This is followed by recommendations to perform the predetermined substitutions required by the skeletal strategy.
- *Means-ends strategy.* The tutor first recommends writing the primitive equation that contains the goal variable and the fewest number of unknowns. Once such an equation exists, the tutor recommends that students write additional primitive equations that can be used to replace the remaining unknown variables with either known values or expressions containing known values and the goal variable.
- *Diagram strategy.* The tutor recommends labelling the diagram with an expression involving only known values and the goal variable (no new variables are recommended). At any point in the solution process, a very limited number of objects can be labelled in this way. Thus, constraints are generated in an order that is determined by the givens of the problem, which gives the diagram strategy a decidedly working-forward flavor.

If an equation or expression is written that contains one or more variables, the tutor tries to disambiguate the meaning of the variables on the basis of context and usage. In truly ambiguous cases (e.g. in the equation $a + b = c$, a and b are ambiguous), the tutor brings up a menu containing English phrases to allow the student to select the intended meaning. The meanings of all bound variables are displayed in a separate window.

Step 4: Combine constraints. To support the combination of constraints, the tutor supplies a substitution operator which can be used to eliminate an unwanted variable in the target equation by replacing it with either its numeric value or an equivalent symbolic expression. For example, given the equations $160 = l * w$ and $l = 20 - 2x$, the substitution operator returns $160 = (20 - 2x) * w$. As outlined above, the tutor offers help on all three strategies for constraint combination.

Step 5: Solving the final equation. The tutor supplies students with a powerful symbolic calculator that solves any equation containing a single variable. The focus in the tutor, then, is on deriving a solvable equation, not on solving it.

Initial Evaluation

We performed an initial evaluation of the system to assess the validity of the tutor as a pedagogical tool. Seventeen beginning algebra students (aged 13 years) spent an average of 10 hours solving eight different kinds of problems with the syslems represented near and far transfer problems. For example, if a subject were trained on coin mixture problems, then the liquid mixture problems would represent near transfer problems, and the two types of frame problems would represent far transfer problems. Near transfer is defined by the sharing of some (but not all) of the underlying primitive equations that define a problem schema. Far transfer is defined by disjoint sets of underlying primitive equations.

Procedure

Subjects first took a short pretest involving algebraic manipulation, e.g. solving linear equations and simplifying expressions. They then read short (about 10 pages) booklets describing one of the three strategies. Each booklet demonstrated the use of the strategy on a single problem. The experimenter then demonstrated the use of the tutor on two sample problems. Both demonstrations involved application of the strategy. Then subjects solved problems drawn from one of the four schemas using the tutor until a training criterion was achieved: two consecutive problems with a total of two or fewer errors and no help requests. Then subjects saw a sequence of transfer problems. (In the transfer phase of the experiment, the strict enforcement of strategies was lifted.) First, subjects solved two more problems drawn from the training schema, but involving different configurations of givens and goals than they saw during training (within-schema transfer). Next, subjects solved two near-transfer problems. Finally, subjects solved four far-transfer problems, two of each type. Prior to solving one of the pairs, subjects were trained on a series of 10 one-step problems that were intended to familiarize them with the underlying primitive equations. This special training was counterbalanced across serial position.

Results

Figure 5 shows the mean number of trials to criterion for each of the six conditions. An analysis of covariance using pretest score as a covariate showed a marginal main effect for type of strategy ($F = 2.7$, $df = 2$, $p < .1$). A post-hoc t-test revealed that the diagram strategy took significantly fewer trials to reach criterion than the skeletal strategy ($t = 2.32$, $df = 16$, $p < .05$). Besides the marginal main effect, the analysis of covariance yielded a marginal interaction between strategy and level of enforcement ($F = 2.3$, $df = 2$, $p < .15$). At the risk of overinterpretation, this result suggests that, whereas generative strategies like the means-ends and diagram strategies benefit from loose enforcement, the more routinized skeletal strategy does not. Other analyses of covariance revealed that loose enforcement results in fewer errors per problem during training (2.8 vs. 4.1; $F = 4.5$, $df = 1$, $p < .05$) and marginally more help requests (2.3 vs. .8; $F = 3.9$, $df = 1$, $p < .07$). Comparisons of the strategies on these dependent measures yielded no significant effects.

Figure 6 shows comparisons of the strategies on the transfer problems in terms of help requests per problem. Other dependant measures, such as time per problem and errors per problem, yielded no significant effects. We have collapsed across levels of enforcement, as this training factor also had no effect on transfer performance. What we see in terms of help requests per problem, however, is the emergence of the means-ends strategy as the superior strategy on both near and far transfer problems. Analyses of covariance and subsequent t-tests revealed that subjects taught the means-ends strategy asked for help significantly fewer times than subjects taught the skeletal strategy in both the near transfer ($t = 3.2$, $df = 16$, $p < .05$) and far transfer ($t = 2.7$, $df = 16$, $p < .05$) problems. Exposure to one-step problems drawn from the far-transfer schema prior to exposure to full-blown far-transfer problems had no measurable effect on performance.

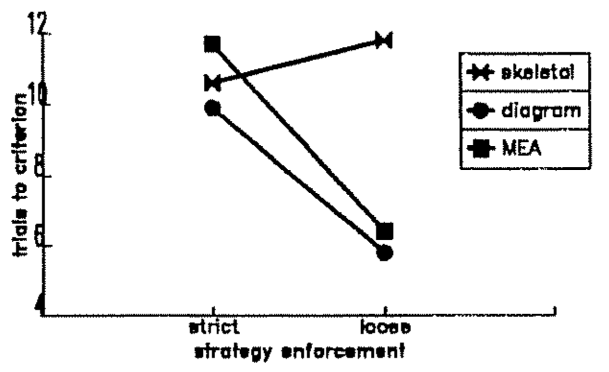

Figure 5. Trials to criterion.

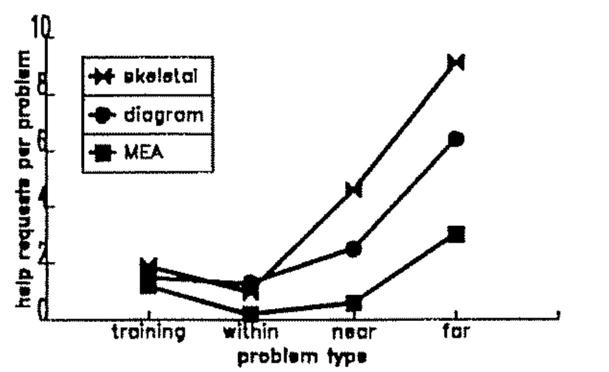

Figure 6. Help requests per problem.

General Discussion

We have demonstrated the superiority of the diagram and means-ends strategies over a formalization of the standard textbook strategy in terms of both learning and transfer. The diagram strategy leads to the fastest criterion performance, and the means-ends strategy leads to the fewest requests for help in both near and far transfer situations. It is interesting

Copies may be requested from:

IBM Thomas J. Watson Research Center
Distribution Services F-11 Stormytown
Post Office Box 218
Yorktown Heights, New York 10598