learning environments, as well as for suggesting modifications to intelligent tutoring systems to make them more like the individualized teaching systems they have the potential to be.

References

Kyllonen, P.C., & Christal, R.E. (in press). Cognitive modeling of learning abilities: A status report of LAMP. In R. Dillon &J.W. Pellegrino (Eds.), *Testing: Theoretical and applied issues.* San Francisco: Freeman.

Regian, J.W., & Schneider, W. (1986). Assessment procedures for predicting and optimizing skill acquisition. Paper presented at ETS conference on diagnostic monitoring, July.

Schneider, W., & Shiffrin, R.M. (1977). Controlled and automatic human information processing: I. Detection, search, and attention. *Psychological Review*, 84, 1-66.

Watson, D., Clark, L.A., & Tellegen, A. (1988). Development and validation of brief measures of positive and negative affect: The PANAS scales. *Journal of Personality and Social Psychology*, 54, (6), 1063 - 1070.

The Algebra Word Problem Tutor

Mark K. Singley, John R. Anderson, Jody S. Gevins, Donn Hoffman
Department of Psychology
Carnegie Mellon University
Pittsburgh PA 15213

Abstract

As an extension of our recent work on tutoring basic algebraic manipulation skills (Lewis, Milson, & Anderson, in press), we have built an algebra word problem tutor that provides mathematical modelling tools for students and also provides corrective help and feedback on errors as students solve problems. The tutor is designed to cover all problem types typically found in a first-year algebra course, and is easily extended to word problems in other domains, such as physics and calculus. The purpose of our project is threefold: 1) to develop an educational tool that will help students learn what has been traditionally regarded as the most difficult part of algebra, 2) to explore the psychological question of what makes word problem solving so difficult, and 3) to extend the model tracing paradigm (Reiser, Anderson & Farrell, 1985) to a new domain that presents some interesting complications. In this report, we describe our task analysis of algebra word problem solving, the subsequent design of our system, and finally the results of a preliminary evaluation involving beginning algebra students.

Introduction

Algebra word problems are notoriously difficult to solve for the average high school student. In one study conducted in the United States (see Mayer, 1982), only 17% of students aged 17 could solve this problem:

> Lemonade costs 95 cents for one 56 ounce bottle. At the school fair, Bob sold cups holding 8 ounces for 20 cents each. How much money did the school make on each bottle?

Needless to say, such results are quite troublesome to mathematics educators, especially so in the current educational climate where problem solving ability has gained greater prominence as a pedagogical goal. The challenge to educators is to understand the sources of difficulty in word problem solving and to somehow overcome them. Our intention is to apply what has been called the *model tracing* paradigm (Reiser, Anderson, & Farrell, 1985) to the tutoring of algebra word problems.

Task Analysis of Algebra Word Problem Solving

The first step in building a model-tracing tutor is to perform a thorough task analysis to determine the component steps of problem solving, as well as some overall strategy for generating and sequencing those components. As we performed our preliminary task analysis of algebra word problem solving, and examined the analyses of others (Paige & Simon, 1966; Hinsley, Hayes, & Simon, 1977; Shalin & Bee, 1985; Hall et al., 1986), we began to see why this task is so difficult. Solving a word problem is like assembling a jigsaw puzzle, where the pieces are yet to be defined. Students must first derive a representation of the elements of the problem, and then do a considerable amount of search to put the elements together. One of the many reasons why word problem solving is so difficult is that in most cases these two phases of problem solving (representation and search for solution) are nearly independent. Mistakes in the representation phase can lead to impasses in the solution phase, but rarely do these impasses reveal anything about the nature of the underlying misconception.

We will develop our task analysis in the context of an example problem, the classic picture frame problem:

> A picture frame measures 20 cm by 14 cm. 160 sq cm of picture shows. What is the width of the frame?

Generally, to solve an algebra word problem, one must bring to bear a set of constraints which serve to express the answer of the problem in terms of known quantities (we view equations as quantitative constraints among variables). We will examine the five steps a hypothetical ideal student might take in solving such a problem. In our description, the steps are presented in a certain well-defined order, where one stage of processing finishes before the next begins. In fact, the students we have observed interleave these components a fair amount.

Step 1: Define the problem situation. As a first step in solving this problem, it is helpful to represent the qualitative relationships between the various elements of the problem. This set of qualitative relationships (e.g. the picture is centered inside the frame, the frame has equal width on all four sides, etc.) has been called the *problem situation* (Reusser, 1985; Hall et al., 1986). In line with the recent work of Larkin & Simon (1987), we view a pictorial diagram as being a powerful and compact representation of this set of relationships. Once constructed, the diagram not only provides useful cues for generating constraints but also acts as a kind of search control mechanism for combining constraints.

Step 2: Map known quantities onto the problem situation. Once the qualitative relationships between the elements have been established, it is then necessary to map the known quantities onto elements of the problem situation. If a diagram is being used, this amounts to labeling 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 labeling 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. The amazing thing is that beginning algebra students can solve these problems at all, given that this search process has been largely hidden from them. 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. Students must learn particular solution patterns for particular problems.

Through our task analysis, we have defined two possible strategies for the generation and combination of constraints. One is means-ends analysis, which has been well-documented as a novice strategy for solving systems of equations (Larkin, McDermott, Simon, & Simon, 1980). In means-ends analysis, one first writes an equation that contains the goal variable. Given this equation, subgoals are set to find values for the remaining unknowns. Additional constraints are generated which contain the unknowns and are substituted into the equation containing the goal. The problem is solved when the goal variable is stated in terms of knowns. Needless to say, this strategy can lead to protracted episodes of substitutions and quite a bit of search that has a somewhat mindless quality. In addition, many extra variables are introduced which tend to burden working memory. Sweller, Mawer, & Ward (1983) have claimed that means-ends analysis interferes with learning because it makes heavy demands on working memory and leaves no capacity for the learning mechanisms. Another strategy we have defined makes direct use of the diagram. This so-called diagram strategy involves the repeated labeling of objects in the diagram using only the goal variable and known quantities until an equation can be written that makes use of the labels and contains the goal variable. For example, Figure 1 shows the labeling of a diagram for the picture frame problem using this strategy. At this point, an equation can be written that states the area of the picture (160) in terms of the length of the picture ($20 - 2x$) and the width of the picture ($14 - 2x$). The virtue of this strategy is that it introduces no extra variables and thus places a smaller burden on working memory. In addition, the diagram provides at least heuristic guidance about which object to label next, so the search problem is somewhat mitigated. The drawback of this strategy is that the termination point is unclear. One must continually monitor the diagram to see when the final equation can be written.

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 our other algebra work. We feel 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 (Heid & Kunkle, 1988; Singley, 1988).

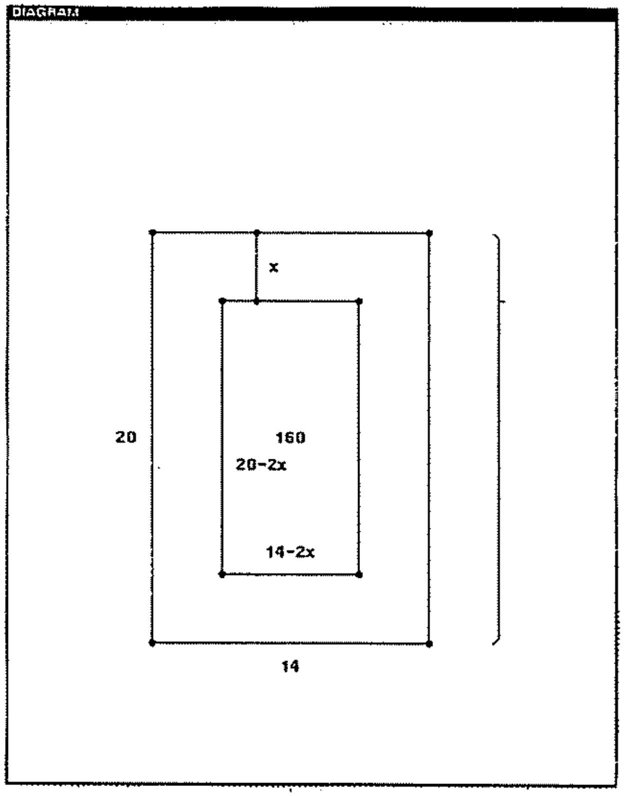

Figure 1 Labeling the diagram using the diagram strategy

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.

Design of the Tutor

The Algebra Word Problem Tutor is designed to support each of the five components of word problem solving skill outlined above. Our intention has been to extend what has been called the *model tracing* paradigm (Reiser, Anderson, & Farrell, 1985) to the tutoring of algebra word problems. The model tracing paradigm is based primarily on two premises: 1) the majority of the learning that takes place in the acquisition of cognitive skill occurs during the solving of problems and not in the perusal of textbooks, and 2) one needs to have an accurate cognitive process model of the student solving problems if one wants to tutor the student effectively. The strength of the model tracing

paradigm is that, when the student makes a mistake, the tutor can localize the problem to a single primitive action, which can be interpreted as a piece of a larger plan.

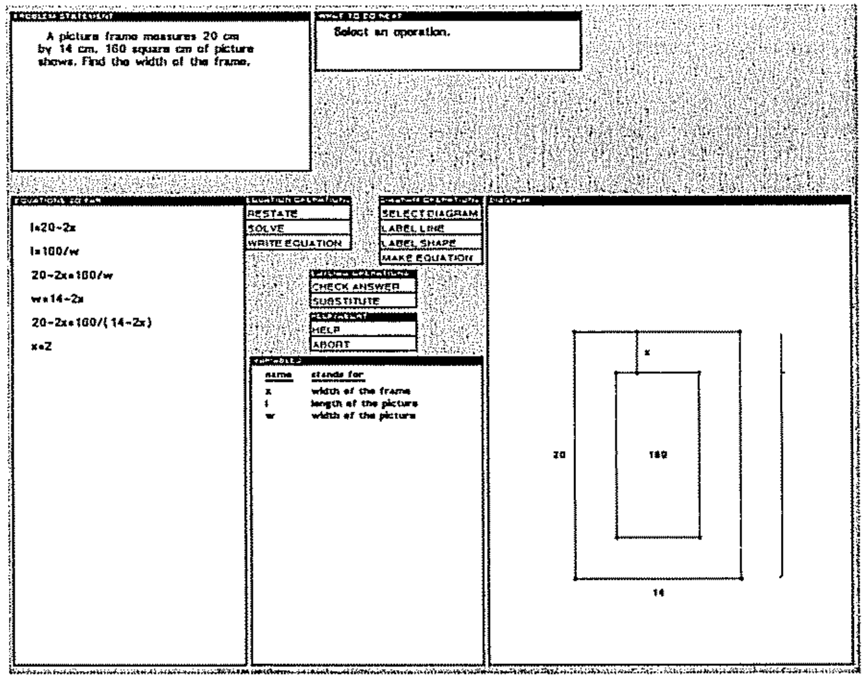

Figure 2 Interface of the Algebra Word Problem Tutor

The tutoring of algebra word problems presents a particularly interesting set of challenges to the model tracing approach. Most notably, in comparison with other mathematical skills, such as simple algebraic manipulation, word problems have a greater cognitive component, i.e. more operations involving the manipulation of internal representations and less involving the manipulation of external objects. The challenge, then, is to somehow reify the representations and goal structures necessary for solution. As mentioned above, the problem situation can sometimes be represented in a fairly straightforward way with pictorial diagrams. In other cases, however, the representation from which equations can be drawn is less obvious.

Another challenge is that, like algebraic manipulation and unlike domains like LISP programming and geometry, the emphasis in word problem solving is more on the solution itself and less on the path to the solution. This means that it is admissible and perhaps even advisable to combine or skip steps. This means that a tutor for algebra word problems must be extremely flexible in terms of the range of solutions recognized. Another requirement is that the tutor must adapt to the grain size defined by the student. For example, the quantitative constraints outlined earlier may be generated one at a time and combined in a piecemeal fashion, or all the constraints may be generated at once and written as a single large equation. A recent evaluation of the algebra manipulation tutor showed that students do better when allowed to define their own path as well as the grain size of the solution.

Figure 2 shows the interface of the tutor following solution of the picture frame problem using the means-ends strategy. Although the means-ends strategy is shown, the tutor is very flexible and allows the student to pursue any desired strategy. We now describe the tools provided by the tutor for the five problem solving components.

Step 1: Define the problem situation. To represent the situation of the problem, the tutor currently allows students to select from a menu of pictorial icons which contains not only the correct icon but also a set of carefully constructed foils. The foils are constructed by systematically violating each of the qualitative relationships that as a whole define the problem situation. For example, Figure 3 shows the icon menu presented to students solving the picture frame problem. In panels one through five, the following qualitative relationships are violated: 1) the frame has four sides, 2) the frame has equal width on all sides, 3) the picture has a rectangular shape, 4) the picture is enclosed by the frame, and 5) the picture is smaller than the frame. The virtue of this formal analysis of foils is that, if a student makes a wrong choice, the tutor can offer specific remediation based on the qualitative relationship that is violated. One drawback of such a scheme is that it is recognition-based and does not require students to build the diagram for themselves. Future work is directed toward allowing students to generate the diagram for themselves using a high-level icon-based language.

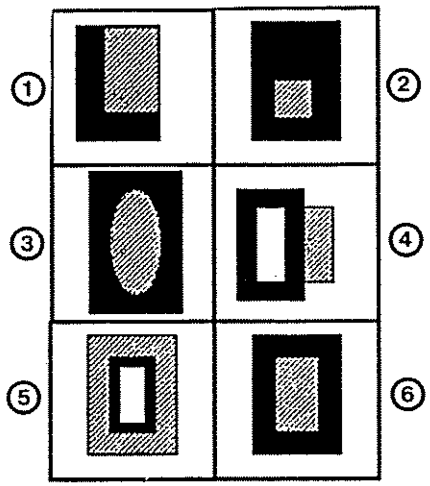

Figure 3 Icon menu for the picture frame problem

Step 2: Map known quantities onto the problem situation. Once the diagram is displayed, students are encouraged to label its lines and shapes with either numbers, variables, or expressions. This is done simply by first telling the system the kind of object to be labeled (line or shape), selecting the particular object with the mouse, and entering the new label 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 1: 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 (often in support of the means-ends strategy) or by labeling the diagram with an expression (often in support of the diagram strategy). In Figure 2, we see that the student has not labeled 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: 1) In the case of the means-ends strategy, if no constraints at all have been generated, the tutor 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. 2) In the case of the diagram strategy, the tutor recommends labeling 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 labeled 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 both the means-ends and diagram 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 are currently engaged in an initial evaluation of the system. This evaluation has two independent yet complementary goals: first, to demonstrate the validity of the tutor as a pedagogical tool, and second, to use that tool to perform basic research on algebra word problem solving. We have analyzed data from four beginning algebra students (aged 13 years) who spent an average of 9.4 hours solving eight different kinds of problems with the system (triangle, rectangular frame, two-sided frame, round-trip motion, liquid mixture, dry mixture, gambling, coin). These students received no formal instruction on word problem solving outside of the tutor. A paper and pencil test involving all eight problem types was given before and after exposure to the system, and performance was measured in terms of both getting the right answer and writing an equation that could be solved to get the right answer. Interestingly, the tutor had little effect on performance in terms of getting the right answer, with students exhibiting means of 10% and 16% on the pretest and posttest, respectively. This is hardly surprising, given that the tutor really did not cover algebraic solution. However, in terms of writing a solvable equation, performance jumped from a mean of 0% on the pretest to 50% on the posttest. With further enhancements of the system, we hope to achieve virtually perfect posttest performance.

As mentioned previously, a classic psychological question involves determining the locus of difficulty in word problems. The algebra word problem tutor offers an exciting new opportunity to study this question. Specifically, it is possible to partition the various components of word problem solving skill and measure each separately. We have determined error frequencies for our four students working with the tutor on the three middle problem solving components described earlier (steps two through four). (The first component was not measured because the diagram selection facility was not in place at the time of the evaluation, and the last component was not measured because algebraic solution is performed directly by the system.) We have found that our students average .4, .5, and .5 errors per problem on the components of mapping known quantities, generating constraints, and combining constraints, respectively. This suggests that word problem solving is truly a complex skill and that the difficulty cannot be localized to a specific component.

Conclusion

Our project involves the development and initial testing of an algebra word problem tutor that we feel has great potential as an educational and scientific tool. Our experimental work is geared toward both validating the system in terms of educational effectiveness and exploring some fundamental psychological issues surrounding algebra word problem solving.

References

Hall, R., Kibler, D., Wenger, E., & Truxaw, C. (1986). *Exploring the episodic structure of algebra story problem solving* (Technical Report No. 86-24). Irvine: University of California.

Held, K., & Kunkle, D. (1988). Computer-generated tables: Tools for concept development in elementary algebra. In *Ideas for algebra: National Council of Teachers of Mathematics yearbook.* Reston, Va.: NCTM.

Hinsley, D., Hayes, J. R., & Simon, H. A. (1977). From words to equations. In P. Carpenter and M. Just (Eds.), *Cognitive processes in comprehension.* Hillsdale, N.J.: Lawrence Erlbaum Associates.

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

Larkin, J., & Simon, H. A. (1987). Why a diagram is (sometimes) worth ten thousand words. *Cognitive Science, 11,* 65-99.

Lewis, M., Milson, R. & Anderson, J. R. (in press). The teacher's apprentice: Designing an intelligent authoring system for high school mathematics. To appear in *Proceedings of Educational Testing Service Conference on Artificial Intelligence.* Princeton.

Mayer, R. E. (1982). Memory for algebra story problems. *Journal of Educational Psychology, 74,* 199-216.

Paige, J. M., & Simon, H. A. (1966). Cognitive processes in solving algebra word problems. In B. Kleinmuntz (Ed.), *Problem solving: Research, method, and theory.* New York: Wiley.

Reiser, B. J., Anderson, J. R., & Farrell, R. (1985). Dynamic student modelling in an intelligent tutor for LISP programming. In *Proceedings of the International Joint Conference on Artificial Intelligence.* Los Angeles.

Reusser, K. (1985). *From situation to equation: On formulation, understanding and solving situation problems* (Technical Report No. 143). Boulder: University of Colorado.

Shalin, V., & Bee, N. V. (1985). *Analysis of the semantic structure of a domain of word problems* (Technical Report). Pittsburgh: University of Pittsburgh.Learning Research and Development Center.

Singley, M. K. (1986). Developing models of skill acquisition in the context of intelligent tutoring systems. Unpublished doctoral dissertation. Pittsburgh: Carnegie-Mellon University.

Singley, M. K. (1988). The relationship between operator selection and application in mathematics problem solving. In *Proceedings of ITS '88 Conference.* New York: Association for Computing Machinery.

Sweller, J., Mawer, R., & Ward, M. (1983). Development of expertise in mathematical problem solving. *Journal of Experimental Psychology: General, 112,* 639-661.