Richard Catrambor

School of Psychology
Georgia Institute of Technology
Atlanta, GA 30332
rc7@prism.gatech.edu

Abstract

A proposal is made for representing the knowledge learners acquire from examples in terms of subgoals and methods. Furthermore, it is suggested that test problems can also be represented in terms of the subgoals and methods needed to solve them. Manipulations of examples can influence the particular subgoals and methods learned. Thus, transfer can be predicted by the overlap in the learned subgoals and methods and those required to solve a novel problem. A subgoal is an unknown entity, numerical or conceptual, that needs to be found in order to achieve a higher-level goal of a problem. A method is a series of steps for achieving a particular subgoal. The experiment presented here suggests that elaborations in example solutions that emphasize subgoals may be an efficient way of helping a learner to recognize and achieve those subgoals in a novel problem, that is, to improve transfer. It is argued that conceptualizing problem-solving knowledge in terms of subgoals and methods is a psychologically plausible approach for predicting transfer and has implications for teaching and design of examples.

Introduction

This paper makes two basic claims. One, people can learn subgoals for solving problems in a particular domain from the examples they study, and two, a person's success at solving a novel problem is partly a function of whether he or she possesses the necessary subgoals to solve that problem. A subgoal is an unknown entity, numerical or conceptual, that must be found in order to achieve a higher-level goal of a problem. A method is a series of steps for achieving a particular subgoal. No explicit mechanism is proposed here for *how* subgoals are learned. Rather, the emphasis is on identifying subgoals and methods and elucidating the conditions under which they will be learned. This emphasis was chosen since it seems important to gather empirical evidence in a variety of domains concerning the factors that influence whether subgoals and methods are learned and to integrate these findings into the simple scheme proposed here, particularly if this scheme can successfully predict transfer performance and improve the design of examples. Initial experiments have provided some support for the above approach for examining transfer (Catrambone, 1991; Catrambone & Holyoak, 1990). This paper explores one particular way of designing examples to increase the likelihood of conveying subgoals to learners: Briefly labeling or elaborating the subgoals in examples. Work in perceptual category learning has suggested that transfer is improved when subjects are led to focus on useful distinctions (categories) during training (e.g., Medin, Dewey, & Murphy, 1983). Subgoal learning may perhaps be viewed as a form of category learning and thus, may also benefit from manipulations that serve to highlight subgoals.

Credits

This research was supported by Office of Naval Research Grant N00014-91-J-1137 and National Institutes of Health Biomedical Research Support Grant 2-SO7-RR07024-24.

It is argued that successful learners acquire useful subgoals and methods for solving problems in a particular domain. The meaning of "useful" is operationally defined as those subgoals and methods that can be used to solve some corpus of target problems. If a learner has learned a particular subgoal, then he or she is more likely to recognize the need to find that subgoal in a novel problem. Even if the method for achieving that subgoal in the novel problem is new or requires a modification of an already-learned method, the learner who recognizes the need to achieve a particular subgoal will at least have some guidance for recruiting useful background knowledge to achieve the modification compared to a learner who does not possess the subgoal. Thus, subgoals provide a control structure to guide problem solving and transfer.

For example, consider a learner who studies a worked-out physics mechanics problem involving a block resting on an inclined plane. Perhaps the problem asks what the angle of incline has to be in order for the block to start sliding down the plane given that other things such as the block's mass and the coefficient of friction were known. A learner could memorize the steps for solving this problem such as picking a coordinate system with the x-axis parallel to the plane (rather than parallel to the problem solver's perspective) and performing the trigonometry for dealing with the angle of incline. However, the learner might not realize that most of these steps collapse into methods for achieving various subgoals. These subgoals might include high-level ones such as finding the opposing forces and equating them as well as lower-level subgoals such as splitting the forces into those acting along the x- and y-axes. Thus, when given a "simpler" problem such as determining how hard one would have to push a block resting on a flat table in order to just get the block to start moving, the learner could be at a complete loss. Many of the steps from the example, such as those involving the trigonometry, are no longer needed, except in a trivial way. If the learner had learned only a series of steps from the example, he or she would have little guidance in determining which ones should be eliminated or adapted for the new problem. Conversely, if the learner did have his or her knowledge organized hierarchically into a series of subgoals, then the learner would be in a better position to separate the steps into those that were necessary for the current subgoals (such as equating the opposing forces) and those that might *not* be necessary (such as choosing a coordinate system that is not horizontal and vertical).

Prior work has demonstrated that when learners study one or two worked-out example problems, such as algebra problems dealing with two workers completing some task together, they can solve isomorphic problems (Reed, Dempster, & Ettinger, 1985). For instance, subjects might study an example in which the workers' rates are given and the goal is to determine how long one of the workers works given that the other worker's time is known. The difficulty emerges when learners are faced with non-isomorphs, that is, problems requiring a modification of the old procedure. For example, in Reed et al.'s study, a non-isomorphic test problem might express one worker's rate in terms of the other worker's rate. Subjects typically had difficulty solving non-isomorphs. Part of the difficulty appeared to be that subjects formed subgoals such as "find the two numbers associated with worker A and multiply them together" rather than subgoals such as "find each worker's rate" and "find each worker's time." In problems which did not have two numbers associated with a worker, perhaps because the rate was a variable, subjects were often unable to solve the problem.

The poor transfer performance that is typically found in studies may partly be due to a failure to identify the necessary knowledge components that should be conveyed to learners to help them solve novel problems. It is claimed here that these components can be viewed as subgoals. The goal of the current experiment was to manipulate the subgoals conveyed to learners via examples and examine transfer effects. Cognitive architectures such as ACT* (Anderson, 1983), PUPS (Anderson & Thompson, 1986), and SOAR (Laird, Newell, & Rosenbloom, 1987) have mechanisms for predicting transfer. For example, ACT* predicts transfer of procedures by measuring overlap of old and new production rules (see also Singley & Anderson, 1989). However, these models are quite complex and many aspects of the models, such as chunking and proceduralization, are not necessary to derive useful predictions about learning from examples. Research utilizing the approach advocated here should allow the development of a relatively simple model to predict learning from examples as well as provide constraints on more complete models of cognition. This approach, if successful in terms of making useful predictions, could be used by researchers and textbook writers fairly easily as a practical guide for designing better examples (to improve learning) and test items (to improve diagnosticity).

In some sense, the greatest value of the approach described here may be its emphasis on the effort required by the teacher or researcher to identify the useful subgoals and methods in a domain. That is, assuming that examples do convey subgoals and methods, it is necessary to first determine what the useful ones are, and then to create examples to convey them. It is not claimed that the useful subgoals and methods can be identified in any algorithmic way. The approach advocated here for determining subgoals and methods is to identify a target set of problems that one wants students to be able to solve. Then solutions to these problems should be written out. These solutions are then analyzed in terms of the steps for achieving certain unknowns, either numerical or conceptual, in a hierarchical fashion. For example, in the physics problems discussed earlier, a high-level subgoal would be to equate the forces along the x-axis. This high-level subgoal will be achieved by first satisfying a lower-level subgoal such as finding the various forces acting along the x-axis. This subgoal can be satisfied by first satisfying the subgoal of identifying all the forces in the problem and breaking them into their x- and y-components. This subgoal decomposition can continue until one reaches the level of mathematical operations. It is a judgment call as to how much decomposition one feels is necessary to represent subgoals and methods at a useful level. In any case, the subgoals and methods identified by the researcher or teacher are at least open to inspection by others and can be debated.

Experiment

The experiment presented here examined subjects' ability to solve permutation and combination problems in probability. This domain was chosen because prior work indicated that learners were strongly influenced by superficial features of training and test problems. For example, Ross (1989; Experiment 1B) had subjects study example probability problems, such as ones involving permutations, and then solve several test problems. The entities (e.g., computers, cars, scientists) in the examples and test problems were manipulated. For instance, Table 1a presents a permutation problem involving the determination of the probability of scientists choosing particular computers. The equation presented to subjects in Ross (1989) in order to solve this example was $1/[n(n-1)...(n-r+1)]$. Test problems included some subjects finding the probability of students choosing particular cars while other subjects had to find the probability of particular students being assigned to particular cars (i.e., inanimate objects "choosing" humans).

Corresponding roles are held by the humans and inanimate objects in the example in Table 1a and the test problem involving students choosing cars. That is, the mathematical roles of the humans (scientists and students) and the inanimate objects (computers and cars) are identical: the inanimate objects provide the values for n and for r while the humans provide the value only for r. The second type of test problem, however, has reversed object correspondences: the humans provide the values for n and r while the inanimate objects provide the value only for r. Ross found that subjects were more successful solving the first type of test problem than the second, presumably because their problem solving was guided to some degree by an object mapping approach. That is, if humans provide the value of n in the examples, the subjects were likely to assign them this role in test problems even if it was really an inanimate object that should be providing the value of n.

Ross also had subjects study and solve combination problems. An example combination problem might ask for the probability of the seven hooks nearest the classroom door being picked by the seven tallest students in a class (see Table 1b). The equation presented to subjects in order to solve combination problems of this sort was $[h!(j-h)!]/j!$, where h is the number of objects (e.g., students) doing the choosing and j is the number of objects in the pool from which objects are chosen (e.g., hooks). Again, Ross found that if the roles of the objects were switched from the training examples to the test problems, subjects would not solve the test problems correctly (that is, they would switch the values for j and h).

While permutation and combination problems can be solved using quite different equations, a further examination of these problem types indicates that they are similar at a fundamental level. Both types of problems can be analyzed by considering the individual event probabilities that contribute to an overall probability. For example, in the permutation problem involving scientists choosing computers (Table 1a), the overall probability can be calculated by explicitly considering each of the individual probabilities. This approach is demonstrated in the solution provided in Table 2c. The combination problem in Table 1b involving students and coathooks can be analyzed in a similar way:

Probability of one of the seven tallest students getting a hook near door = 7/17
Probability of one of the remaining six tallest students getting a hook near door = 6/16
Probability of one of the remaining five tallest students getting a hook near door = 5/15, etc

So, $\frac{7}{17} * \frac{6}{16} * \ldots * \frac{1}{11} = \frac{7!}{17*16*\ldots*11}$ = overall probability

Table 1
Permutation and Combination Problems

a.) The supply department at IBM has to make sure that scientists get computers. Today, they have 11 IBM computers and 8 IBM scientists requesting computers. The scientists randomly choose their computer, but do so in alphabetical order. What is the probability that the first 3 scientists alphabetically will get the lowest, second lowest, and third lowest serial numbers, respectively, on their computers?

b.) The Happy House Nursery School has had 17 hooks put up in the hall for the coats of their 14 students, with each student using one hook. The students are each randomly assigned a hook as they come in one morning. What is the probability that the 7 hooks closest to the classroom door are assigned to the 7 tallest students?

Table 2
Solution Types Used in Experiment for Problem in Table 1a

a.) Equation Solution

The equation needed for this problem is $\frac{1}{n * (n - 1) * \ldots * (n - r + 1)}$. In this problem n = 11 and r = 3. So, $\frac{1}{11 * 10 * 9} = \frac{1}{990}$ = overall probability

b.) Equation+Elaboration Solution:

The equation needed for this problem is $\frac{1}{n * (n - 1) * \ldots * (n - r + 1)}$. This equation allows one to determine the probability of the above outcome occurring. In this problem n = 11 and r = 3. The 11 represents the number of computers that are available to be chosen while the 3 represents the number of choices that are being focused on in this problem. The equation divides the number of ways the desired outcome could occur by the number of possible outcomes. So, inserting 11 and 3 into the equation, we find that $\frac{1}{11 * 10 * 9} = \frac{1}{990}$ = overall probability

c.) Subgoal Solution:

The equation needed for this problem is $\frac{1}{n * (n - 1) * \ldots * (n - r + 1)}$. In this problem n = 11 and r = 3. However, another way of approaching the problem is to think of it in the following way:

Probability of the first scientist (who comes first alphabetically) getting the computer with the lowest serial number = 1/11.
Probability of second scientist getting second lowest serial number = 1/10.
Probability of third scientist getting third lowest serial number = 1/9.

So, $\frac{1}{11} \cdot \frac{1}{10} \cdot \frac{1}{9} = \frac{1}{990}$ = overall probability

d.) Numerator/Denominator-Subgoal Solution:

The equation needed for this problem is $\frac{1}{n * (n - 1) * \ldots * (n - r + 1)}$. In this problem n = 11 and r = 3. However, another way of approaching the problem is to think of it in the following way:

Of the 11 computers that the first scientist has to choose from, only 1 of them would be acceptable, since only 1 has the lowest serial number. So, the probability of the first scientist getting the computer with the lowest serial number = 1/11.

Of the 10 remaining computers that the second scientist has to choose from, only 1 of them would be acceptable, since only 1 has the lowest remaining serial number. So, the probability of the second scientist getting a computer with the lowest remaining serial number = 1/10.

Of the 9 remaining computers that the third scientist has to choose from, only 1 of them would be acceptable, since only 1 of has the lowest remaining serial number. So, the probability of the third scientist getting the computer with the lowest remaining serial number = 1/9.

So, $\frac{1}{11} \cdot \frac{1}{10} \cdot \frac{1}{9} = \frac{1}{990}$ = overall probability

The permutation problem solution presented in Table 2c might help learners form the subgoal of finding each event probability (e.g., the probability of the first scientist getting the computer with the lowest serial number, the probability of the second scientist getting the computer with the second-lowest serial number, etc). In particular, it should help learners form the subgoal of finding the entity being chosen and using the value associated with that entity as the starting value for the denominator. Thus, subjects who study examples using this solution approach are predicted to be more likely to be able to find event probabilities in cases where the roles of the inanimate objects and humans are different from the examples (see Table 3b) compared to subjects who study examples using only the permutation equation (Table 2a). This latter group presumably learns only to place numbers into an equation. Furthermore, learners who study examples emphasizing each event probability might be more likely to represent the event probabilities correctly in combination problems for which the numerators are no longer simply "1" (such as the combination problem in Table 1b). That is, because these subjects are assumed to be more likely to be able to focus on the individual event probabilities, they may have a better chance of representing them correctly compared to learners who do not have this focus. However, this generalization may be too difficult for learners when solving combination problems. Since the solution presented in Table 2c does not emphasize finding the numerator, but rather just involves placing a "1" in each numerator, subjects who study this solution type might not form the subgoal of finding the numerator from the number of "acceptable" outcomes. Subjects may require example solutions that *explicitly* highlight how the numerator is chosen in the permutation examples in order to form a subgoal for dealing with the numerators. A solution approach that might accomplish this is presented in Table 2d.

Another concern in this study was whether subjects might misinterpret the combination problems as permutation problems. If that were to happen, then subjects would be likely to use "1" in the numerators of the probabilities for all problems regardless of the examples they studied. In order to investigate this issue, some subjects were given combination problems that contained a clarification indicating that order did not matter. Subjects who study examples demonstrating the solution approach in Table 2d might outperform subjects who study examples using the approach in Table 2c on combination problems containing a clarification. This would occur because the latter group will presumably not have the subgoal for finding the number of acceptable outcomes to put in the numerator for each probability and thus, they would be more likely to simply use "1" in the numerator.

Method

Subjects. Subjects were 90 students from introductory psychology classes at the Georgia Institute of Technology.

Materials and Procedure. All subjects studied three isomorphic example problems dealing with permutations that involved humans choosing objects (including the problem in Table 1a). Two factors were manipulated. The first was the type of solution provided to the examples subjects studied. The examples either provided a solution using the permutation equation (the Equation-Only group; see Table 2a for an example), the permutation equation plus an elaboration of what the numbers in the equation represented (the Equation-Elaboration group; Table 2b), a subgoal-oriented solution that emphasized the rationale for the denominator in each probability (the Denominator-Subgoal group; Table 2c), or a subgoal-oriented solution that contained a rationale for the denominator *and* numerator in each probability (the Numerator/Denominator-Subgoal group; Table 2d).

The second factor manipulated was whether the test combination problems contained a clarification about order. The clarification consisted of the following sentences added on to the end of the first and second combination problems (see Table 3c, d), respectively: "It does not matter which of the particular seven hooks closest to the door these students get, just as long as it is any one of the seven closest" and "It does not matter which of the particular six front seats the pitchers get, just as long as it is any one of the six in the front."

After studying the examples subjects solved the four problems in Table 3. Subjects could not refer back to the examples when working on the problems. The first problem was isomorphic to the examples (see Table 3a). The second problem was a permutation problem isomorphic to the examples but had humans playing the roles of n and r and objects playing only the role of r (Table 3b). The third and fourth problems were combination problems. The solutions to these problems involve both a numerator and denominator that start at some value and then are decremented. The first combination problem involves humans picking objects (Table 3c) while the second combination problem involves objects picking humans (Table 3d).

Results and Discussion

Each problem was scored for whether a subject used the correct starting number in the denominator and whether the denominator was decremented appropriately. For instance, the solution to the second permutation problem is 1/11 * 1/10. If subjects wrote 1/14 * 1/13, confusing the roles of the chairs and secretaries, they would be scored as having the incorrect starting number of the denominator but the correct number of decrements. For combination problems, the numerator and denominator were both scored for whether the correct starting number was used and whether they were decremented appropriately. The frequencies for the various categories listed above were analyzed using the likelihood ratio chi-square test (G^2; Bishop, Fienberg, & Holland, 1975).

There was no significant difference in performance among the groups on the first permutation problem that was isomorphic to the training examples and had humans and objects playing the same roles as in the examples (Equation-Only: 73%, Equation-Elaboration: 83%, Denominator-Subgoal: 91%, and Numerator/Denominator-Subgoal: 91%), $G^2(3)=3.77, p=.29$.

The second permutation problem had objects picking humans, a reversal from the training examples (see Table 3b). The major error subjects made was to use the wrong starting value for the denominator. Many subjects used 14 (the number of chairs) as the starting point rather than 11 (the number of secretaries). It was expected that the two subgoal groups would be more likely to use the correct starting point for the denominator than the equation groups. However, while the subgoal groups tended to perform the best, the most striking result was the poor performance of the Equation-Elaboration group (Equation-Only: 32%, Equation-Elaboration: 13%, Denominator-Subgoal: 48%, and Numerator/Denominator-Subgoal: 41%), $G^2(3)=7.63$, $p=.054$. It is not clear why the Equation-Elaboration group did so poorly.

Table 3
Test Problems

a.) As part of a new management policy, the Campbell Company is allowing the 20 company-owned vacation cottages to be used for vacations by their 14 plant managers. If the managers, in order of seniority, randomly choose a cottage from a list, what is the probability that the manager with the most seniority gets the most lavish cottage, and the manager with the second most seniority gets the second most lavish, and the manager with the third most seniority gets the third most lavish, and the manager with the fourth most seniority gets the fourth most lavish cottage?

b.) The secretaries at city hall are supposed to get new chairs this week. Today, city hall received 14 new chairs and there are 11 secretaries requesting them. For inventory purposes, the property manager wants to assign the chairs in the order that they are unpacked. So, starting with the chair that is unpacked first, he randomly chooses a secretary to receive it, and continues until all the secretaries have chairs. What is the probability that the first 2 secretaries alphabetically will get the first and second chairs that are unpacked, respectively?

c.) The Happy House Nursery School has had 17 hooks put up in the hall for the coats of their 14 students, with each student using one hook. The students each choose a hook at random as they come in one morning. What is the probability that the 7 tallest students get the 7 hooks closest to the classroom door?

d.) The Nashville Gnats Baseball team has a bus that has 30 seats. There are 25 players that are going on a road trip to play in a nearby town. To avoid arguments, the manager randomly chooses a player for each seat, starting with the seats in the front. What is the probability that the 6 pitchers get the 6 front seats?

Table 4
Performance (Percent Correct) on Various Aspects of Combination Problems as a Function of Examples Studied and Presence of Clarification

	Group							
	Clarification				No Clarification			
Problem	Eq Only (n=11)	Eq + Elab (n=12)	Denom Subgoal (n=12)	Num/Denom Subgoal (n=11)	Eq Only (n=11)	Eq + Elab (n=11)	Denom Subgoal (n=11)	Num/Denom Subgoal (n=11)
Combination Problem #1 (people choose objects)								
Correct Start for Denominator	82	70	75	80	91	100	91	70
Correct Denominator Decrement	82	60	92	90	82	91	91	80
Correct Start for Numerator	0	0	25	40	9	0	20	10
Correct Numerator Decrement	0	0	33	40	0	0	18	10
Combination Problem #2 (objects choose people)								
Correct Start for Denominator	27	30	33	36	27	36	9	27
Correct Denominator Decrement	82	70	92	64	82	82	82	64
Correct Start for Numerator	0	0	33	46	9	0	9	9
Correct Numerator Decrement	0	0	33	36	9	0	9	9

The third and fourth problems were combination problems (see Tables 3c and d). It was hypothesized that subjects might interpret these combination problems as permutation problems and, as a result, the subgoal manipulation would not have as strong an effect as it could. In order to investigate this possibility, half of the subjects in each instructional group received modified versions of the two combination problems that included a clarification about order. If subjects needed the clarification in order to be more likely to interpret these problems correctly, then the instructional manipulation might only produce an effect for subjects who receive the clarification.

The first combination problem involved humans picking objects. Thus, as in the examples, humans provide the value for the denominator. The presence of a clarification did not seem to differentially influence performance for the various instructional groups on the likelihood of choosing the correct starting value for the denominator (see Table 4; Clarification: $G^2(3)=.49, p=.92$, No-Clarification: $G^2(3)= 5.29, p=.15$). The clarification also did not differentially affect performance on decrementing the denominator (Clarification: $G^2(3)=4.04$, $p=.26$, No-Clarification: $G^2(3)=.91, p=.82$). These results are not surprising since the clarification has little relationship to choosing the denominator and decrementing it appropriately.

Subjects were much less successful choosing the correct numerator on the first combination problem. Unlike the examples, the numerator is not simply "1," rather it is a value starting at 7 and decrementing to 1. Most subjects simply used "1" as the numerator, as in the training examples. For non-clarification subjects, there was no significant difference in the groups' success in choosing the correct starting value for the numerator as a function of the examples they studied, $G^2(3) = 3.21, p = .36$. However, for clarification subjects, the subgoal groups outperformed the equation groups with the Numerator/Denominator group performing the best, $G^2(3) = 11.25, p = .01$. This pattern of results is repeated for decrementing the numerator.

Subjects were less likely to choose the correct starting value for the denominator in the second combination problem, presumably due to the switch in roles of humans and objects; in this problem objects were picking humans. As in the prior problem, though, the presence of a clarification did not seem to differentially influence performance for the various instructional groups on choosing the correct starting point for the denominator (Clarification: $G^2(3) = .24, p = .97$, No-Clarification: $G^2(3) = 2.58, p = .46$). The clarification also did not differentially affect performance on decrementing the denominator (Clarification: $G^2(3) = 3.21, p = .36$, No-Clarification: $G^2(3) = 1.45, p = .69$).

As in the first combination problem, subjects were not very successful choosing the correct numerator. For non-clarification subjects there was no significant difference in the groups' success in choosing the correct starting value for the numerator, $G^2(3) = 1.80, p = .62$. However, for clarification subjects, the subgoal groups outperformed the equation groups with the Numerator/Denominator group performing the best, $G^2(3) = 12.03, p = .007$. This pattern of results is repeated for decrementing the numerator.

Conclusions

The results from this study are important because they suggest that training subgoal recognition facilitates their flexible application. This conclusion is supported by the fact that the Numerator/Denominator-Subgoal group was the most likely to recognize that the numerators of the individual event probabilities in the combination problems were not simply 1 but rather were a value that represented the number of acceptable outcomes. Even though these subjects had studied examples that only used 1 in the numerators, they were sensitized to the role this value played--presumably via the elaboration during training--and were able to generalize when faced with novel problems. This supports the claim that if useful subgoals can be conveyed to learners then they will be able to solve novel problems, that is, problems that involve modified or new methods for achieving those subgoals.

General Discussion

The approach for predicting transfer described in this paper views problem-solving knowledge in terms of subgoals and methods. An important task then is to develop rules for predicting when these subgoals and methods will be learned from examples and to develop a more theoretically-guided motivation for the rules. It has been argued that the relatively simple scheme presented here can account for transfer effects without considering the additional complexities of a full model of cognition. Clearly this simplicity is a virtue only if the approach is reasonably accurate and can be applied fairly easily to a variety of domains by different researchers and teachers. While full models of cognition are extremely valuable, a simpler approach to predicting transfer, if successful, can help improve instruction in a direct way. The quality of examples and the diagnosticity of tests can be improved if instructors and textbook writers become sensitive to the subgoals and methods that students need to learn in a particular domain. The exercise of identifying these components is valuable in its own right since it would make researchers and instructors aware of the building blocks learners need to acquire in order to solve novel problems successfully. Examples can then be devised that convey the subgoals and methods. Although converging evidence is certainly needed, the present results are consistent with predictions that learners who acquire a more hierarchical subgoal structure will be more successful on novel problems than learners who learn a linear series of steps.

References

Anderson, J.R. (1983). *The architecture of cognition*. Cambridge, MA: Harvard University Press.

Anderson, J.R., & Thompson, R. (1986). *Use of analogy in a production system architecture*. Paper presented at the Illinois Workshop on Similarity and Analogy, Champaign-Urbana, IL, June.

Bishop, Y.M.M., Fienberg, S.E., & Holland, P.W. (1975). *Discrete multivariate analysis: Theory and practice*. Cambridge, MA: MIT Press.

Bovair, S., Kieras, D.E., & Polson, P. (1990). The acquisition and performance of text editing skills: A production system analysis. *Human Computer Interaction, 5(1)*, 1-48.

Catrambone, R. (1991). *Using labels to help learners acquire subgoals for solving novel problems*. Paper presented at the Fourth European Conference for Research on Learning and Instruction, Turku, Finland.

Catrambone, R., & Holyoak, K.J. (1990). Learning subgoals and methods for solving probability problems. *Memory & Cognition, 18(6)*, 593-603.

Laird, J.E., Newell, A., & Rosenbloom, P.S. (1987). Soar: An architecture for general intelligence. *Artificial Intelligence, 33(3)*, 1-64,

Mawer, R.F., & Sweller, J. (1982). Effects of subgoal density and location on learning during problem solving. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 8*, 252-259.

Medin, D.L., Dewey, G.I., & Murphy, T.D. (1983). Relationships between item and category learning: Evidence that abstraction is not automatic. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 9(4)*, 607-625.

Newell, A., & Simon, H.A. (1972). *Human problem solving*. Englewood Cliffs, NJ: Prentice-Hall.

Reed, S.K., Dempster, A., & Ettinger, M. (1985). Usefulness of analogous solutions for solving algebra word problems. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 11*, 106-125.

Ross, B. (1989). Distinguishing types of superficial similarities: Different effects on the access and use of earlier problems. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 15*, 456-468.

Singley, M.K., & Anderson, J.R. (1989). *The transfer of cognitive skill*. Cambridge, MA: Harvard University Press.

Acknowledgement

I thank Alana Anoskey for her help in collecting and coding the data and Sharon Carter and two anonymous reviewers for their helpful comments.