

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Effects of Background on Subgoal Learning

Permalink

<https://escholarship.org/uc/item/2482k4d5>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 17(0)

Author

Catrambone, Richard

Publication Date

1995

Peer reviewed

Effects of Background on Subgoal Learning

Richard Catrambone
Georgia Institute of Technology
School of Psychology
Atlanta, GA 30332-0170
USA
rc7@prism.gatech.edu

Abstract

It is hypothesized that when a set of steps in an example solution are labeled, the label can serve as a cue to the learner to group those steps and to attempt to determine their purpose. The resulting subgoal that represents the steps' purpose can aid transfer to novel problems that involve the same subgoal but require new or modified steps to achieve it. The present experiment tested the label-as-grouping-cue hypothesis by examining transfer performance by learners with different math backgrounds who studied examples that used either no labels or labels that varied in meaningfulness. Learners with a stronger math background transferred equally well regardless of the meaningfulness of the label, and better than learners not receiving labels in their examples, while learners with weaker math backgrounds transferred successfully only when they studied examples using meaningful labels. This result is consistent with the claim that the presence of a label, rather than only its semantic content, can be sufficient to induce subgoal learning if the learner has sufficient background knowledge.

Organizing Problem Solving Knowledge by Subgoals

A good deal of research has examined the transfer success people have after studying training materials such as those containing step-by-step instructions (Kieras & Bovair, 1984; Smith & Goodman, 1984), examples (e.g., Ross, 1987, 1989), or both (Fong, Krantz, & Nisbett, 1986). Although there have been some exceptions (e.g., Fong et al., 1986; Zhu & Simon, 1987), the usual finding from such research is that people can carry out new procedures or solve new problems that are quite similar to those on which they were trained, but have difficulty when the novel cases involve more than minor changes from what they had previously studied.

This transfer difficulty seems to stem from a tendency by many learners to form representations of a solution procedure that consist of a linear series of steps rather than a more structured hierarchy. An advantage of a hierarchical organization is that it can provide guidance for adapting the procedure for novel cases. One potentially useful hierarchical organization for a solution procedure would be a set of goals and subgoals with methods for achieving them (e.g., Card, Moran, & Newell, 1983; Catrambone & Holyoak, 1990; Newell & Simon, 1972; Singley & Anderson, 1989). Problems within a domain typically share

the same set of subgoals, although the steps for achieving the subgoals might vary from problem to problem. For instance, physics mechanics problems typically share the subgoals of identifying all "systems" in the problem and identifying all forces acting on the object of interest regardless of whether the problems involve objects on inclined planes or blocks suspended over pulleys (Heller & Reif, 1984).

Consider a student facing a novel problem, that is, one in which the steps are not the same as those seen in a previously-studied example. If the student has memorized only a rote set of steps for the overall solution procedure, he or she will have little guidance as to which steps need to be modified, as well as what new steps might need to be created, in order to solve the problem. Conversely, a student who learned a solution procedure organized by subgoals and methods--a set of steps for achieving a subgoal--could attempt to apply those subgoals to the novel problem. This approach has two advantages. First, the learner would know which steps from the learned procedure are relevant for achieving a particular subgoal. Thus, if those exact steps can not be carried out in the current problem, the learner knows that *those* steps need to be modified. Second, if the learner is attempting to achieve a particular subgoal and realizes that a modification to the old steps will not achieve it, then the subgoal can help constrain the memory search for other relevant information for achieving that subgoal (Anzai & Simon, 1979). Thus, the search space for useful information would be reduced.

Factors Influencing Subgoal Learning

Anzai and Simon (1979) offered an account of subgoal learning in the context of a person learning to solve the Tower of Hanoi problem. They argued that subgoal acquisition is greatly aided when the search space (e.g., possible moves in the Tower of Hanoi problem) is simplified. When the search space is simplified, working memory load is reduced. This aids subgoal formation because a *subgoal* is formed when a learner is working towards a certain goal (perhaps derived from task instructions) and notices that a set of steps places him or her in a situation to be able to carry out additional steps that ultimately achieve the goal. The learner will be better able to notice the result of the first set of steps, and be able to chunk that sequence of steps, if working memory load has been reduced (see also Sweller, 1988).

Labels. A key component in Anzai and Simon's model is the presence of a perceptual system that allows the learner to observe various external features of the problem situation. For example, in the case of the Tower of Hanoi problem, one feature would be a particular disk being located next to a smaller disk. However, in learning tasks that are less obviously perceptually oriented, such as learning to solve word problems in probability, physics, or algebra, simple perceptual features are less likely to play a key role in subgoal formation. Rather, cues in worked examples will play a larger role. These cues may take the form of text and diagrams in the problem that direct the learner to relevant aspects of the problem and relevant prior knowledge (cf. Ward & Sweller, 1990). A label can serve as a cue by leading the learner to group a set of steps in the example solution and thus, to increase his or her chances of recognizing that a particular outcome is the result of the execution of those steps. That is, the recognition of the grouping is hypothesized to lead the learner to try to uncover the purpose of the group of steps. This "purpose" can be conceptualized as a subgoal.

Background Knowledge. Anzai and Simon (1979) suggested that one way a learner can simplify the search space, and thus, to reduce working memory load, is to use prior knowledge of certain facts that can be applied to the domain. In the case of the Tower of Hanoi, such a fact might be that move repetitions are inefficient. In domains such as probability or physics, relevant background knowledge might include the ability to recognize what a set of steps calculate. Ausubel (1968, p 148-149) suggested that the value of "organizers" hinges upon the learner possessing relevant background information so that the pieces of information being organized already have some meaning. For instance, if a student learning mechanics is told that one part of a solution procedure is to determine the components of force along the x and y axes, this organizer for the subsequent steps will be of minimal use if the learner knows little or nothing about coordinate systems or trigonometry.

Testing the Label-as-Grouping-Cue Hypothesis

In the probability examples used in the current study, the ultimate goal of each is to calculate a probability. The solution procedure for achieving this goal involves a number of steps, a subset of which constitutes a sequence of multiplication and addition operations that can be grouped under the subgoal "find the total frequency of the event."

Consider the "No Label" solution to the probability example in Table 1 involving the Poisson distribution.¹ A learner could study this example and memorize the steps for solving a problem that involves the same set of steps, even

¹The Poisson distribution is often used to approximate the binomial for events occurring with small probabilities. The

Poisson equation is $P(X=x) = \frac{[(e^{-\lambda})(\lambda^x)]}{x!}$, where λ is the average (the expected value) of the random variable X.

if the new problem involved farmers and tractors instead of lawyers and briefcases. After studying the No Label solution, the learner's knowledge for the part of the solution procedure that involves finding λ , the average, might be represented as:

Goal: Find λ

- Method:*
1. Multiply each category (e.g., owning exactly zero briefcases, owning exactly one briefcase, etc) by its observed frequency.
 2. Sum the results.
 3. Divide the sum by the total number of lawyers to obtain the average number of briefcases per lawyer.

This representation would serve the learner well for problems that involve calculating λ in the same way as the example. However, this representation fails to capture the fact that the first line of the No Label solution also involves calculating a total frequency. Finding the total frequency is a subgoal that might be achieved in a variety of ways depending on the givens in the problem. A novel problem that requires finding total frequency in a different way than in the example might cause problems for the learner with the above representation. For instance, consider the problem in Table 2b. In this problem the total frequency is calculated by adding a set of simple frequencies. This is a less-complex method than was used in the example, but the learner might not be able to construct it because the subgoal for finding the total frequency, and an instance of a method for achieving it, were never isolated. If the learner had formed the following representation, then his or her chance of solving the problem in Table 2b might be better since this representation identifies the steps involved in finding the total:

Goal: Find λ

- Method:*
1. *Goal: Find total number of briefcases*
Method:
 - a. Multiply each category by its observed frequency.
 - b. Sum the results to obtain the total number of briefcases.
 2. Divide the total number of briefcases by the total number of lawyers to obtain the average number of briefcases per lawyer.

Catrambone (1995) found that learners studying the "Meaningful Label" solution in Table 1 were more likely than those studying the No Label solution to find the total frequency as measured by their success at solving problems such as the ones in Table 2. This was taken as initial evidence that the former group had learned the subgoal to find a total. While the results of Catrambone (1995) were consistent with the claim that a label aids subgoal learning by leading a learner to group a set of steps, they did not constitute a strong test of the account. It is possible that part of the transfer advantage enjoyed by the Meaningful

Label group could have been due to the fact that the label itself provided information beyond serving as a cue to group a set of steps. That is, the label indicated that the total number of briefcases was being found. Thus, instead of the label leading learners to group a set of steps and to form a subgoal to represent the steps' purpose, it may simply have provided them with this fact: finding the total number of things is something that one does when solving Poisson problems.

One way to tease apart these possible explanations is to provide learners with labels, such as the label "Ω" used in the "Less-Meaningful Label" solution in Table 1, that contain no explicit information about the domain and examine whether transfer performance is as good as transfer performance by learners who study examples with more meaningful labels. A second way to get at this issue is to examine the effects of learners' background on subgoal formation.

A learner with a weak math background might look at a series of addition and multiplication steps for the No Label

solution in Table 1 (i.e., "1(180) + 2(17) + 3(13) + 4(9)") and not group them and therefore not notice that they calculate a total. Even a learner with a reasonable math background (but little or no training in probability) might be predicted to be less likely to form the subgoal of finding the total frequency in this situation compared to a learner studying the Meaningful Label solution in Table 1 in which the steps were labeled with "total number of briefcases owned." However, if it is merely the *presence* of a label, rather than its *content*, that is sufficient to lead a learner to group a set of steps and form a subgoal to represent their purpose, then the Less-Meaningful Label solution in Table 1 should also be effective in helping a learner to form the subgoal to find a total, at least if the learner has a reasonably strong math background. Conversely, a learner with a weaker math background might be less likely to be able to determine the purpose of the steps labeled with "Ω" in the Less-Meaningful label solution and thus, be less likely to form the subgoal to find a total. This learner would require a more meaningful label in order to form this subgoal.

Table 1: Example with No Label , Meaningful Label, and Less-Meaningful Label Solutions.

A judge noticed that some of the 219 lawyers at City Hall owned more than one briefcase. She counted the number of briefcases each lawyer owned and found that 180 of the lawyers owned exactly 1 briefcase, 17 owned 2 briefcases, 13 owned 3 briefcases, and 9 owned 4 briefcases. Use the Poisson distribution to determine the probability of a randomly chosen lawyer at City Hall owning exactly two briefcases.

No Label Solution:

$$E(X) = \frac{1(180) + 2(17) + 3(13) + 4(9)}{219} = \frac{289}{219}$$

$$= 1.32 = \lambda = \text{average number of briefcases owned per lawyer}$$

$$P(X=x) = \frac{[(e^{-\lambda})(\lambda^x)]}{x!}$$

$$P(X=2) = \frac{[(2.718^{-1.32})(1.32^2)]}{2!} = \frac{(.27)(1.74)}{2} = .235$$

Meaningful Label Solution:

$$E(X) = \frac{1(180) + 2(17) + 3(13) + 4(9)}{219} = \frac{\text{total number of briefcases owned}}{219} = \frac{289}{219}$$

$$= 1.32 = \lambda = \text{average number of briefcases owned per lawyer}$$

(rest of solution identical to No Label solution)

Less-Meaningful Label Solution:

$$E(X) = \frac{1(180) + 2(17) + 3(13) + 4(9)}{219} = \frac{W}{219} = \frac{289}{219}$$

$$= 1.32 = \lambda = \text{average number of briefcases owned per lawyer}$$

(rest of solution identical to No Label solution)

Table 2: Sample Test Problems.

a.) A number of celebrities were asked how many commercials they made over the last year. The 20 celebrities made a total of 71 commercials. Use the Poisson distribution to determine the probability that a randomly chosen celebrity made exactly 5 commercials.

Solution (not seen by participants):

$$E(X) = \frac{71}{20} = 3.55 = \lambda = \text{avg \# of commercials per celebrity}$$

$$P(X=5) = \frac{[(2.718^{-3.55})(3.55^5)]}{5!} = \frac{(.029)(563.82)}{120} = .136$$

b.) Over the course of the summer, a group of 5 kids used to walk along the beach each day collecting seashells. We know that on Day 1 Joe found 4 shells, on Day 2 Sue found 2 shells, on Day 3 Mary found 5 shells, on Day 4 Roger found 3 shells, and on Day 5 Bill found 6 shells. Use the Poisson distribution to determine the probability of a randomly chosen kid finding 3 shells on a particular day.

Solution (not seen by participants):

$$E(X) = \frac{4 + 2 + 5 + 3 + 6}{5} = \frac{20}{5} = 4.0 = \lambda = \text{average number of shells per kid}$$

$$P(X=3) = \frac{[(2.718^{-4.0})(4.0^3)]}{3!} = \frac{(.018)(64)}{6} = .195$$

Experiment

Method

Participants. Participants were 150 students recruited from several Atlanta-area colleges who received course credit or payment for their participation. In order to participate in the experiment, a student either had to have taken no college level calculus courses or to have had between two and four college-level mathematics courses beyond introductory calculus. However, no student could have taken a probability course prior to participating in the experiment.

Materials and Procedure. All participants initially studied a cover sheet that briefly described the Poisson distribution along with a simplified notion of a random variable.

Participants were randomly assigned to one of three groups. Of the 50 participants per group, 25 had a stronger (calculus) math background and 25 had a weaker (no calculus) math background. The Meaningful Label group studied three examples demonstrating the weighted average method for finding λ in which the steps for finding the total frequency were given a label that was assumed to have meaning to the participants and made mathematical sense given the steps that preceded it (see the Meaningful Label Solution in Table 1 for an example). The Less-Meaningful Label group studied examples in which the steps for finding the total frequency were labeled with Ω which was assumed to have little meaning for the participants in the context of the examples (see the Less-Meaningful Label Solution in Table 1). The No Label group studied examples in which

the steps for finding the total frequency were not labeled (see the No Label Solution in Table 1).

After studying the examples, participants were asked to describe how to solve problems in the domain. After writing their descriptions, participants solved six problems. The first two required the use of the weighted average method for finding λ (isomorphic to the example in Table 1). The next four problems required new ways of finding the total frequency: either by recognizing that the value was given directly in the problem (see Table 2a for an example) or by adding simple frequencies (see Table 2b). Participants were told not to look back at the examples when writing their descriptions or solving the test problems. Only the results from the transfer task have been analyzed so far.

Predictions

Most participants were predicted to find λ correctly in the isomorphic test problems since the same sets of steps used in the examples could be applied to those problems.

If the presence of a label, rather than its semantic content, is sufficient to promote grouping, then learners with adequate background knowledge should be able to determine the purpose of the grouped steps and thus, form a subgoal to represent that purpose regardless of the meaningfulness of the label. Therefore, participants with a calculus background should be equally likely to learn the subgoal to find a total in the Meaningful and Less-Meaningful Label conditions. Thus, these two groups should find λ correctly on the novel test problems about equally often and should outperform the No Label group since these problems involved new ways of finding the total frequency. However, for participants with

the weaker math background, the semantic content of the label is predicted to play a larger role in helping them to understand the purpose of the labeled steps. Thus, for these learners, those receiving the more meaningful label should be more likely to form the subgoal to find a total compared to those receiving the less-meaningful label. Therefore, the Meaningful label condition should produce more success than the Less-Meaningful Label condition at finding λ on the novel problems.

Results

Participants were given a score of 1 for a given problem if they found λ correctly and a score of 0 otherwise. The scores for the two problems that were isomorphic to the training examples, Problems 1-2, were summed, thus creating a score from 0-2 for performance on those problems. Similarly, the scores for the four novel problems, Problems 3-6, were summed, thus creating a score from 0-4 for performance on those problems.

A two-way analysis of variance was carried out with condition and math background as the factors. Table 3 presents the average scores on the test problems as a function of group. There was a significant effect of condition, $F(2, 144) = 9.49, p = .0001, MS_e = 3.18$, and math background $F(1, 144) = 19.31, p < .0001$. There was also a significant interaction between these factors, $F(2, 144) = 3.12, p = .047$. The most typical mistake that students made on these problems was to write in the solution area that not enough information was given to solve the problem. If only participants with a calculus background were considered, pairwise comparisons indicated that the two label groups did not perform differently ($p > .7$) while both showed a tendency to outperform the No Label group ($p = .04$ vs Meaningful Label; $p < .08$ vs Less-Meaningful Label; one-tailed). If only participants without a calculus background were considered, the Meaningful Label group outperformed the other two groups (both p 's $< .0005$) while there was no significant difference between the Less-Meaningful and No Label groups ($p > .7$).

Discussion

Students frequently learn a solution procedure as a series of steps with little or no higher-level organization (Reed, Dempster, & Ettinger, 1985). As a result, while they can solve new problems that involve the same steps as a previously-studied example, they have difficulty with

problems that require a change in the steps, even though the conceptual structure from the example to the problem is preserved.

A guiding assumption of the present research is that transfer performance will be enhanced if a solution procedure is structured by subgoals and a method for achieving each one rather than just a single linear set of steps for the entire procedure. Presumably there is a continuum of structuredness depending on the number of subgoals into which a procedure is broken.

The results from the present experiment are consistent with the hypothesis that the presence of a label, rather than its semantic content, can be sufficient to induce a learner to form a subgoal, at least when the learner has adequate background to take advantage of the label. Presumably, a label serves as a cue to learners to group a set of steps and to retrieve information from long-term memory in order to explain why those steps belong together. The subgoal that is hypothesized to be formed as a result of this process can aid transfer to novel problems. While the present results do not include converging evidence of subgoal learning, such as evidence from learners' descriptions of how to solve the problems, prior studies have found such converging evidence (Catrambone, 1994; Catrambone, 1995).

It is worth noting that while the present results provide support for the theoretical claim that subgoal learning can be aided by the presence of a label, regardless of its meaningfulness, there is also a potential educational implication. A subgoal that is formed in response to a label that makes mention of superficial features from the example might become tied to those features. For instance, the subgoal formed by Meaningful Label participants in the experiment might have been "find the total number of objects." Conversely, a subgoal formed in response to a more abstract label might be less likely to be tied to superficial features. For instance, the subgoal formed by Less-Meaningful Label participants with a stronger math background might have been "find the total." This latter subgoal is more general and closer to being formally correct. One implication of forming a subgoal that is tied to superficial features is that the learner is confusing superficial and structural features of the domain. Future work could test this possibility by constructing test problems that systematically manipulate the relationship between superficial and structural features and observing the degree to which the features guide learners' performance (cf. Ross, 1987, 1989).

Table 3: Scores on Novel Test Problems as a Function of Condition and Math Background.

	Meaningful Label	Less-Meaningful Label	No Label
Calculus	3.04	2.88	2.08
No Calculus	2.72	0.80	0.64

(Maximum possible score = 4)

Acknowledgements

This research was supported by Office of Naval Research Grant N00014-91-J-1137.

References

- Anzai, Y., & Simon, H.A. (1979). The theory of learning by doing. *Psychological Review*, 86 (2), 124-140.
- Ausubel, D.P. (1968). *Educational psychology: A cognitive view*. New York: Holt, Rinehart and Winston, Inc.
- Card, S.K., Moran, T.P., & Newell, A. (1983). *The psychology of human-computer interaction*. Hillsdale, NJ: Erlbaum.
- Catrambone, R. (1995). Aiding subgoal learning: Effects on transfer. *Journal of Educational Psychology*, 87 (1), 5-17.
- Catrambone, R. (1994). Improving examples to improve transfer to novel problems. *Memory & Cognition*, 22 (5), 606-615.
- Catrambone, R., & Holyoak, K.J. (1990). Learning subgoals and methods for solving probability problems. *Memory & Cognition*, 18 (6), 593-603.
- Fong, G.T., Krantz, D.H., & Nisbett, R.E. (1986). The effects of statistical training on thinking about everyday problems. *Cognitive Psychology*, 18, 253-292.
- Heller, J.I., & Reif, F. (1984). Prescribing effective human problem-solving processes: Problem description in physics. *Cognition and Instruction*, 1 (2), 177-216.
- Kieras, D.E., & Bovair, S. (1984). The role of a mental model in learning to operate a device. *Cognitive Science*, 8, 255-273.
- 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 (1), 106-125.
- Ross, B. (1987). This is like that: The use of earlier problems and the separation of similarity effects. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 13 (4), 629-639.
- 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 (3), 456-468.
- Singley, M.K., & Anderson, J.R. (1989). *The transfer of cognitive skill*. Cambridge, MA: Harvard University Press.
- Smith, E.E., & Goodman, L. (1984). Understanding written instructions: The role of an explanatory schema. *Cognition and Instruction*, 1 (4), 359-396.
- Sweller, J. (1988). Cognitive load during problem solving: Effects on learning. *Cognitive Science*, 12, 257-285.
- Ward, M., & Sweller, J. (1990). Structuring effective worked examples. *Cognition and Instruction*, 7 (1), 1-39.
- Zhu, X., & Simon, H.A. (1987). Learning mathematics from examples and by doing. *Cognition and Instruction*, 4 (3), 137-166.