Generalization by Studying Examples Versus Generalization by Applying Examples to Problems

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

Abstract

Two views of problem solving procedure generalization are compared in an experiment: the Generalization by Applying Examples (GenApp) and Generalization by Studying Examples (GenStudy) views. The results suggest that learners can acquire a sufficiently general approach for solving novel problems by studying appropriately-designed examples that encourage one to form subgoals to represent a solution procedure. Learners who are led to form a more rote procedure show much less transfer. No evidence was found for generalization through application.

Introduction

Learners have difficulty solving problems that involve more than minor changes to the procedure demonstrated by training problems or examples (e.g., Bassok, Wu, & Olseth, 1995; Catrambone, 1995, 1996, in press; Novick & Holyoak, 1991; Reed, Dempster, & Ettinger, 1985; Ross, 1987, 1989). People tend to form solution procedures that consist of a long series of steps rather than more meaningful representations that would enable them to successfully tackle new problems (Singley & Anderson, 1989).

Such findings are a cause for concern. Presumably one of the jobs of education is to equip people to deal with novel problems and situations, not just a small recognizable set. Yet it appears that this job does not get done. Learners seem to be predisposed, or the environment shapes them to develop the disposition, to have their problem solving guided by sets of memorized steps and by surface features of problems (Chi, Feltovich, & Glaser, 1981; Larkin, McDermott, Simon, & Simon, 1980; Ross, 1987, 1989). Surface or superficial features are those aspects of problems that, when changed, do not affect the solution procedure; that is, they have no necessary relevance to the solution to the problem. Learners often do not realize that seemingly different sets of steps across problems might be calculating the same thing such as the force acting on a particular object.

Students tend to memorize the details of how equations are filled out rather than learning the deeper, conceptual knowledge that is implicit in the details. Thus, if they are given a new problem that seems similar to an old one--at a surface level--they will try to apply a set of steps from the old problem. These steps are invoked when the learner recognizes certain features to be present in a problem. If the steps can not be used, the learner will frequently not know what to do or will carry out an inappropriate procedure.

A more fruitful approach to problem solving would be to organize one's problem solving knowledge in some way that generalizes across problems in a domain. One type of knowledge structure that appears to aid procedural generalization is one organized by subgoals.

As used in the present paper, a subgoal represents a meaningful conceptual piece of an overall solution procedure. Subgoals can be used by a learner to help him or her solve novel problems since problems within a domain typically share the same set of subgoals, although the steps for achieving the subgoals might vary from problem to problem. A subgoal can serve as a guide to which part of a previously-learned solution procedure needs to be modified for a novel problem (Catrambone, 1996, in press).

Designing Better Examples

Earlier studies have demonstrated that if examples are designed in such a way as to encourage subgoal learning, then learners are more likely to correctly solve new problems that involve the same subgoals but require new steps for achieving them (Catrambone, 1995, 1996). This view might be called the "generalization by studying examples" (or GenStudy) view.

For instance, consider the permutation example and "equation-oriented" solution in Figure 1a. After studying this example about computers, a learner might think that the way to solve such "choice" problems is to find the number of things being picked, decrement that number by the number of times things are being picked, multiply those numbers together, and then divide 1 by that result. However, such an approach would be wrong for problems in which the roles of humans and objects were reversed. Consider the problem in Figure 1b about chairs. The correct answer to this problem is $\frac{1}{11*10}$. One must consider *what* is being picked and not just assume that the number of objects forms the basis for the denominator. In the computer problem it is computers that are being picked and thus, the number of computers (or things) supplies the starting value for the denominator. Such an approach to the chair problem would lead to an incorrect answer of $\frac{1}{14*13}$. In the chair problem it is the number of secretaries that supplies the starting value for the denominator since secretaries are being assigned to chairs (or, chairs are "picking" secretaries).

Suppose though that the solution studied to the computer problem was presented in the following "subgoal-oriented" 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, $\frac{1}{11} * \frac{1}{10} * \frac{1}{9} = \frac{1}{990}$ = overall probability

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?

Equation-Oriented Solution:

The equation needed for this problem is $\frac{1}{n * (n - 1) * ... * (n - r + 1)}$. 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

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, she 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 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? (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.)

d.) 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 four managers with the most seniority get the most lavish, second most lavish, third most lavish, and fourth most lavish cottages, respectively?

Figure 1: Example and Problems.

This subgoal-oriented solution is assumed to help learners form two goals. The first is the goal to find the overall probability. This goal is assumed to be formed because that goal is explicitly stated in the example. The second is the subgoal to find each event probability, for example, 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. This subgoal is assumed to be formed because each individual event probability is explicitly labeled and spatially separate in the subgoal-oriented solution. A learner who studied such an approach might, when faced with the chair problem, be more likely to notice that chairs are picking secretaries rather than the other way around. Thus, this learner might be more likely to provide the correct answer. In addition, such a learner might also have a better chance at figuring out that the numerators for the individual event probabilities in the problem in Figure 1c are not simply "1" but rather are numbers that indicate the number of "acceptable choices."

GenApp Vs. GenStudy

One problem with prior studies supporting the GenStudy view (e.g., Catrambone, 1995, 1996) is that participants were given, after studying examples, one or two isomorphic problems to solve before being given novel problems to solve (where "novel" means that the problem had a change in roles compared to the examples and may have involved a change in the steps needed to achieve the subgoals). It is possible that the attempt to apply the examples to the isomorphic problems led participants to form generalizations of the solution procedure which then helped them solve the novel problems. This alternative might be called the "generalization by applying examples" (or GenApp) view.

The GenApp view is supported by the findings of Ross and Kennedy (1990). In a typical experiment they had learners study four probability principles (e.g., permutations, combinations) that were each illustrated through a worked example. After studying the principles and examples learners attempted to solve two problems for each principle. The first test problem for each principle either did or did not contain a cue indicating which prior training example was relevant for solving the problem. The second test problem for each principle did not contain a cue.

Ross and Kennedy (1990) found that when learners received a cue on the first test problem for a particular principle, they were more likely to correctly solve the *second* test problem for that principle compared to cases in which the first test problem was uncued. More specifically, this benefit manifested itself in terms of an increased likelihood in using the correct principle for the second test problem as well as instantiating the variables correctly. For instance, if the example involved humans picking objects (e.g., scientists choosing computers), the problems would involve objects "picking" humans (e.g., as a particular computer is unpacked, a randomly chosen scientist is assigned to use it). Learners who were cued to the relevant example when working on the first test problem for a particular principle were more likely to get the roles for humans and objects correct when working on the second test problem for that principle (there was no difference between cued and uncued performance on the *first* test problem with respect to getting the roles correct). Ross and Kennedy argued that differences between the problem and the cued example led learners to form a generalization as they attempted to apply the example to the problem. This generalization affected performance on the second test problem.

One difficulty with the Ross and Kennedy (1990) study is that the examples were not designed to help learners form a generalization of the solution procedure. Rather, training consisted of a statement of the probability principle, a study example, and one means of working out the example (similar to the equation-oriented solution to the computer example in Figure 1a).

The aim of the present experiment was to pit the GenStudy and GenApp views against each other. Participants studied a single permutation example (Figure 1a) with either the equation-oriented solution or the subgoal-oriented solution. The first test problem was a permutation problem that had humans and objects playing either the same roles as in the example or playing reversed roles. So, for half of the participants, the first test problem they received was the chair problem which had reversed roles relative to the example (see Figure 1b). For the other half of the participants, the first test problem was just like the training example and is shown in Figure 1d. The second and third test problems were the same for all participants. The second problem was a permutation problem with reversed roles, that is, objects choosing humans (similar to the problem in Figure 1b). The third problem was a combination problem also with reversed roles (see Figure 1c).

Consider the fate of two hypothetical learners--one who studied the equation-oriented solution to the computer example and one who studied the subgoal-oriented solution--when faced with the novel problem in Figure 1c in which the make-up of the individual events, as well as the roles of humans and objects, are different from the example. The answer to this problem is

$$\frac{6}{25} * \frac{5}{24} * \frac{4}{23} * \frac{3}{22} * \frac{2}{21} * \frac{1}{20} = \frac{6*5*4*3*2*1}{25*24*23*22*21*20}$$

The successful learner must be sensitive to the fact that the numerator for each individual event probability is not simply "1" and that the denominator is not automatically based on the number of objects. A learner with an equation-oriented approach has little guidance for making such observations; however, a learner with the subgoal-oriented approach might be able to figure all this out since he or she is more likely to focus on the individual events.

Predictions

According to the GenApp view, the following predictions should be made:

1) Participants whose first test problem was the reversed roles permutation test problem should do better on the *second* reversed roles permutation problem compared to participants whose first test problem was the same-roles permutation problem. This is because the first group would be led, in the process of applying the example, to form a better generalization than the second group and thus could use this superior generalization to deal with the reversed roles in the second permutation problem.

2) Participants whose first test problem was the reversed roles permutation test problem should do better on the reversed roles combination problem--at least with respect to role assignment, that is, putting the correct values in the denominator--compared to participants whose first test problem was the same-roles permutation problem. Once again, the former group would have a better generalization to use when solving the combination problem.

3) No particular prediction would be made about differential performance between the equation-oriented and subgoal-oriented groups.

According to the GenStudy view, the following predictions should be made:

1) For participants who receive the reversed-roles problem as their first test problem, those receiving the subgoal-oriented solution in the example will do better on that problem compared to participants who studied the equation-oriented solution.

2) Subgoal-oriented participants will solve the second test problem--the reversed-roles permutation problem--with more success than equation-oriented participants.

3) Subgoal-oriented participants will perform better on the reversed-roles combination problem compared to equation-oriented participants.

Experiment

Method

Participants. Participants were 120 students from introductory psychology classes at the Georgia Institute of Technology who participated in the experiment for course credit. None of them had taken a probability course prior to participating in the experiment.

Materials and Procedure. Participants received a booklet containing one training example (a worked-out permutation example involving humans picking objects; see Figure 1a) and three test problems.

Two factors were manipulated: 1) subgoal orientation of the example, and 2) order of test problems. With respect to the subgoal orientation, half of the participants were in the subgoal-oriented condition which meant that that the solution to the studied example was designed to help participants to form the needed subgoals for solving both permutation and combination problems. The other half of the participants studied the equation-oriented solution that encouraged a more rote approach. With respect to test problem order, the first test problem was either a same-roles permutation problem in which humans pick objects (as in the example) or a reversed-roles permutation problem in which objects pick humans. The second and third test problems were the same for all participants: the first was a reversed-roles permutation problem and the second was a reversed-roles combination problem. Thus, there were four groups with 30 participants per group.

Participants were asked to study the example carefully since after studying it they would be asked to solve some problems. They were told they could not look at the example when working on the problems. This restriction was intended to increase the likelihood that participants would pay attention to the example and how it was solved.

Participants worked at their own pace and were asked to show all their work. In general, participants took about 30 minutes to complete the experiment.

Each permutation problem was scored for whether a participant used the correct denominator. For instance, the solution to the chair problem is $\frac{1}{11*10}$. If a participant wrote $\frac{1}{14*13}$, confusing the roles of the chairs and secretaries, the denominator would be scored as incorrect. For the combination problem, the numerator and denominator were both scored as correct or incorrect. Two raters independently scored the problems and agreed on scoring 96% of the time. Any disagreements were resolved by discussion.[1]

Results

Table 1 presents the percentage of subjects in each condition who found the denominator correctly in each problem as well as the percentage who found the numerator correctly in the combination problem. These percentages are compared

---

[1] Logically one could write the correct starting value for the denominator (for permutation and combination problems) or the correct starting value for the numerator (for combination problems) but fail to decrement the value appropriately. In practice though, if a participant found the correct initial value (e.g., 14 for the denominator in the chair problem or 6 for the numerator in the baseball problem), he or she almost invariably did the decrementing appropriately. Thus, for each problem there is a single score for the denominator and, for the combination problem, a single score for the numerator.

in various ways below in order to test the predictions of the GenApp and GenStudy views.

Performance on First Permutation Problem with Same Roles as Example: As the first row of results in Table 1 indicates, participants in both the equation-oriented and subgoal-oriented conditions found the denominator in this problem with little difficulty. This simply demonstrates that participants could mimic the steps shown in the example.

Performance on First Permutation Problem with Reversed Roles from Example: The second row of results in Table 1 shows that, as predicted by the GenStudy view, subgoal-oriented participants outperformed equation-oriented participants. $\chi 2(1,\ N = 60) = 4.44, p = .035$.

Performance on Second Permutation Problem: As predicted by the GenStudy view, subgoal-oriented participants outperformed equation-oriented participants, $\chi 2(1,\ N = 120) = 9.85,\ p = .0017$. Furthermore, the two equation-orientation groups did not differ from each other. The GenApp view predicted a difference between these two groups since the 1st-problem-has-reversed-roles-from-example group should have been led to a generalization that should have helped their performance on the 2nd permutation problem relative to the other equation-oriented group. This did not occur. Such a difference also failed to appear between the two subgoal-oriented groups.

Performance on Combination Problem Denominator: As predicted by the GenStudy view, subgoal-oriented participants outperformed equation-oriented participants, $\chi 2(1, N = 120) = 10.11, p = .0015$.

Performance on Combination Problem Numerator: As predicted by the GenStudy view, subgoal-oriented participants outperformed equation-oriented participants, $\chi 2(1, N = 120) = 9.02, p = .0027$.

Discussion

The overall performance differences among the groups can be summarized as follows: the subgoal-oriented groups outperformed the equation-oriented groups on all aspects of the novel problems (role reversals and using a non-"1" numerator for the combination problem). There was no evidence of improved generalization by any group as a function of having attempted to solve a reversed-roles problem first.

The results suggest that generalization can occur from properly designed examples and that a learner does not necessarily have to apply an example to a problem in order to form useful generalizations. While getting learners to form useful generalizations is an important pedagogical goal, it apparently can be achieved in more than one way. Carefully-designed examples seem to be one effective way to make this happen.

Table 1: Percentage of Participants Correctly Finding Denominators for Each Test Permutation and Combination Problem and Percentage Correctly Finding Numerator for Combination Problem

	Condition			
	Equation-Oriented		Subgoal-Oriented	
Problem Feature	1st Problem has Same-Roles as Example	1st Problem has Reversed-Roles from Example	1st Problem has Same-Roles as Example	1st Problem has Reversed-Roles from Example
Denominator for 1st Permutation Problem (Same Roles)	93.3	n/a	90.0	n/a
Denominator for 1st Permutation Problem (Reversed Roles)	n/a	46.7	n/a	73.3
Denominator for 2nd Permutation Problem (Reversed Roles)	40.0	46.7	70.0	73.3
Denominator for Combination Problem (Reversed Roles)	43.3	50.0	73.3	76.7
Numerator for Combination Problem	26.7	23.3	53.3	50.0

There was no evidence of a generalization being formed due to applying an example to an initial test problem for participants in the equation-oriented condition. This is a bit surprising since this condition was meant to be similar to Ross and Kennedy's (1990). However, there are methodological differences between the present experiment and those in Ross and Kennedy that may account for the lack of an effect of application. For instance, Ross and Kennedy used examples to illustrate four probability principles during training and were explicitly cueing (or not cueing) a relevant example for the first test problem for each principle. Perhaps the potential confusion about which principle is relevant for a test problem played a role in the generalization process. Such confusion was not an issue in the present experiment since only one principle was illustrated during training. A second difference was that no explicit cue to a relevant example was used in the present study. However, there was certainly an implicit cue since participants were told that the example would help them solve the test problems.

Such methodological differences may need to be explored systematically in order to determine if there are situations in which the application of an example aids generalization as much as studying examples *designed* to encourage generalization.

Conclusions

While the results support a subgoal-oriented approach to designing example solutions, they do not provide guidance as to the specific subgoals that should be taught. That is, another researcher or teacher working with permutation and combination problems might determine that a different set of subgoals than those used here are better for students to learn. The aim of this study was to show that a particular set of good subgoals--as determined by a task-analysis and by the researcher's intuition--can be conveyed to learners through examples. It would be useful though to develop constraints on how one determines what are "good" subgoals for problems in a particular domain. Cognitive modeling tools such as ACT-R (Anderson, 1993) may provide constraints within a unified theory that can help one determine the subgoals that should be taught to learners. On the other hand, from a pragmatic point of view, it has been this researcher's observation that forcing oneself to solve a reasonably large number of problems within a domain, and taking careful notes on how one went about solving the problems, can produce a useful list of subgoals, and other types of information, that can then be taught to learners through paper-and-pencil examples, animations, and other types of teaching materials (Catrambone, Stasko, & Byrne, 1996).

References

Anderson, J.R. (1993). *Rules of the mind.* Hillsdale, NJ: Erlbaum.

Bassok, M., Wu, L.L., & Olseth, K.L. (1995). Judging a book by its cover: Interpretive effects of content on problem-solving transfer. *Memory and Cognition, 23(3)*, 354-367.

Catrambone, R. (1995). Aiding subgoal learning: Effects on transfer. *Journal of Educational Psychology, 87* (1), 5-17.

Catrambone, R. (1996). Generalizing solution procedures learned from examples. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 22*(4), 1020-1031.

Catrambone, R. (in press). The subgoal learning model: Creating better examples so that students can solve novel problems. *Journal of Experimental Psychology: General.*

Catrambone, R., Stasko, J., & Byrne, M. (1996). *Do algorithm animations help students learn?* Paper presented at the Basic Research Symposium of the CHI Human Factors in Computing Systems Conference (Vancouver, April).

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

Larkin, J., McDermott, J., Simon, D.P., & Simon, H.A. (1980). Expert and novice performance in solving physics problems. *Science, 208*, 1335-1342.

Novick, L.R., & Holyoak, K.J. (1991). Mathematical problem solving by analogy. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 17* (3), 398-415.

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.

Ross, B.H., & Kennedy, P.T. (1990). Generalizing from the use of earlier examples in problem solving. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 16* (1), 42-55.