

Learning Subgoals and Methods for Solving Problems

Richard Catrambone

Keith J. Holyoak

The University of Michigan
Department of Psychology

The University of California at Los Angeles
Department of Psychology

In prior studies (Catrambone & Holyoak, 1987) we demonstrated that subjects studying examples dealing with the Poisson distribution tended to learn a series of steps for solving problems rather than learning meaningful chunks. We argued that these chunks should consist of relevant subgoals and methods for achieving those subgoals. If a person learns subgoals and methods for a particular problem domain, then they should be better able to solve novel problems in that domain. Novel problems are defined as 1) having a subgoal order that is different from previously studied problems and/or 2) requiring modifications of old methods for achieving one or more subgoals.

One encouraging finding from our earlier experiments, from a pedagogic point of view, was that if subjects studied examples that demonstrated two different solution procedures, then they were more likely to learn some of the subgoals in the problem space. However, subjects still had difficulty modifying old methods. Two limitations with the earlier experiments were that subjects only studied four examples and the procedures differed from each other in several ways. The first limitation made transfer more difficult to achieve and the second limitation made comparisons among groups less straightforward than we would have liked. The purpose of the present paper is to describe an experiment that gave subjects more extensive training and carefully manipulated just the methods subjects learned for finding subgoals. We wished to examine whether more extensive training with two or more methods for finding a subgoal would help learners to figure out how to modify those methods in new problems.

Our work has suggested that people's procedural knowledge for a domain such as solving probability problems can be represented in terms of subgoals and methods. These subgoals and methods may be based on the "true" subgoals and methods in the domain, or they may correspond to superficial features of example problems. There may be no agreement over what are the "true" structural features for a problem domain, but there is certainly more agreement among experts than novices in a domain about the domain's relevant features (e.g., Adelson, 1981; McKeithen & Reitman, 1981).

Related Work

Various researchers have shown that learners focus on superficial features of examples and need to be guided to focus on the deeper aspects. Anderson and his colleagues (Anderson, Farrell, & Sauer, 1984; Pirolli & Anderson, 1985) demonstrated that students learning how to write recursive functions in the LISP programming language relied heavily on the syntactic elements of examples in order to create their first few functions. These students essentially mapped each part of the current problem onto an example. This approach worked because the early target problems were designed to be isomorphic with the examples. However, students must eventually be weaned from heavy reliance on pure example mapping and induced to focus on subgoals and methods in order to be able to solve more difficult or novel problems in a domain.

McKendree (1986) found that geometry students performed best on new problems when they received error feedback on examples that focused on subgoals rather than feedback that simply told the student that he or she had made an error. In addition, an earlier protocol study of tutors

CATRAMBONE, HOLYOAK

indicated that a very important role of feedback was to make the novice's goals explicit and to help keep track of these goals during the task (McKendree, Reiser, & Anderson, 1984).

Charney and Reder (1986) showed that diverse examples are important in helping people learn procedures and when to apply them (i.e., to learn selection rules). Their subjects studied examples in order to learn various commands on a personal computer. Charney and Reder found that varied examples, which demonstrated a range of situations in which a command could be used, helped new users grasp the utility and applicability of the command.

The studies cited above suggest that people learning information in a new domain need to be directed to focus on the relevant aspects of training materials. Examples must be carefully chosen to help the user gain a functional understanding of the subgoals, methods, and when to apply the methods (Charney & Reder, 1986; VanLehn, 1982, 1985). In addition, learners must be able to modify methods in new situations. This last issue was an important focus in the current experiment which dealt with subjects learning how to solve problems dealing with the Poisson distribution.

The Poisson Distribution and Some Examples

The Poisson distribution is often used to approximate binomial probabilities for events that occur with some small probability p . The Poisson equation is: $P(X=x)=[e^{-\lambda}\lambda^x]/x!$, where λ is the expected value--the average--of the random variable X . One use of the Poisson distribution is illustrated in Table 1.

The subgoals and methods (in parentheses) for this "quarry" problem could be listed as follows:

- 1) find λ (find total frequency of event and divide by total number of trials)
- 2) find expected probability for each X (plug each X into the Poisson equation)
- 3) find expected frequency for each X (multiply each $P(X)$ by the total number of trials).

In prior experiments (Catrambone & Holyoak, 1987) we found that subjects could study a problem such as the one in Table 1 and then solve an isomorphic problem such as one dealing with the number of errors made per game by a baseball team's infielders. However, these subjects had great difficulty solving the "birthday" problem in Table 2 that had a different subgoal order and required a modified method to find λ .

In the birthday problem the solver must realize that λ is simply the frequency of the event (number of people in the room) divided by the number of trials (days of the year). Once λ is found, then it can be put into the Poisson equation along with the value of X that is desired in order to find the expected probability for that value of X . Most subjects did not realize that λ was a subgoal that could be found using a modification of the method used in the "quarry" problem (where the event frequency was not given directly). In addition, few subjects even realized they could find the expected probability for a particular value of X . The protocols of most of the subjects indicated that all they really knew how to do was apply a series of steps they had learned from the "quarry" problem. They typically did not differentiate the steps into methods for reaching particular subgoals.

CATRAMBONE, HOLYOAK

Table 1
"Quarry" Problem Using the Poisson Distribution

A horizontal quarry surface was divided into 30 squares about 1 meter on a side. In each square the number of specimens of the extinct mammal *Ditolestes motissimus* was counted. The results are given in the table below. Fit a Poisson distribution to \underline{x} , that is, give the expected frequencies for the different values of \underline{x} based on the Poisson model.

Number of Specimens per Square	Observed Frequency
0	16
1	9
2	3
3	1
4 or more	1
Total	30

SOLUTION:

$$E(X) = [0(16)+1(9)+2(3)+3(1)+4(1)]/30 = 22/30 = .733 = \lambda$$

=average number of specimens found in each square

$$P(X=x) = [(e^{-.733})(.733^x)]/x! = [(.48)(.733)^x]/x!$$

Fitted Poisson Distribution:

x	Expected Frequency
0	.48 * 30 = 14
1	.352 * 30 = 11
2	.1289 * 30 = 4
3	.0315 * 30 = 1
4 or more	.0058 * 30 = 0

Table 2
"Birthday" Problem Using the Poisson Distribution

Suppose you took a random sample of 500 people and found out their birthdays. A "success" is recorded each time a person's birthday turns out to be January 1st. Assume there are 365 days in a year, each equally likely to be a randomly chosen person's birthday. Fit a Poisson distribution to \underline{x} (the number of people born on January 1st) and find the predicted likelihood that exactly 3 people from the sample are born on January 1st.

SOLUTION:

$$\lambda = 500/365 = 1.37 = \text{average number of people born on any given day}$$

$$P(X=3) = [(e^{-1.37})(1.37^3)]/3!$$

$$= [(.254)(2.57)]/6$$

$$=.109$$

=likelihood of exactly three people being born on January 1st

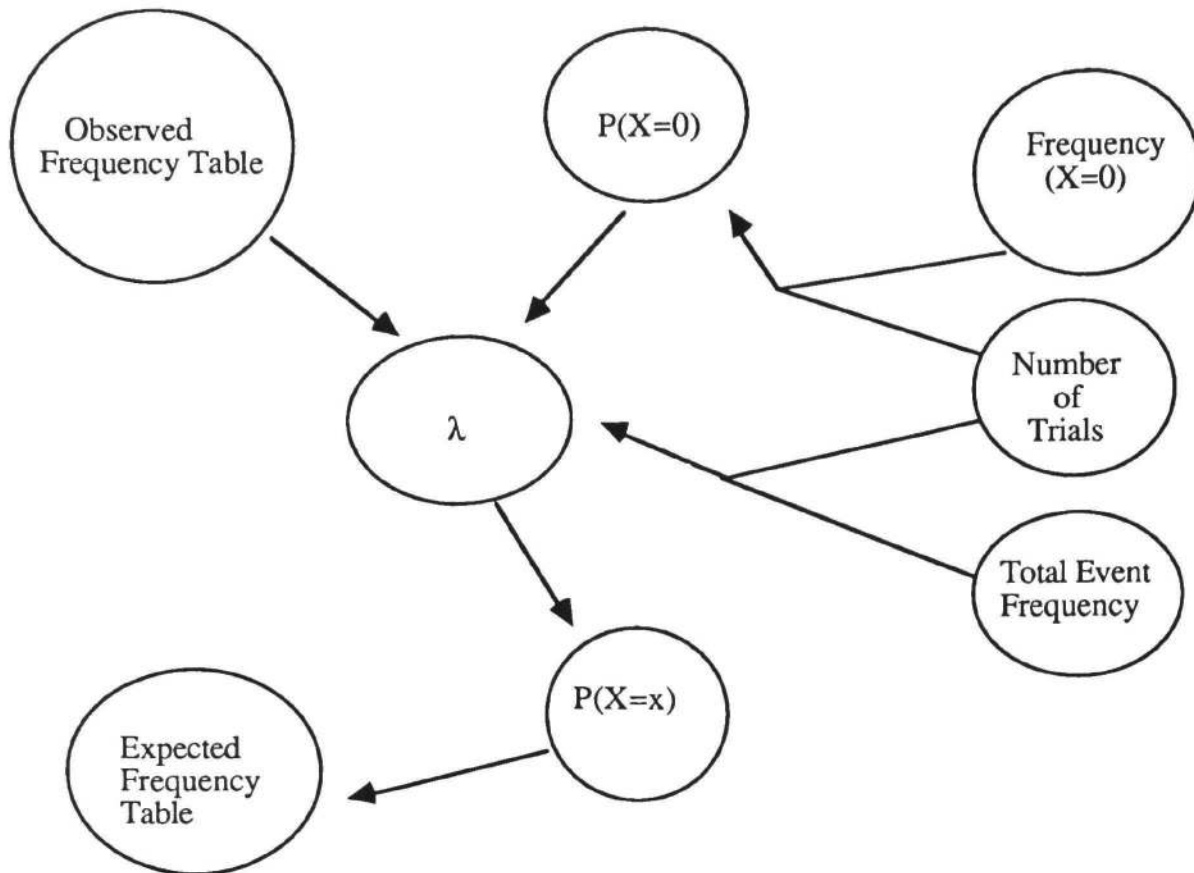


Figure 1
Problem Space for Poisson Distribution Problems

Figure 1 shows the problem space for Poisson distribution problems. The circles represent possible subgoals (or givens) and the arrows indicate which subgoals can be reached from other subgoals. The figure shows that λ can be reached in a few ways.

In our earlier experiments (Catrambone & Holyoak, 1987) we found that subjects who studied problems only of the quarry type recognized the subgoal of finding λ in the birthday problem about 30% of the time and the subgoal of finding $P(X=x)$ about 35% of the time. In addition, we found that subjects who studied problems of the quarry type and another type (in which $P(X=0)$ was given and the problem asked for the value of λ) recognized the subgoal of finding λ about 50% of the time and finding $P(X=x)$ 98% of the time. This suggests that studying multiple solution procedures helped make subgoals (such as finding λ) more apparent to learners. However, one thing these subjects could not do very well was figure out how to find λ . Half of the subjects knew that λ needed to be found but did not know how to get it other than to use the frequency table method demonstrated in the quarry problem. They were not able to modify this method.

Current Experiment

In the current experiment, we solicited paid volunteers from an upper-level probability course at the University of Michigan. These students had been taught about random variables and were in the process of learning the binomial distribution. Students participated in the present experiment,

CATRAMBONE, HOLYOAK

which dealt with the Poisson distribution, while learning the binomial distribution but before learning the Poisson. We sought to manipulate more systematically the methods for finding λ that subjects studied and to increase the number of examples of each method. An extensive discussion of the design of the experiment is in Catrambone and Holyoak (in preparation). Below is an outline of the relevant aspects of the design.

Procedure

Subjects studied 10 examples which demonstrated different methods for finding λ . Methods for finding λ are listed below. Some subjects learned two methods and other subjects learned three methods. No subject studied methods #4 or #5. In the test phase, subjects solved 10 problems using all five λ -methods.

The five methods for finding λ are:

- 1) λ given directly in the problem
- 2) Calculated from an observed frequency table
- 3) Calculated from the Poisson equation if $P(X=0)$ is given
- 4) Calculated from the Poisson equation if the frequency of $X=0$ is given and the total number of trials is given
- 5) Calculated by dividing the frequency of the event by the total number of trials

Results and Discussion

Figure 2 presents the percentage of subjects that used λ -methods #4 and #5 correctly in the test phase as a function of the number of λ -methods they studied in the training phase. Subjects in both groups were quite good using λ -method #4. This result is not too surprising since λ -method #4 is quite similar to λ -method #3 and all subjects studied λ -method #3. The only difference in the methods is that #4 requires the solver to find $P(X=0)$ rather than being given it directly. The prior studies had indicated that subjects were quite good at going from a frequency to a probability. A more interesting result is that subjects used λ -method #5 correctly about 75% of the time. This method is a simpler version of λ -method #2. In #2 the solver must find the frequency of the event whereas in #5 the frequency is given directly. In the prior studies, subjects were typically quite unsuccessful in adapting λ -method #2 into #5. The highest success rate of any of the earlier groups was 53%. The performance of subjects in the current experiment was significantly better than that (75% vs 53%, $z = 2.44$, $p < .02$).

This would suggest that the more extensive practice with example problems benefitted subjects in the current experiment. The manipulation of two versus three types of methods did not seem to make a great deal of difference. It would seem that extensive practice with at least two methods is the critical feature for being able to modify the old methods. The particular methods that are learned also matters (see Catrambone & Holyoak, in preparation, for discussion of this issue).

The prior experiments combined with the current one suggest that subjects who studied example problems that illustrated multiple methods for finding a particular subgoal were able to recognize that subgoal more successfully than subjects who only studied one procedure. There are several explanations for this result. One explanation is related to the work of Sweller and his colleagues (Mawer & Sweller, 1982; Sweller & Levine, 1982; Sweller, Mawer, & Ward, 1983) that indicates that when problem solvers are led to focus on a single ultimate goal, they concentrate on reducing the distance between the initial state and the goal state and thus fail to recognize subgoals and

methods in the domain. The one-procedure subjects in the earlier experiments learned a single concrete goal when they studied the example problems such as taking a frequency table as input and, through a stereotyped series of operations, producing an expected frequency table as output. They did not form intermediate goals such as finding λ or a single $P(X)$. The multiple-procedure subjects, by virtue of seeing the same unknown (such as λ) reached in more than one way, were able to recognize that unknown as a subgoal. The extensive practice on two or more methods by subjects in the current experiment helped them to modify old methods better than subjects in the earlier studies.

Extensive practice with multiple methods appears to be an important component in enabling learners to recognize subgoals or modify methods, at least to some degree. That is, learners are reasonably successful in altering methods if they have had a fair amount of practice with two or more. It will be important to continue to examine the relative contributions of variety and amount of practice on learners' ability to isolate subgoals and methods and to modify the methods for a variety of domains. In addition, it will be important to study how practice allows people to modify methods.

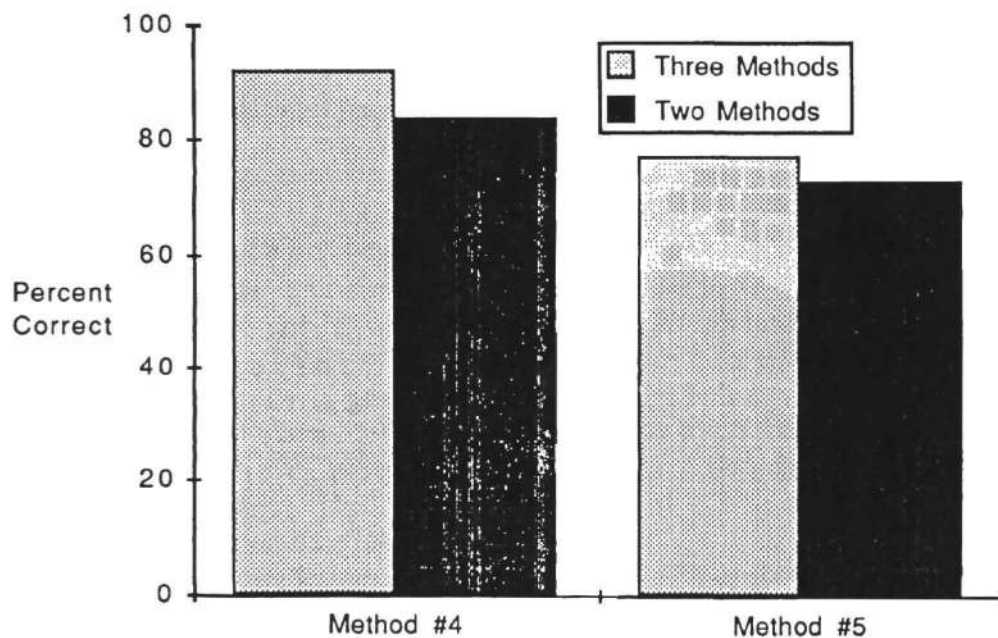


Figure 2
Success at Using New λ -Methods as a Function of the Number of
Methods Studied During Training

CATRAMBONE, HOLYOAK

References

- Adelson, B. (1981). Problem solving and the development of abstract categories in programming languages. Memory & Cognition, 9, 422-433.
- Anderson, J.R., Farrell, R., & Sauers, R. (1984). Learning to program in LISP. Cognitive Science, 8, 87-129.
- Catrambone, R., & Holyoak, K.J. (1987). Transfer in problem solving as a function of the procedural variety of training examples. In Proceedings of the Ninth Annual Conference of the Cognitive Science Society. Hillsdale, NJ: Erlbaum, 36-49.
- Catrambone, R., & Holyoak, K.J. (in preparation). The impact of procedural variety of training examples on problem-solving transfer.
- Charney, D.H., & Reder, L.M. (1986). Initial skill learning: An analysis of how elaborations facilitate the three components. (Tech. Rep. No. ONR-86-1). Pittsburgh: Carnegie-Mellon University.
- Mawer, R., & Sweller, J. (1982). The effects of subgoal density and location on learning during problem solving. Journal of Experimental Psychology: Learning, Memory, and Cognition, 8, 252-259.
- McKeithen, K.B., & Reitman, J.S. (1981). Knowledge organization and skill differences in computer programmers. Cognitive Psychology, 13, 307-325.
- McKendree, J. (1986). Impact of feedback content during complex skill acquisition. Ph.D. Dissertation, Carnegie-Mellon University.
- McKendree, J., Reiser, B.J., & Anderson, J.R. (1984). Tutor goals and strategies in the instruction of programming skills. In Proceedings of the Sixth Annual Conference of the Cognitive Science Society, Boulder, CO, 252-254.
- Pirolli, P.L., & Anderson, J.R. (1985). The role of learning from examples in the acquisition of recursive programming skill. Canadian Journal of Psychology, 39, 240-272.
- Sweller, J., & Levine, M. (1982). Effects of goal specificity on means-ends analysis and learning. Journal of Experimental Psychology: Learning, Memory, and Cognition, 8, 463-474.
- Sweller, J., Mawer, R.F., & Ward, M.R. (1983). Development of expertise in mathematical problem solving. Journal of Experimental Psychology: General, 112, 639-661.
- VanLehn, K. (1982). Bugs are not enough: Empirical studies of bugs, impasses, and repairs in procedural skills. The Journal of Mathematical Behavior, 3, 3-71.
- VanLehn, K. (1985). Arithmetic procedures are induced from examples, (Tech. Report No. ISL-12). Xerox Palo Alto Research Center.