REPORT DOCUMENTATION PAGE

1a REPORT SECURITY CLASSIFICATION	1b. RESTRICTIVE MARKINGS
Unclassified	

2a. SECURITY CLASSIFICATION AUTHORITY	3. DISTRIBUTION/AVAILABILITY OF REPORT
	Unlimited
2b.	

4. PERFORMING ORGANIZATION REPORT NUMBER(S)	5. MONITORING ORGANIZATION REPORT NUMBER(S)
	AEOSR-TR- 93 0009

6a. NAME OF PERFORMING ORGANIZATION	6b. OFFICE SYMBOL (If applicable)	7a. NAME OF MONITORING ORGANIZATION
Department of Psychology		Air Force Office of Scientific Research/NL

6c. ADDRESS (City, State and ZIP Code)	7b. ADDRESS (City, State and ZIP Code)
University of Illinois 06 S. Wright Street Urbana, IL 61801	Bolling AFB, DC 20332-6448

8a. NAME OF FUNDING/SPONSORING ORGANIZATION	8b. OFFICE SYMBOL (If applicable)	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER
AFOSR	NL	AFOSR 89-0447

8c. ADDRESS (City, State and ZIP Code)	10. SOURCE OF FUNDING NOS			
110 Duncan Avenue, Suite B115 Bolling AFB DC 20332-0001	PROGRAM ELEMENT NO.	PROJECT NO.	TASK NO.	WORK UNIT NO
	61102F	2313	A4	

11. TITLE (Include Security Classification)

Reminding-based learning

12. PERSONAL AUTHOR(S)

Brian H. Ross

13a. TYPE OF REPORT	13b. TIME COVERED	14. DATE OF REPORT (Yr., Mo., Day)	15. PAGE COUNT
Annual Technical	FROM 1/21/92 TO 1/20/93	February 1, 1993	15

16. SUPPLEMENTARY NOTATION

17. COSATI CODES			18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)
FIELD	GROUP	SUB. GR	
5	09		

19. ABSTRACT (Continue on reverse if necessary and identify by block number)

In learning new cognitive skills involving problem solving, novices are often reminded of earlier problems. The use of earlier problems is a common means of problem solving and affects the learning of the skill. This project has three aims in understanding this learning. First, the representation of the resulting generalizations is being examined. Generalizations learned from remindings are likely to be conservative, in that they may be more tied to the examples than many current theories allow. A main aim of the project is to distinguish and test different forms of this conservatism. Second, the development of problem solving expertise is examined by focusing on differences in how typical and atypical problems are solved. Third, the effects of such reminding-based learning in everyday problem solving is examined to extend the findings and test some theoretical ideas that are difficult to investigate in more formal domains. This report provides an overview of this work and the progress on these objectives during the last year.

20. DISTRIBUTION/AVAILABILITY OF ABSTRACT	21. ABSTRACT SECURITY CLASSIFICATION
UNCLASSIFIED/UNLIMITED ☑ SAME AS RPT. ☐ DTIC USERS ☐	Unclassified

22a. NAME OF RESPONSIBLE INDIVIDUAL	22b. TELEPHONE NUMBER (Include Area Code)	22c. OFFICE SYMBOL
Tangney	(202) 767-5021	NL

DD FORM 1473, 83 APR   EDITION OF 1 JAN 73 IS OBSOLETE.   UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS

ANNUAL TECHNICAL REPORT: REMINDING-BASED LEARNING

In this annual report, I have repeated the overview and background needed to understand the progress made since the last report. Although this strategy makes this report partially redundant with the 1992 report, it allows the reader to understand this report without referring back to the earlier report for clarification.

OBJECTIVES

When learning a new cognitive skill, novices spend much of their time solving problems. In doing so, it is common for novices to think back to an earlier problem that the current problem reminds them of and use this earlier problem to help solve the current problem (e.g., Ross, 1984, 1987, 1989a). This use of the earlier problem not only affects performance on the current problem, but also provides the learner with additional knowledge that can be accessed and used on later problems. The aim of this research is to understand the nature of the learning that results from this use of earlier problems. Little is known about this crucial source of learning.

Such within-domain analogies occur frequently during learning. In the view presented in this work, a generalization is formed from making an analogy between problems. Rather than positing a separate generalization process that operates upon completed instances, the generalization may be a byproduct of the analogy. In using the earlier problem to help solve the current problem, comparisons must be made and some aspects generalized over. Remindings, by setting up the analogy, may determine what pairs of problems are compared and, hence, what generalizations are made. The learning comes about because, while the noticing might be based on a variety of similar aspects (including superficial ones) between the problems, the comparisons forces the generalization of many of the aspects. My earlier work (e.g., Ross & Kennedy, 1990) has shown that the use of earlier problems allows novices to begin to form generalizations across problems. Thus, this means of learning may be one way in which novices can begin to develop more expert-like knowledge structures. However, much work remains to be done to understand this learning. The research I have been conducting has three goals.

The first goal is to understand the nature of the resulting generalization, the information included in its representation. If the generalization results from the reminding, it is likely to be a conservative one, in that it will be somewhat tied to the problems from which it arises. In fact, most theories of learning assume some conservatism, but my research (a) distinguishes among different types of assumptions of conservatism, (b) relates these distinctions to current theories, and (c) tests these distinctions. Thus, this work examines the specificity of what is learned, as well as its generality.

The second goal of this research is to examine the implications of these ideas for the development of expertise. Research on expertise in mathematics and physical science domains suggests that experts have problem schemas that allow them to categorize problems, as well as associated procedures for solving problems of that type (e.g., Chi, Feltovich, & Glaser, 1981). Despite the importance of these schemas, little research examines how they are learned. A common idea is that they may develop from the comparison of problems.

However, this idea leads to two questions? One, how do people know which problems to compare? Two, why are people comparing problems (i.e., what is the nature of this comparison)? The reminding-based learning view suggests that people compare problems when one makes them think back to another and they do so in order to use the earlier one to solve the current problem. In addition, this view suggests that people may develop problem schemas that are influenced by the superficial aspects, because these aspects are known to affect remindings. Thus, part of this project is concerned with the development of such problem schemas and the possibility that some schemas may critically depend upon superficial contents, even in experts.

The third goal of this project is to begin to extend this work to more everyday problem solving situations. This extension will not only allow the application and test of these ideas in an important new setting, it will also force the extension of this work to important situations that are hard to experiment with in more formal domains.

<center>STATUS OF PROJECTS</center>

In this section, I will provide some details of the status of the projects underway. For each goal, I will first give a brief summary of the findings and then provide further details.

1. Conservatism of learning

As described earlier, this project examines the representation of the resulting generalization. The focus here has been on asking how irrelevant aspects may affect the access and use of the generalization.

Summary. A common theme across many learning models is that generalizations are accessed solely by those features that are contained in the generalization (i.e., the features common across the instances). One experiment shows that people may first access instances to then provide access to the generalizations. That is, a manipulation that affects the access of instances improved performance even though if the instance were being used directly, performance would not have been increased (the results are below). This result, should it hold up in replications and extensions, will be difficult for many current theories to account for. In addition, another project has been examining how the representation of the generalization may depend upon the details of how the earlier problem is used. This second project has led to an unanticipated third project that investigates the interactions among the different processes involved in reminding, analogical transfer, and generalization.

Method. The basic paradigm for examining the effect of using an earlier problem is referred to as the cuing method. The typical experiment involves study examples, first test problems (with cues to the study examples), and second test problems. This simple case allows us to isolate the specific influences on each test and to detail the learning on a step-by-step level. More specifically, subjects are instructed in a simplified formal domain, elementary probability theory. They are given a short introduction to some probability concepts and then learn four principles (e.g., permutations). For each principle, an explanation of the principle with the appropriate formula is given. The subject then helps to

solve a word problem that requires the use of the principle. This study problem is presented in a workbook format, with the subjects guided through the solution. After learning these principles (4 mins/principle), subjects are given first test problems to solve (3 mins). In many of the studies, some or all of these first test problems include a cue (e.g., "this is like the earlier golf problem"), which has been shown to result in a generalization (Ross & Kennedy, 1990). After each problem, feedback is given and subjects study the solution. A second test is then given. Of main interest is performance on the second test problem as a function of the experimental manipulation. The second test measure varies for different experiments, but for all but one of the experiments (to be noted) in this section, the measure is how able subjects are to instantiate a provided formula (i.e., they are given the appropriate formula and have to fill in the numbers correctly for the second test problem). Earlier studies have shown this use measure to be very different from the ability to determine the appropriate formula and it appears to be a function of subjects' understanding of what the variables represent. In all of these experiments, principles are rotated through conditions and several test examples are used for each principle to avoid effects due to any particular example.

1a. Can irrelevant distinctive aspects provide access to the generalization? One study provides a test of how instances and the generalizations they give rise to are related. That is, the proposal provides an analysis of a number of ways in which the instances and the generalization might be "connected". One notion of conservatism is that the generalization is still closely tied to the instances. In particular, it is possible that even some superficial aspects that occur in only one of the examples may still provide access to the generalization. This idea, if true, is problematic for many theories because they assume that the generalization includes only features that were common to the instances that gave rise to the generalization. Thus, if the superficial aspect occurred in only one instance, it could not have been in common and should not be part of the generalization. The difficulty with testing this idea is that an experimental manipulation that helps to access an instance may increase performance because the instance is used rather than because the generalization is used. (Similar to the exemplar accounts of many of the prototype effects in categorization). Thus, what is needed is a situation in which direct use of the instance hurts performance, while indirect use of the instance to access the generalization helps performance. That was the purpose of this study, which was included in my first annual report.

That experiment made use of an earlier manipulation (Ross, 1987, 1989b) in which object correspondences were shown to affect how people instantiated a formula. More specifically, if novices learn about one principle in which the variables are attached to certain objects, a later test with the same formula will often lead them to try to instantiate the formula by matching the objects. For example, if permutations were learned with the number of objects (n) being mechanics and the number of objects selected (r) being cars, at a test with the formula and a word problem involving mechanics and cars the novices would be likely to again assign mechanics to the variable n and cars to r. This assignment occurs even if the word problem has been changed to make the reversed assignment correct. Thus, these earlier studies use this manipulation of reversed object correspondences to show that the instantiation of the formula relies on superficial aspects. In fact, the studies in Ross (1989b) show that this superficial similarity does not have to be at the level of mechanics and cars, but rather at the level of animate and inanimate. That is, novices will also reverse

correspondences if the test problem has teachers (rather than mechanics) assigned to classrooms (rather than cars).

The reason why I have gone to this level of detail is that these reversed correspondences allow a test of the issue under consideration here. In particular, assume that we have used the cuing paradigm to get generalizations. Then, we manipulate the superficial similarity of the story line to affect the likelihood of accessing a particular instance. If the instance being accessed has reversed correspondences to the test problem then using it will hurt performance. Thus, if it helps performance the instance must not have been used directly, but rather indirectly to access the generalization. The following table may help to make this concrete. The examples are written with the story content then, in parentheses a type of object correspondence in which different numbers mean unrelated correspondences and R means reversed. In both conditions, the first test problems were cued to increase the probability that a generalization was formed between the study and first test problems (as shown in Ross & Kennedy, 1990). The subjects received one of the two conditions for each principle, but had two principles in each condition.

Study example	First Test	Second Test (with formula)
Golfers (1)	Mechanics (2)	Golfers (1-R)
Golfers (1)	Mechanics (2)	Dancers (1-R)

Consider the difference between the two conditions, which shall be called the golfers and dancers. After the first test, I assume that subjects have some information about each of the first two problems and a generalization formed by comparing the two problems (Ross & Kennedy, 1990, showed that this manipulation led to a generalization). On the second test, the golfers problem is likely to make subjects think back to the earlier study example (as has been shown a number of times in my earlier work). However, if subjects use this study example it will lead to poor second test performance because the correspondences are reversed. Thus, most theories would predict that golfers will lead to worse performance than dancers (if the study example is used) or equal performance (if only the generalization is used). However, if the study example could be used to access the generalization, then performance in golfers might be higher than in dancers.

In the experiment, the results confirmed this last prediction. A study with 24 subjects each learning 4 principles showed that performance in the golfers condition was .71 (i.e., 71% of the formulae were appropriately instantiated) compared to only .53 in the dancers condition. (Technically, the $t(23) = 1.69$, which is not reliable by a two-tailed test. However, this is due to one subject who showed the exact opposite effect and increased the standard deviation by 25%. I have examined the data in a number of ways and any way that reduces that subject's influence leads to a highly reliable effect. In addition, this effect occurs for all 4 of the principles).

A parallel experiment tested the materials to show that in fact the reversed correspondences did not inadvertently help performance. In particular, this exact experiment was conducted on another 24 subjects but instead of receiving the first test problems, they

participated in a filler task for the same amount of time. For the first experiment effect to be attributable to the indirect access of the generalization, it is necessary to check that the golfer problem does not lead to higher performance than the dancers problem when no generalization is likely to be formed. As expected, the golfers problem led to slightly worse performance than the dancers problem (.31 vs. .36). A final experiment replicated the first study with the only change being a 20 minute delay between the first test and the second test. Here too the golfers problem led to higher performance, but the difference was only .40 vs. .35 and was not reliable.

In last year's report, I explained two variations of this experiment I had tried, both of which led to predicted, but unreliable effects. In the last year, I attempted what I thought was a more sensitive design, but again found small effects. This difficulty is partly due to the fact that these results for learning are necessarily more difficult to find than the earlier work on problem solving, because they concern the effects of this problem solving. Although this difficulty could be alleviated by testing (many) more subjects in the experiments, this would force me to curtail much of my other work. (Subject availability is good, but not unlimited.) I have instead decided to try to rethink how to approach this problem during the next couple of months and see if I can find a more sensitive way of examining the same issue. I believe that these studies are important for understanding the nature of the generalization and its hypothesized conservatism. Although this recent lack of statistically significant results requires me to change my plans somewhat, the consistency of the effect is at least suggestive that the main idea may be on track. I hope to be able to continue to address this issue, though with a more sensitive approach. I am disappointed in not having been able to accomplish more on this project.

<u>1b. The effect of problem use on the generalization</u>. In addition to testing the conservatism of the generalization in the way described above, a central idea of the reminding-based view is that the generalization depends upon the earlier problem comparison. This <u>selectivity</u> effect has been investigated by me in the context of category learning (Ross, Perkins, & Tenpenny, 1990), because in formal domains different earlier problems will often lead to very similar generalizations. However, in this study I examine whether the details of <u>how</u> the problem is used (rather than <u>which</u> problem) affects what is learned. As will be seen, it can be viewed as a type of transfer-appropriate processing in problem solving.

(i). Problem categorization: The earliest studies using the cuing paradigm showed that cuing on the first test led to improved performance on the second test in three ways: better problem solving performance, higher probability of accessing the appropriate formula, and better instantiation of a provided formula (Ross & Kennedy, 1990, Exs. 1 and 2). After conducting these studies, I also became interested in problem categorization, which is thought to be a crucial element in developing expertise. Thus, I used the cuing paradigm but instead of a second test problem, subjects were given 12 problems (3 of each principle) and asked to categorize them by principle. That is, at the first test, half of the principles were cued and half were not. If cuing affects problem categorization, then higher categorization should occur for the cued principles. In one study, I gave them formulae to use as the categories for sorting and in another I gave them the principle names. The (unexpected) result was that the cuing had no effect.

In my first annual report, I reported a study that investigated one difference between the cue method and the usual way in which people think of earlier problems. Usually when people are reminded they need to decide whether the reminding is appropriate. That is, they need to decide whether to use the reminding or not, since we often get reminded of earlier problems for inappropriate reasons. It seems to me that this difference may be crucial because a determination of appropriateness may include deciding whether the problems are of the same "type" so that one could be used to help solve another. Thus, such a determination might affect what is learned about problem categories.

The study used the same procedure and materials that led to no cuing effect and made one change. On the principles that subjects are cued, they are told there is a chance that these cues are not correct. Their task is to first decide if the cue is correct or not. They are then given feedback on this response and given the correct cue. In fact, all the cues are correct so they are given exactly the same cues as in the earlier study, but they have to determine appropriateness before using the cued problem. The results in this case showed an effect of cuing, with the cued principles leading to .47 of the second test problems being correctly categorized compared to .30 for the not cued principles, $t(29) = 2.17$.

In my earlier report, I mentioned some difficulty I had had replicating this result (again, getting only small effects). However, I now believe that part of this difficulty was due to the materials. As I mentioned last year, the materials I have used in these studies were not designed originally for a categorization test and they are quite dissimilar across the different principles. The problems were rewritten to be a little more confusing and the results were more promising. In the last year, my graduate RA, Marty Preslar, has finished running this experiment (and is writing it up for a master's thesis). The results are encouraging: the probabilistic cue subjects showed a small, but consistent and significant effect of cuing, .47 vs. .40, $t(47)=2.77$. The subjects receiving the usual cues, showed a slight, nonsignificant, effect in the other direction, .413 vs. .455.

Preslar is leaving the graduate program at Illinois after he gets his masters, but a new student, Matt Kilbane, will be picking up this project. We are making some minor changes in the materials and procedure and will then replicate this experiment. An additional interesting result that we hope to examine here is that almost all of the cued-uncued difference in the probabilistic group occurred for the subjects having the most difficulty (i.e., the "poor" learners).

Assuming it works, we will have good evidence that how the earlier problem is used affects what is learned. This has been an important motivation for a new project I have started (see 1c).

(ii) Direct reports on what was learned: The earlier studies showing an effect of the cuing on what was learned, examined problem solving performance. In the last two years, I have conducted two new studies. In this work, at the second test, the subjects were directly queried on what they knew about each principle. A pilot study with 8 subjects (with responses tape-recorded and transcribed) suggested an interesting result. The cuing led to no better conceptual understanding of the principle (as measured by the explanation of when the principle might be sued and what it did). However, when the subjects were asked to explain

the formula, the cuing condition led to considerably better understanding of what the variables were and how they would be instantiated. This is an interesting result, in that it suggests that cuing may selectively help the novice in learning certain aspects of the principles. In particular, the low level information about variables may be learned so that later problem solving performance is improved. (More speculatively, I believe that this result may give us some insight as to how novices may use some knowledge to bootstrap their learning of other, more conceptual knowledge. In this scheme, low-level knowledge about variables may provide some later clues as to how the principle is working in a conceptual way.)

Last year, Marty Preslar followed up this study with a much larger one (n=24) in which we also asked the subjects to generate the formula. The results have been scored extensively and are quite similar to the pilot work. The explanations are no better in the cuing condition (just as we found with the pilot study). However, cuing does lead to better understanding of the variables (again as in the pilot work) and to better generation of the formula. Preslar presented this work at MPA in May.

1c. Interactions among different processes. The work on selectivity can be viewed in a slightly different way - the generalization depends upon the details of the problem solving. This idea has led me to go back and consider in much more detail how the different processes involved in reminding, application, and generalization may be interacting. The first studies I have conducted in this project are in collaboration with Prof. Gary Bradshaw, a tenured faculty member at University of Colorado who has been visiting at Illinois. (It appears he will be offered a job here, so I am anticipating a continued collaboration.) We have focused on remindings and transfer. All current views of analogical problem solving (including my own) assume the new (target) problem is represented and that this representation leads to a retrieval of some earlier problem (which is then mapped and transferred to the new target problem). However, there is also evidence that people may be reminded during the reading of a problem.

Our studies examine whether such remindings of earlier problems may mean that the target problem's representation is affected by the reminding. The importance of this point is that much of the focus on mapping as a separate process may be misplaced, IF the earlier problem is affecting how the target problem is represented (i.e., much of the transfer occurs DURING the target being represented, not between the represented target and the earlier problem). So, if this problem makes me think of an earlier one, I may fit the current problem to the earlier one. Bradshaw and I have begun to test this idea in a general way, using how simple stories are interpreted. Our results have been encouraging: we find a simple cue (a proper name) reminds people of an earlier story and affects their interpretation of the current one. In addition, we find that this is an encoding effect in that one can see the effect on sentence-by-sentence reading times. We have submitted this preliminary work as a paper to the Cognitive Science Conference (enclosed) and will, with one additional experiment, be writing it up for submission to a journal. After these general interpretation effects, we will begin to look at how the particulars of the target representation may be affected by the reminding. I am quite excited by this project and think it will be an important addition to our understanding of the reminding effects on problem solving and learning.

2. The development of problem solving expertise

As discussed earlier, reminding-based learning provides one perspective on the nature of problem schemas and their development. The work already discussed can be viewed as very early precursors to problem schemas, but no evidence has yet been presented that in fact they do lead to problem schemas.

My focus has been on the specialized schemas that often include superficial information as well. To study this, I have written a number of algebra word problems (e.g., distance, interest, mixture). Each problem has two versions: one with appropriate (typical) contents and one with neutral contents. The first two studies conducted have followed up observations noted in a chapter a number of years ago (Hinsley, Hayes, & Simon, 1977). First, we have protocols of subjects solving these problems in order to ask how they might differentially solve the appropriate and neutral problems. Hinsley et al. argued that the appropriate ones were solved by schema instantiation, while the others were solved by translating each sentence to an equation. This result is quite important, but their results were based on few observations and no quantitative data were provided. However, we were unable to replicate this finding, not because it is incorrect, but because so few of the problems were solved at all (less than 50%), making it impossible to see differences between conditions. We considered several training studies, but opted for what I think will be a more fruitful tack. An honors student, Steven Blessing (now at Carnegie-Mellon, working with John Anderson), has redone this experiment with three major changes. First, we have used high math students (UI students who graduated from a math and science academy). Second, he has included a set of inappropriate problems as well, that is problems of a given type with the story contents of a different type. Third, we have examined the time and proportion correct. Our finding is that the appropriate and neutral conditions do not differ, but they are both more accurate and faster than the inappropriate condition. In addition, he has collected problem solving protocols from another group of subjects.

In the second study originally conducted, other subjects received the same problems one clause at a time and were required after each clause to say what type of problem it is. Hinsley et al. showed people could do this quite readily. Our interest is in asking how this measure may differ for appropriate and neutral contents. The appropriate problems are categorized much earlier.

In the last year, we have conducted several other studies. First, the high math subjects were tested in the clause study with appropriate, neutral, and inappropriate problems. We found that the appropriate were categorized faster than the neutral (as before) and the inappropriate were categorized slower than the neutral. The inappropriate contents often lead to subjects initially making the (content-appropriate) inappropriate categorization, before "seeing" the right category.

Second, the regular subjects (i.e., UI students not from the math academy) participated in the timed problem solving task. We found the same results as with the academy students (though the means were lower).

Together, these studies provide important information about the problem schemas and

their use, particularly concerning the inclusion and use of superficial aspects. Our current hypothesis is that the usual or appropriate contents allow one to more quickly categorize problems (i.e., the clause result), but that even neutral problems are generally categorized correctly by the end of the problem (another result from the clause study). Thus, there is little accuracy or latency difference in problem solving between these two conditions. However, the inappropriate problems are more difficult to categorize and the categorization is often uncertain even after the problem has been read.

However, what we have not been able to see is any evidence that the appropriate materials provide any benefit to how the problem is solved. Our hypothesis is that for these simple problems, the schema is recognized by the end of the reading even for the neutral (as we saw in the clause study). Thus, we wanted some manipulation in which the schema might provide some benefit if it were recognized early. We have decided to use more complex problems, in which we add irrelevant information. The idea is that if the solver has the appropriate schema, s/he may be better able to decide on the relevance of the different information in the problem. This manipulation has been quite successful. First, a group of the academy subjects are faster and more accurate in solving the appropriate problems versus the neutral problems (the first time this difference has come out). Second, Blessing collected some protocols and found that the appropriate problems led to subjects talking more about relevant information compared to the neutral problems (i.e., a higher proportion of the protocol statements concerned information that was relevant to the problem solution).

We are encouraged by these findings and believe that they may allow us to see some benefit of the content on how the problem is solved (not just on how the schema may be accessed). We are currently deciding on the next step, but a likely possibility is to get protocols from some real experts (algebra teachers) and examine them for whether the contents are used throughout the problem solving.

3. Everyday problem solving

In addition to the work on probability theory and formal domains, I have been examining how remindings may be used in less formal situations. In particular, I have been examining how they may affect the simple categories that are learned. I view this work as parallel to the problem solving work, but it sometimes allows me to more cleanly investigate certain issues, especially selectivity effects (Ross, Perkins, & Tenpenny, 1990).

Summary. The work from the last two years (with a graduate student RA, Tom Spalding) has extended this earlier paper to show (1) that reminding-based generalizations occur in a common categorization study paradigm, (2) that remindings serve to focus the learner on to relevant features, (3) that these manipulations lead to differences in perceived frequency, (4) that they also affect people's prediction of the values for missing features, and (5) that the effects of early focus features affects what is learned from later instances. Spalding has also been able to show that these effects hold up with different ways of getting subjects to think back to the earlier examples. These results greatly extend the earlier work and address many potential criticisms of it. I believe that this idea may help to provide one means by which categories can be learned and help relate such learning to performance issues. In addition to a chapter based on these ideas (Ross & Spalding, 1991), Spalding and

I have written a paper that is under review (Spalding & Ross, 1992) and have presented this work at MPA and Psychonomics. It is impossible to explain the most recent work without going through the background earlier work, so I repeat here much from the last report. (If you have read that recently, the newer stuff is the last few paragraphs.)

3a. Reminding-based generalizations in category learning. In this study, reported in the first report, we show that the manipulation of a feature that affects reminding (color), but which subjects know will not be included on tests, still affects what is learned about the category. As an abstract example, the letters a, b, c, d, e, f, g and h stand for features (e.g., a might be "likes gardening"). Subjects then learned about people from two categories. The structure for each category was as follows (the other category had the same structure but e, f, g, h occurred twice and a, b, c, d occurred once):

             Person 1 has features a b e
             Person 2 has features a c f
             "   3   "   "   d b g
             "   4   "   "   d c h

The experimental manipulation was that for half the subjects Persons 1 and 2 were presented in red and Persons 3 and 4 in green. For the other half of the subjects Persons 1 and 3 were in red and Persons 2 and 4 in green. Subjects knew that the tests would not include color. They were shown each of the 8 people (4 for the other category) once per block in an anticipation learning paradigm until a learning criterion was met. A number of different tests were then given. The results were quite clear: those features correlated with color (e.g., a and d for the first half of subjects) were viewed as more representative and important to the category. For instance, if a correlated feature from one category was presented with a non-correlated feature from the other category, .74 of the categorizations favored the correlated feature. Note that in all cases, the four critical features (a, b, c, d) occurred equally often with the category (and half as often with the other category). Thus, it appears that the reminding features of color were used to in some way "organize" the category and notice common features within each part. Last year, we replicated this study with a number of procedural changes, to insure its generality.

3b. Remindings serve to focus the learner. One hypothesis that has not been tested in the formal domains is that remindings help novices to focus on relevant features. That is, one difficulty novices have is that they are not sure what information is relevant. If a novice is reminded of an appropriate earlier problem, however, the commonalities are much more likely to be relevant (compared to the non-common aspects) and thus the novice may tend to focus more on these common aspects. The category paradigm presents a nice opportunity to examine this hypothesis. In this study, 40 subjects were presented three times with 9 members of a single category and told to try to learn about the category for a future test. As in the above experiment, color was correlated with some feature. In this study, it was correlated with the value of the first dimension or the second dimension. Each dimension had three values (e.g., hobby - painting, photography, weaving) and the three colors were presented with the same value each time. In addition, half the subjects l- rned instances with only two dimensions, while half learned instances with four dimensions (the values for the third and fourth dimensions were uncorrelated with other dimensions). The idea was that the

two dimension case was simple enough to learn completely, so any focussing effect of reminding was unnecessary. However, the four-dimension case was much more complicated and remindings would help learners to focus. For the test, subjects ranked how important the three values of dimension 1 and the three values of dimension 2 were (i.e., the 6 values were ranked 1 to 6). The results confirmed these predictions. In the two-dimension condition, those values correlated with color led to only a .1 rank difference (out of a possible 3). In the four-dimension condition, the difference was 1.2, with 18 of 20 subjects showing the effect. The interaction with the two-dimension condition was reliable as well.

We conducted several follow-ups to these studies during the last two years. Most important were four results using the one category paradigm. First we varied the frequencies (9, 12, or 15 times) of presented features, for both focussed and non-focussed. After study, subjects were presented with each feature and asked to judge its frequency. There were significant effects for frequency and for focussed vs. non-focussed. Focussed features were rated as more frequent at each level of presented frequency. This experiment shows that the effect is not due to the forced choice nature of earlier studies. In addition, frequency effects are often thought to relate closely to many basic memory effects, suggesting that the manipulation could have a number of implications for how the features are remembered and used.

Second, we conducted a more straightforward prediction experiment. In this study, each "person" had four dimensional features. In addition, the values on dimensions 1 and 3, as well as on 2 and 4, were perfectly correlated. The color was consistent with one or the other pair. At test, we would present, with no color, two of the dimensions filled in (e.g., the values for 1 and 2) and the subjects had to fill in the values for the other two dimensions (3 and 4, in this instance). Consistent with the earlier findings, subjects were much better able to fill in the values consistent with the focussed features, suggesting that they picked up the correlation for those feature pairs that were consistently colored (with an average difference of 4.46 out of a maximum of 24).

Third, we tested the generality of the idea by seeing if some salient feature, other than color (which is qualitatively distinct), would lead to the same effect. To test this, we ran another experiment like part of the focussed experiment, but used the similarity of names rather than color to be the reminding feature. Although the effect was not quite as large, we found an advantage of .686 out of 3 maximum, with $t(57)=2.78$.

Fourth, and new this year, we found that the "focus" features that are in common between compared instances have later effects. In particular, they can later affect what later instances are compared and thus, what is learned from these instances. This work is important for two reasons. One, it shows that the effects may occur throughout learning, not just during the early learning. Two, it brings the category work much closer to some of the work I have been doing in problem solving domains.

In addition to this study, Spalding and I have written the paper mentioned earlier and are about done with a handbook chapter (Ross & Spalding, 1993). We have also begun to examine how the features that are correlated over instances may be learned and used, not just for categorization, but for various types of inference (such as prediction). Again, the

motivation is to bring the category work to bear on issues we have been examining in the problem solving domain. For example, problem categories (as discussed in section 2 of this report) are viewed as essential in problem solving, but it is not enough to categorize a problem. One must then use the associated procedures for solving the problem. Similarly, within the concept learning work, we are going beyond the question of how an instance is categorized to investigate how the categorization is used. These studies add not only to the idea of reminding-based learning, but also to how categories might be learned.

General Summary

I have presented a brief summary of the results from the projects addressing the three object.. .s. As is evident, my work continues to address each of these objectives I have tried to outline for each where I see the work going in the next year. The projects that seem most ripe for advances are the problem categorization work (1b), the interaction among processes (1c), and the concept learning work (3). The work on problem solving expertise (2) has gone well, but I think we will probably finish up some of the work I have mentioned and then try to put the finished studies together. The work on irrelevant distinctiveness (1a) I still find exciting, but as I mentioned, I think it needs a new approach.

References

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.

Hinsley, D. A., Hayes, J. R., & Simon, H. A. (1977). From
words to equations: Meaning and representation in
algebra word problems. In M. A. Just & P. A.
Carpenter (Eds.), Cognitive processes in comprehension.
Hillsdale, NJ: Erlbaum.

Ross, B. H. (1984). The effects of remindings on learning a
cognitive skill. Cognitive Psychology, 16, 371-416.

Ross, B. H. (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, 629-639.

Ross, B. H. (1989a). Remindings in learning and instruction.
In. S. Vosniadou & A. Ortony (Eds.) Similarity and
analogical reasoning. Cambridge: Cambridge University
Press.

Ross, B. H. (1989b). 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.

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, 42-55.

Ross, B. H., Perkins, S. & Tenpenny, P. L. (1990) Reminding-based
category learning. Cognitive Psychology, 22.

Ross, B. H. & Spalding, T. L. (1991). Some influences
of instance comparisons in concept formation. In
D. Fisher, M. Pazzani,& P. Langley (Eds.) Concept formation:
Knowledge and experience in unsupervised learning.
San Mateo, CA: Morgan Kaufmann.

PUBLICATIONS ACKNOWLEDGING AFOSR SUPPORT

Ross, B. H., Perkins, S. & Tenpenny, P. L. (1990) Reminding-based category learning. Cognitive Psychology, 22.

Ross, B.H. (1990). The access and use of relevant knowledge: A specific case and general issues. In R. Freedle (Ed.) Artificial intelligence and the future of testing. Hillsdale, NJ: Erlbaum.

Ross, B. H. & Spalding, T. L. (1991). Some influences of instance comparisons in concept formation. In D. Fisher, M. Pazzani,& P. Langley (Eds.) Concept formation: Knowledge and experience in unsupervised learning. San Mateo, CA: Morgan Kaufmann.

(also Cognitive Science Technical Report from Beckman Institute).

Murphy, G. L. & Ross, B. H. (1992). Predictions from uncertain categorizations. Under review.

Spalding, T. L. & Ross, B. H. (1992). Comparison-based learning: Effects of comparing instances during category learning. Under review.

Ross, B.H. & Spalding, T.L. (1993). Concepts and categorization. To appear in R. J. Sternberg (Ed.) Handbook of Perception and Cognition, Vol. 12: Thinking and problem solving. San Diego, CA: Academic Press.

Bradshaw, G. L. & Ross, B.H. (1993). Reminding and transfer during encoding. Under review for Cognitive Science Conference.

PRESENTATIONS BASED UPON AFOSR FUNDED WORK

Reminding-based learning, AFOSR Workshop on Cognition, Arlington,
   VA, November, 1989.

Some influences of instance comparisons in concept formation.
   Invited speaker, Workshop on Computational Approaches to
   Concept Formation, Palo Alto, CA, January, 1990.

Reminding-based problem solving and learning. Cognitive Science
   Colloquium speaker, Georgia Institute of Technology.
   Atlanta, GA, April, 1990.

Reminding-based generalizations in categorization. Midwestern
   Psychological Association, Chicago, May, 1990.
   (Spalding & Ross, presented by Spalding).

Reminding-based problem solving and learning. Department of
   Psychology Colloquium speaker, McMaster University.
   Hamilton, Ontario. December 1990.

Focussing effects of reminding-based category learning.
   Midwestern Psychological Association, Chicago, May 1991.
   (Spalding & Ross, presented by Spalding).

Conservatism in learning. GUV meeting, Bloomington
   IN, September, 1991.

Concept formation and the learning of feature
   correlations. (T. Spalding & B. Ross), MPA, May 1992.

Learning from the use of earlier examples: Evidence about
   what is learned.   (M. Preslar & B. Ross), MPA, 1992.

Categorizing and solving algebra word problems. GUV
   meeting, Ann Arbor, MI, October 1992.

Instance comparisons during category learning.
   (T. Spalding & B. Ross), Psychonomics Society,
   St. Louis, November, 1992. [poster]

Reminding-based learning. University of Michigan,
   Ann Arbor, MI, December, 1992.

Remindings during learning. University of Chicago. Feb 1993.