Effects of Goal Specificity and Explanations on Instance Learning and Rule Learning

Rosemary J. Stevenson and **Bruce W. Geddes**
Human Communication Research Centre,
Department of Psychology,
University of Durham,
Durham, DH1 3LE, U.K.
[rosemary.stevenson, b.w.geddes] @durham.ac.uk

Abstract

We distinguish between instance learning and rule learning (e.g. Shanks & St. John, 1994). Instance learning involves memorizing learning instances while rule learning involves the abstraction of an underlying rule. Instance learning and rule learning can be explained by a dual space model of learning (Klahr & Dunbar, 1988; Simon & Lea, 1974). In relation to Simon and Lea's model, instance learning can be said to occur in instance space while rule learning makes use of both instance space and hypothesis space. We describe an experiment to test the view that whether instance learning or rule learning occurs depends on the learning goal and on whether or not the subjects explain what they are doing. Subjects were asked to learn a dynamic computer control task guided by either a specific or a non-specific goal. During learning, subjects also carried out a secondary task. They either described what they were doing during learning or explained what they were doing. We predicted that giving descriptions would favour instance learning and prevent rule learning irrespective of the learning goal, since giving descriptions forces subjects to focus on the task itself. Giving explanations should favour rule learning when subjects are given a non-specific goal, since both the non-specific goal and giving explanations focus on the reasons for the computer's behaviour. Giving explanations should not lead to rule learning when subjects have a specific goal since the specific goal forces subjects to focus on a search of instance space and to neglect the hypothesis space. The results confirmed these predictions. They support the view that goal specificity guides learning by directing attention to either instance space or both instance space and rule space, and that giving explanations encourages the revision of hypotheses in the light of the evidence.

Introduction

A number of researchers have proposed that there are two separate learning systems, one that learns instances and another that learns rules (Shanks & St. John, 1994). Instance learning comes about through the memorization of learning instances (e.g. Broadbent, Fitzgerald & Broadbent, 1986; Dienes & Fahey, 1995), and is best modeled by a PDP connectionist network (e.g. Cleeremans, 1993). By contrast, rule learning comes about through hypothesis generation and testing (e.g. Simon & Lea, 1979; Klahr & Dunbar, 1988) and may be best modeled by a serial symbolic network. In this paper, we describe an experiment to show that whether subjects learn specific instances or general rules depends on the cognitive processes induced by the learning goal, specific or non-specific, and on whether or not the learners explain what they are doing while learning.

In a previous study Geddes and Stevenson (in press) showed an effect of goal specificity on whether instance learning or rule learning occurred. Geddes and Stevenson used one of Berry and Broadbent's (1984) dynamic control tasks, which showed an apparent dissociation between learning and awareness. The task required subjects to interact with a 'computer person' called Clegg and try to get him to become and stay *Very Friendly*. Clegg initiated the interaction by displaying one of twelve attitudes (e.g. *Polite, Very Friendly, Loving*) on the computer screen, after which the subject had to respond by typing in another attitude. The attitudes reflected an intimacy scale from low to high and Clegg's response to the subject's choice of attitude was retaliatory. If Clegg had typed *Polite*, and the subject responded with *Friendly*, then Clegg would retaliate with the attitude *Loving*. Clegg's attitude on each trial was a simple numerical function of the subject's response on that trial and Clegg's previous output. Subjects successfully learned to carry out this task, but when questioned about the experiment afterwards, they were unable to describe what they were doing or what the underlying rule was.

In Geddes and Stevenson's study, one group of subjects was given a specific learning goal, comparable to the learning goal used in Berry and Broadbent (1984). Subjects were instructed to make Clegg polite and stay polite. However, in contrast to Berry and Broadbent, Geddes and Stevenson gave a second group of subjects a non-specific learning goal. These subjects were instructed to find out the pattern that explained Clegg's behaviour.

All the subjects had 30 learning trials, after which they were tested on what they had learned. In the first test, subjects in both goal groups were given 30 trials to learn a novel specific goal - to make Clegg very friendly. The results showed that non-specific goal subjects performed better than specific goal subjects with the novel specific goal (52% correct responses vs. 41%). In a second test, all subjects predicted Clegg's response, given a sequence of three responses. For example, a subject might be told "You were *very cool*, then Clegg was *very rude*, You were then *polite*. What did Clegg do next?" Some of these prediction questions described 'old' situations, which the subject had encountered during learning. Others described 'new' situations, which the subject had not seen before. Non-specific goal subjects made correct predictions in both old and new situations while specific goal subjects only made correct predictions in old situations. In a third test, subjects were asked to describe the rule that governed Clegg's behaviour. While 79% of the non-specific goal subjects gave either correct or partially correct rule descriptions, over 80% of the specific goal subjects gave wrong descriptions.

Thus, subjects given a non-specific goal learned the abstract rule underlying Clegg's behaviour while subjects given a specific goal remembered specific responses. These results are consistent with other evidence suggesting that the learning goal can have profound effects on learning, whether it be instance learning (Whittlesea & Dorken, 1993) or rule learning (Owen & Sweller, 1986; Sweller, 1988; Vollmeyer & Burns, 1995; Vollmeyer, Burns & Holyoak, 1996).

"Dual space" models of learning explain rule learning and instance learning within a single framework (Klahr & Dunbar, 1988; Simon & Lea, 1974). Simon and Lea proposed that the problem space is separated into two spaces: a rule space and an instance space. People search instance space when seeking the solution to a specific goal. Geddes and Stevenson suggested that one way in which instance space is searched to reach a specific goal is through means-ends analysis (Newell & Simon, 1972)[1] Means-ends analysis involves successive reductions of the difference between the learner's current state and the goal state until the goal is reached. Heuristic strategies are usually employed to bring about difference reduction. For example, subjects may decide to give an extreme response or to give a response midway between Clegg's last response and the target response. But regardless of the precise way in which instance space is searched, what gets learned are the specific states encountered on the route to the goal. In hypothesis testing, people search both rule space and instance space. Explicit hypotheses are generated in rule space, which are then tested by experiments that generate states in instance space. In these circumstances, subjects learn rules that explain the system being studied. Klahr and Dunbar (1988) have adopted a similar model, consisting of hypothesis space (comparable to Simon and Lea's rule space) and experiment space (comparable to Simon and Lea's instance space). As in Simon & Lea's model, hypotheses are generated and modified in hypothesis space and tested in experiment space. On the basis of these models, we suggest that a specific goal induces a search through instance space while a non-specific goal induces a search of both instance space and hypothesis space. (See also Vollmeyer, Burns & Holyoak, 1996.)

The experiment reported here has two aims. First, it tested the view that the learning goal influences cognitive activities by directing attention to one or both of the two problem spaces. Second, it examined the impact of explanations on learning. To test our interpretation of the cognitive processes induced by each goal, subjects carried out a secondary task during the learning phase. The secondary task required the subjects to talk aloud while learning and it was either compatible or incompatible with the hypothesized learning processes associated with each goal. Half the subjects described *what* they were doing. The remaining subjects explained *why* they were doing what they were doing. We hypothesized that giving descriptions would be compatible with the learning processes employed by specific goal subjects. These subjects should readily be able to describe what they were doing to reach the goal. Thus, giving descriptions should facilitate instance learning induced by a specific goal since giving descriptions maintains the subjects' search of instance space. However, giving descriptions should impede rule learning by non-specific goal subjects because the need to describe what is being done should deflect attention away from the rule space. On the other hand, giving explanations should be compatible with the learning processes employed by non-specific goal subjects, since the rule to be learned is the one that explains the computer's behaviour. Thus, giving explanations should facilitate rule learning induced by a non-specific goal, by reinforcing the subjects' search of rule space.

To examine the impact of explanations on learning under specific and non-specific goals, we focused on the explanation conditions. The performance of the non-specific goal subjects in Geddes and Stevenson's study was very good, but it was not optimal. As Vollmayer, Burns and Holyoak (1994) point out, inducing subjects to engage in hypothesis testing does not necessarily mean that they will use optimal strategies of hypothesis testing. Vollmayer et al found that instructing subjects in efficient hypothesis testing enhanced performance on a novel specific goal, irrespective of whether the initial learning goal was specific or non-specific. Explanation learning has also been shown to be a powerful mode of learning (e.g. Chi, de Leeuw, Chiu, & LaVancher, 1994; VanLehn, & Jones, 1993) and so might also be expected to increase the efficiency of hypothesis testing. In our study, we expected that explanations would make it clear to the learner when he or she did not fully understand the rule and thus motivate a further search of hypothesis space to modify or refine the current hypothesis. Consequently, we predicted that non-specific goal subjects who gave explanations would outperform non-specific goal subjects who gave descriptions.

Method

Subjects

Forty eight student volunteers from Durham University served as subjects. Their ages ranged from 18 to 24 years.

[1] Geddes and Stevenson also included a third group of subjects who were given both the specific and the non-specific goals. The results suggested that these subjects engaged in implicit instance learning in which only correct trials were memorized. The specific goal group, by contrast, appeared to engage in both implicit instance learning and explicit means-ends analysis. They, therefore, memorized both correct and incorrect trials.

Twenty four were assigned to the specific goal group and 24 to the non-specific goal group. Within each group, half the subjects were assigned to the description condition and half to the explanation condition.

Design

All subjects were required to complete 30 learning and 30 test trials. The goal groups were defined by the nature of the goal in the 30 learning trials, either specific ('Make Clegg polite') or non-specific ('Find the underlying pattern'). Half the subjects in each goal group performed the descriptions secondary task, the remaining subjects performed the explanations secondary task. In the test trials, all subjects were given a new specific goal ('Make Clegg very friendly'), with no secondary task. After the test trials, all subjects were given two further (unexpected) tests of learning: predicting Clegg's next response from a sequence of three responses and answering questions designed to elicit descriptions rule underlying Clegg's behaviour.

Learning and Test Trials. Subjects were told that they would be meeting a computer person named Clegg and would communicate with Clegg through the screen and keyboard. Clegg would express his attitude towards them by displaying one of twelve descriptions (*Very Rude, Rude, Very Cool, Cool, Indifferent, Polite, Very Polite, Friendly, Very Friendly, Affectionate, Very Affectionate, Loving*). Following this, subjects responded to Clegg by choosing one of the above descriptions. This was done by typing in the first letter or letters of that description (e.g. VP for *Very Polite*). Once subjects had responded, Clegg would display his new attitude (produced by the equation described below). It would then be the subject's turn to enter their next attitude, and so on. The list of possible responses was displayed on a piece of paper attached to the bottom of the screen for permanent reference.

In addition to the above instructions, each group of subjects was given specific instructions concerning their learning goal and their secondary task. Subjects in the Specific Goal Group were told "Your aim is to shift Clegg to the *Polite* level and maintain him at that level" Subjects in the Non-specific Goal Group were told "Your aim is to establish under what pattern Clegg is reacting" To remind subjects of their respective goals, the goal of their task was permanently displayed on a piece of paper attached to the bottom of the screen. Subjects in the description condition were told to describe aloud what they were doing during the learning trials. Subjects in the explanation condition were told to explain why they were doing what they were doing. The rest of the experiment was identical for all subjects.

On each trial Clegg's and the subject's responses were displayed on the screen. These scrolled up the screen so that it was possible to see the previous six trials on the screen at any one time. The equation relating Clegg's responses to those of the subject's was identical to the non-salient rule used by Berry & Broadbent (1984). The descriptions were given a value from 1 (*Very Rude*) to 12 (*Loving*) and Clegg's response was determined by the equation:

$$CNR = (2 \times SOR) - COR + Z,$$

where CNR = Clegg's new response, SOR = subject's old response, COR = Clegg's old response and Z = a random number with the value of -1, 0 or +1. The random element in the equation ensures that subjects must exercise continuous control over the computer person. It also means that there is no unique input associated with any one output. If subjects reached their target output then simply re-entering the same input is unlikely to keep them on target (Berry & Broadbent, 1984). To allow for the random element in the equation producing Clegg's response, the responses of subjects in the specific goal group were scored as correct if they were either on the target or one response either side of the target. That is, a response from Clegg of *Indifferent, Polite,* or *Very Polite* was scored as correct. The test trials were identical to the learning trials for the Specific Goal Group except that the goal was changed. As was the case in the learning trials, a response either on the target or one step either side of the target was scored as correct, to allow for the random element in the equation.

Prediction Questions. There were 15 prediction questions, 5 new, 5 old correct and 5 old wrong. For each question, a typical trial situation was presented. The subject's and Clegg's behaviour was displayed on the screen, below this the subject's new behaviour was displayed - e.g. You were *Very Cool,* Clegg was *Very Rude,* You were then *Polite*. Subjects then had to predict what Clegg's response would be. The five 'new' situations were generated randomly from a list of all possible trial situations that the subject had not encountered during either the learning trials or the testing trials. The five '*Old-wrong'* situations were randomly selected from all the trials the subject had got wrong during the test phase. The five '*Old-correct' situations* were randomly selected from all the trials the subject had got correct during the test phase. To produce five *Old-wrong* and five *Old-correct* questions meant that the subject must get at least five wrong or five correct respectively during the test trials. The program controlling the experiment allowed for the possibility of this not occurring and would have substituted any uncreated questions with *New* questions.

Rule descriptions. Two questions tested the subjects' ability to describe the rule underlying Clegg's behaviour. One was "How did you get Clegg to behave as you wanted him to?" This question was designed to be sensitive to any procedural knowledge that may have been acquired during learning. The other question was "Could you try to describe what sort of pattern you thought Clegg was using to respond to your behaviour?" This question was designed to be sensitive to declarative knowledge.

Procedure

Subjects were randomly allocated to one of the two goal groups. Within each goal group, half the subjects were instructed to describe what they were doing during the learning trials; the other half were instructed to explain what they were doing during the learning trials. The remaining instructions were identical for the two goal groups apart from one sentence. This sentence dictated the goal of that particular group for the learning trials.

On completion of the learning trials, all subjects were instructed on the learning goal for the test trials and then the test trials started. Clegg initiated both learning and test trials by displaying one of the three adjectives centered on *Polite*. Following the test trials, subjects were instructed on the prediction questions. The instructions described the questions and gave an example of a prediction situation. The instructions also explained that each question was unrelated to the previous one. After completing the prediction questions subjects were given a pen and paper and were asked to answer the two general questions appearing on the paper.

Results

Learning Trials

Learning trials were scored as correct for the Specific Goal subjects if they obtained a response from Clegg of *Indifferent, Polite* or *Very Polite*. This scoring takes into account the random element of the equation producing Clegg's behaviour. The mean number of correct learning trials for the specific goal subjects who gave descriptions was 14.67 (49%) and for the specific goal subjects who gave explanations was 13.42 (45%). Thus, learning was comparable in the two secondary task conditions.

Test Trials (Novel Specific Goal)

For both goal groups, correct trials were identified in the same way as in the learning trials. Table one shows the percent correct test trials for group. A 2 (learning goal) by 2 (secondary task) analysis of variance revealed significant main effects of learning goal ($F=4.69$, $df=1,44$, $p<.03$) and secondary task ($F=15.01$, $df=1,44$, $p<.001$). However, these effects were modified by a significant interaction ($F=23.61$, $df=1,44$, $p<.001$). Non-specific goal subjects outperformed the specific goal subjects in the explanations condition only.

Table 1: Percent correct responses on the test trials as a function of secondary task.

	Secondary Task	
	Descriptions	Explanations
Learning Goal		
Specific	49	45
non-specific	38	74

Prediction Questions

Predictions were scored as correct if the response predicted by the subjects was one above, the same as, or one below the response expected from Clegg in each situation. The response expected from Clegg was calculated by using the equation from the learning phase of the experiment, but not including the random element of the equation, since the scoring process took it into account. All subjects produced sufficient correct and incorrect responses in the test trials to have 5 old correct and 5 old wrong prediction questions. The data for old-correct and old-wrong situations were combined in the results and the percent correct responses for old and new situations are shown in Figure one.

A two (goal group) by 2 (secondary task) by 2 (question type) analysis of variance was performed on the data. All three main effects were significant. Non-specific goal subjects performed better than specific goal subjects ($F=13.86$, $df=1,44$, $p<.001$); all subjects performed better in the explanation condition than in the description condition ($F=27.01$, $df=1,44$, $p<.001$); and performance was better on old predictions questions than on new ones ($F=29.74$, $df=1,44$, $p<.001$). There was also a significant interaction between learning goal and secondary task ($F=15.55$, $df=1,44$, $P<.001$): While performance of the two learning goal groups was comparable in the description condition, non-specific goal subjects outperformed the specific goal subjects in the explanation condition. The three way interaction between learning goal, secondary task and question type failed to reach significance ($F=3.30$, $df=1,44$, $p<.08$), However, since we had predicted a differential effect of question type according to learning goal in the explanation condition but not in the description condition, we conducted two separate analyses of the interaction between goal and question type in each of the secondary task conditions. The results showed that, as predicted, the interaction was not significant in the description condition ($F<1$): Both goal groups were better at old than at new prediction questions. However, the interaction was significant in the explanation condition ($F=4.35$, $df=1,22$, $p<.05$): The non-specific goal group performed equally well on both old and new questions, while the specific goal group performed best on the old questions.

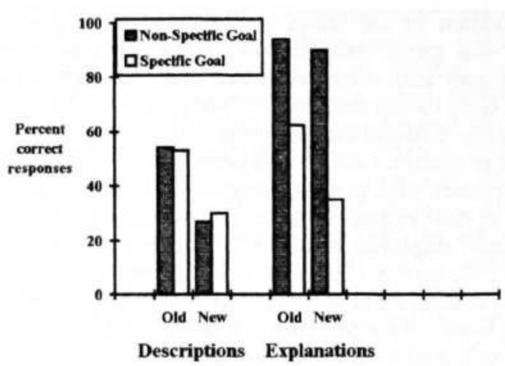

Figure 1: Percent correct predictions test as a function of whether subjects give descriptions (left hand panel) or explanations (right hand panel) during learning.

Rule Descriptions

Subjects' answers to the two questions about the rule (asking how to control Clegg and asking what was Clegg's underlying pattern) were treated together as subjects generally answered only one of the questions and included information in that answer that was relevant to both questions. The answers were judged by two judges and placed into one of three categories; *No information or Wrong*, *Partially Correct*, *Correct*. Answers were categorized as *No information or Wrong* if subjects gave no relevant information about the pattern Clegg was following or about how they controlled Clegg, and if part of the answer gave wrong information. Answers were categorized as *Partially Correct* if subjects mentioned Clegg's tendency to move along the scale beyond the subject's response (away from his own); mentioned any other information that described this approximate characteristic of Clegg's behaviour; made one precise possible prediction of Clegg's behaviour; or mentioned how Clegg's behaviour clustered around a continuous behaviour of the subjects. Answers were categorized as *Correct* when subjects mentioned Clegg's tendency to move along the scale, beyond the subject's response (away from his own) AND described the distance along the scale that Clegg would move (i.e. roughly double the distance the subject was from Clegg). Answers that made 3 or more precise possible predictions of Clegg's behaviour were also classified as *Correct*. The results are shown in Table 2.

As can be seen in Table 2, subjects who described their actions during learning gave mostly wrong answers to the request for rule descriptions. Fisher exact probability tests comparing the number of answers in the *No information or wrong* category and in the *Correct* category showed that wrong answers predominated (specific goal group $p < 0.001$; non-specific goal group $p < 0.001$). However, subjects who explained their actions during learning gave diametrically opposite results depending on their learning goal. All the specific goal subjects gave wrong answers to the questions while all the non-specific goal subjects gave correct answers.

Table two: Percentage of correct, partially correct and wrong rule descriptions as a function of secondary task. (**SG**=Specific Goal; **N-SG**=Non-Specific Goal; **Desc**.=Description condition; **Exp**.= Explanation condition.)

		Correct	Partially Correct	Wrong
SG	Desc.	0	27	75
	Exp.	0	0	100
N-SG	Desc.	17	8	76
	Exp.	100	0	0

Discussion

As predicted, the results showed that explanations facilitated rule learning by the non-specific goal group, while either having a specific goal or giving descriptions fostered instance learning. Such results suggest that learners use a combination of empirical learning and rule learning, since the rule learning we observed was closely tied to the learning instances. In the concept learning literature, Wisniewski and Medin (1995) have proposed a model in which empirical learning and theory driven learning interact. Machine learning researchers have also developed systems that combine both empirical and explanation based learning (e.g. Lebowitz, 1986).

However, our results pose a problem for the concept learning models: how to explain the influence of learning goal or secondary task on the acquisition of instances on the one hand and rules on the other. The strongest evidence for this dissociation between instances and rules comes from the prediction questions. Only the non-specific goal subjects who gave explanations made correct predictions in both old and new situations, consistent with performance based on a rule. The remaining three groups gave more correct prediction in old situations than in new ones, consistent with the retrieval of stored instances. According to Wisniewski and Medin's interactive model, people will learn instances when they have no prior knowledge to inform learning. However, in our study, we can assume that all subjects had roughly the same prior knowledge available to them. Subjects who had a non-specific goal and gave explanations presumably used their prior knowledge of mathematics to help them form, test and refine hypotheses. But this prior knowledge was not used by subjects who had a specific goal or who gave descriptions. The dual space models of Klahr and Dunbar (1988) and Simon and Lea (1974) give the best account of this observation, since in these models, learning can be directed to one or both problem spaces as a function of learning goal and type of verbalization. In the absence of such direction, it is likely that relevant prior knowledge guides the learner to use the hypothesis space as well as the instance space, as was observed Wisniewski and Medin (1995).

The dramatic improvement in the non-specific goal subjects who gave explanations testifies to the powerful effects of explanations on learning (e.g. Chi, et al, 1989; VanLehn, & Jones, 1993). The non-specific goal subjects who gave explanations in the present study learned considerably better than the non-specific goal subjects in the Geddes and Stevenson (in press) study. For example, 100% of the non-specific goal subjects who gave explanations in the present study gave correct rule descriptions, while only 76% of Geddes and Stevenson's subjects gave either complete or partial descriptions. In the educational literature, Ng and Bereiter (1995) have identified three kinds of learners who each spontaneously adopt a different learning goal. Learners with performance goals focus on completing the learning tasks. Such learners can be equated with what Stevenson and Palmer (1994) call 'learning through problem solving'. Learners with instructional goals focus on the manifest learning objectives; they use their background knowledge to help them understand the material but do not use the new material to restructure prior knowledge. This kind of learning can be equated with what Stevenson and Palmer call 'learning through memorization'. Finally, learners with knowledge building goals focus on going beyond the instructional material in pursuit of wider learning goals. Only these learners use the new material to restructure prior knowledge as well as using prior knowledge to understand the new material. This kind of learning can be equated with what Stevenson and Palmer call 'learning through understanding'

While these three kinds of learning are not mutually exclusive, we may speculate that specific goal subjects and subjects who gave descriptions were learning through problem solving; they searched instance space for a route to the goal. We may also speculate that the non-specific goal subjects in Geddes and Stevenson's' study were learning through memorization. They used prior knowledge in conjunction with the initial learning instances to construct a possible hypothesis but may have done little revision of the hypothesis in the light of subsequent learning trials. Finally, the non-specific goal subjects who gave explanations in the present study seem to have been learning through understanding. Giving explanations seems to have encouraged them to modify and refine their hypotheses until the underlying rule was correctly acquired. Our findings, therefore, suggest ways in which learners can be guided to learn more effectively, since we have shown that goal orientation and the use of explanations can be modified to the advantage of the learner.

Acknowledgments

The Human Communication Research Centre is funded by the Economic and Social Research Council of Great Britain. Bruce Geddes is supported by a postgraduate research studentship from the Engineering and Physical Sciences Research Council of Great Britain. We thank two anonymous reviewers for helpful advice.

References

Berry, D. C. & Broadbent, D. E. (1984). On the relationship between task performance and associated verbalisable knowledge. *The Quarterly Journal of Experimental Psychology, 36A*, 209-231.

Broadbent, D. E., Fitzgerald, P., & Broadbent M. H. P. (1986). Implicit and explicit knowledge in the control of complex systems. *British Journal of Psychology, 77*, 33-50.

Chi, M.T.H., de Leeuw, N., Chiu, M-H. & LaVancher, C. (1994) Eliciting self-explanations improves learning. *Cognitive Science, 18*, 439-478.

Cleeremans, A. (1993). *Mechanisms of Implicit Learning: Connectionist Models of Sequence Processing*. MIT Press.

Dienes, Z., & Fahey, R. (1995). The role of specific instances in controlling a dynamic system. *Journal of Experimental Psychology: Learning, Memory & Cognition, 21*, 848-862.

Geddes, B.W. & Stevenson, R.J. (in press). Explicit learning of a dynamic system with a non-salient pattern. *Quarterly Journal of Experimental Psychology*.

Klahr, D., & Dunbar, K. (1988). Dual space search during scientific reasoning. *Cognitive Science, 12, 1-55*.

Lebowitz, M. (1986). Integrated learning: Controlling explanation. *Cognitive Science, 10*, 219-240.

Newell, A., & Simon, H.A. (1972). *Human Problem Solving*. Engelwood Cliffs, NJ: Prentice-Hall.

Ng. E. & Bereiter, C. (1995). Three Levels of Goal Orientation in Learning. In A. Ram & D.B. Leake (Eds.) *Goal-Driven Learning*. Cambridge, Mass., London England, MIT Press.

Owen E., & Sweller, J. (1985). What do students learn while solving mathematics problems? *Journal of Educational Psychology, 77*, 272-284.

Reber, A. S. (1989). Implicit learning and tacit knowledge. *Journal of Experimental Psychology: General, 118*, 219-235.

Shanks, D. R., & St. John, M. F., (1994). Characteristics of dissociable human learning systems. *Behavioral and Brain Sciences*, 17, 367-447.

Simon, H. A. & Lea, G. (1974). Problem solving and rule induction: A unified view. In L.W. Gregg (Ed.) *Knowledge and Cognition* (pp. 105-128). Potomac, Maryland: Lawrence Erlbaum Associates.

Stevenson, R.J. & Palmer, J.A. (1994). Learning: Principles, Processes and Practices. London: Cassell.

VanLehn, K. & Jones, R.M. (1993) Learning by explaining examples to oneself: A computational model. In S. Chipman and A.L. Meyrowitz (Eds.) *Foundations of Knowledge Acquisition*. Kluwer Academic Publishers.

Vollmeyer, R., & Burns, B.D. (1995). Does hypothesis-instruction improve learning? in J.D. Moore & J.F. Lehman (Eds.) *Proceedings of the Seventeenth Annual conference of the Cognitive Science Society* (pp. 771-776). Mahwah, NJ; Hove, UK: LEA.

Vollmeyer, R., Burns, B.D., & Holyoak, K.J. (1996) The impact of goal specificity on strategy use and the acquisition of problem structure. *Cognitive Science, 20*, 75-100.

Whittlesea, B.W.A., & Dorken, M.D. (1993). Incidentally, things in general are particularly determined: An episodic-processing account of implicit learning. *Journal of Experimental Psychology: General, 122*, 227-248.

Wisniewski, E.J. & Medin, D.L. (1995). Harpoons and Long Sticks: The Interaction of Theory and Similarity in Rule Induction. In A. Ram & D.B. Leake (Eds.) *Goal-Driven Learning*. Cambridge, Mass., London England, MIT Press.