Goal Specificity and Learning: Reinterpretation of the Data and Cognitive Theory

Adisack Nhouyvanisvong (adisack@andrew.cmu.edu)
Department of Psychology
Carnegie Mellon University
Pittsburgh, PA 15213

Kenneth R. Koedinger (koedinger@cmu.edu)
Human-Computer Interaction Institute
Carnegie Mellon University
Pittsburgh, PA 15213

Abstract

In this paper, we review the literature on the relation between solving nonspecific goal problems and learning. Research has shown that reduced goal-specificity facilitates learning of rules and principles of the target domain. Researchers have accounted for this effect using a cognitive load theory (Sweller, 1988) and a dual space theory of problem solving (Vollmeyer, Burns, & Holyoak, 1996). Other researchers have shown that learning can be both facilitated by nonspecific as well as specific goals and account for their findings using goal appropriateness theory (Miller, Lehman, & Koedinger, under review). We judge each theoretical account by evaluating their consistencies with unified theories of cognition and other empirical data. We note the shortcomings of the each theory and incorporate elements of each to explain all the data.

Introduction

The specificity of the goal of a problem can be quite varied. The goal can be very specific. For example, in order to get lunch, a person short on cash might set a specific goal to go to the bank on First Street before going to lunch. A student solving a geometry problem might be asked to find the value of a particular angle in a given diagram. On the other hand, the goal can be relatively nonspecific. The person short on cash might set a nonspecific goal to go get cash before lunch. The geometry student might be asked to find the values of *all* angles in the figure.

A number of researchers have explored the relation between goal specificity and learning addressing the question: Do students learn more from solving problems with specific goals or non-specific goals? This question is related to the more general debate about discovery learning: Do students learn more when given specific tasks or when given open-ended discovery tasks? Inspection of the research literature on these questions leads unequivocally to the answer "It depends!". The goal of this paper is to clarify what it depends on, namely, under what circumstances do nonspecific goal problems lead to better learning and when do they lead to poorer learning? The approach is first, to systematically investigate the theoretical conjectures different researchers use to explain their apparently conflicting results and second, to sift these conjectures through the filter of unified theories of cognition, ACT-R (Anderson, 1993) and Soar (Newell, 1990).

Reduced-Goal Specificity Effect

Sweller and his colleagues (Sweller, Mawer, & Ward, 1983; Owen & Sweller, 1985; Sweller, 1988) can be credited with originating investigations of the relationship between goal specificity and learning. In their studies, they have consistently found that students learn more in nonspecific goal conditions than in specific goal conditions. For example, a nonspecific goal geometry problem may ask a problem solver to find all the angles in a figure. The goal of such a problem is nonspecific in that it does not ask for a specific value for a specific angle. In the specific goal condition, study participants are instead asked for a specific angle. Sweller et al. (1983) found that nonspecific goal participants acquired appropriate problem schemas better than specific goal participants.

Sweller, Mawer, & Ward (1983) also found that participants who solved nonspecific goal physics problems exhibited behavior more characteristic of expertise than participants who solved specific goal physics problems. Initially, participants in both groups solved a set of specific goal problems by using means-ends analysis. However, only the participants who solved a set of nonspecific goal problems switched to a forward-working strategy on the last set of specific goal problems. In contrast, participants who solved only specific goal problems continued to use means-ends analysis.

Means-ends analysis is characterized by working backward from the goal state and setting subgoals to reduce the difference between the current state and the goal state. Larkin, McDermott, Simon, and Simon (1980) and Simon and Simon (1978) found that novice physics problem solvers used means-ends analysis to solve the problems. On the other hand, expert problem solvers used a working-forward strategy (see also, Koedinger & Anderson, 1990). The experts were able to recognize and choose the appropriate equations that lead to the goal and apply them immediately.

In addition to the expert-like strategy use, Sweller et al. (1983) found that the nonspecific goal participants wrote significantly fewer equations without variable substitution to solve these problems. They also solved these problems with fewer moves. These behaviors are also characteristic of expertise. In contrast, participants in the specific goal group wrote equations that required variable substitution (i.e., equations with two unknown variables), and performed the same number of moves to solve the final set of problems as they did on the initial problems. These results also indicate that the participants who solved nonspecific goal problems developed expertise.

Similar results have been found with geometry and trigonometry problems. In a series of studies similar in design to the physics studies, Sweller et al. (1983) found that reduced goal specificity on geometry problems also facilitated the switch from a means-ends to a forward-working strategy. More participants in the nonspecific goal condition used a forward-working strategy on the final problems than participants in the specific goal condition. In their experiments testing the participants ability to apply the sine, cosine, and tangent trigonometric ratios, Owen and Sweller (1985) found that participants who went through a nonspecific goal acquisition phase committed fewer errors at post-test. More importantly, they committed fewer fundamental errors, i.e., basic errors indicating the lack of understanding of the meaning of each trigonometric ratio, when to apply it, and how to apply it. Furthermore, on a structurally different trigonometry problem given after the post-test, participants in the nonspecific goal group continued to commit fewer fundamental errors.

Taken together, these results appear to provide a strong case, at least under the conditions studied, that learning is more effective when students practice with nonspecific goal problems rather than specific goal problems. Why might this be so?

Cognitive Load Theory

Sweller and colleagues (Owen & Sweller, 1985; Sweller, 1988) have proposed an explanation for the positive effects of nonspecific goal problems on learning as illustrated by links 1, 2, and 3 in Figure 1. Problem solvers tend to solve nonspecific goal problems using a forward-working strategy (link 1) while those solving specific goal problems tend to use a means-ends strategy as was demonstrated in Sweller et al. (1983).

Further, Sweller (1988, Owen & Sweller, 1985) argues that means-ends analysis, with its emphasis on subgoal storage and difference reduction, is taxing to a limited cognitive processing capacity. A forward-working strategy, in contrast, does not require subgoal storage and thus requires less cognitive load (link 2 in Figure 1). The final step in Sweller's argument (link 3) is that since more cognitive resource is required to use a means-ends strategy, there is less resource that can be allocated to the learning of rules. Or, to state it the other way, a forward-working strategy requires less cognitive resource and thus there is more available to learning.

What evidence is there to support Sweller's argument that means-ends analysis places a heavy load on this resource (link 2)? Sweller presented some indirect evidence consisting of performance characteristics of the strategies used. He found that solution times are longer and mathematical errors more frequent when participants use a means-ends strategy (Owen & Sweller, 1985). Sweller (1988) also provided a theoretical argument in the form of a cognitive model that showed means-ends strategy requiring more cognitive resources than a working-forward strategy. In particular his production system model illustrated that more productions, cycles, and conditions must be matched to implement a means-ends strategy compared to a working-forward strategy.

Sweller (1988) also gave some direct evidence for his cognitive load theory. In an experiment with trigonometry problems, in addition to the primary task of solving the problems, participants performed the secondary task of reproducing the problem structure and correct solution path on preceding problems. The hypothesis was that if means-ends analysis requires more cognitive load, fewer resources would be available to perform the secondary task. Sweller found that participants who solved specific goal problems by using means-ends analysis performed more poorly on the secondary task by committing fewer errors in reproducing the structure and solution path of previous problems. This supports the hypothesis that more excess capacity is available when solving nonspecific goal problems than conventional problems (Sweller, 1988).

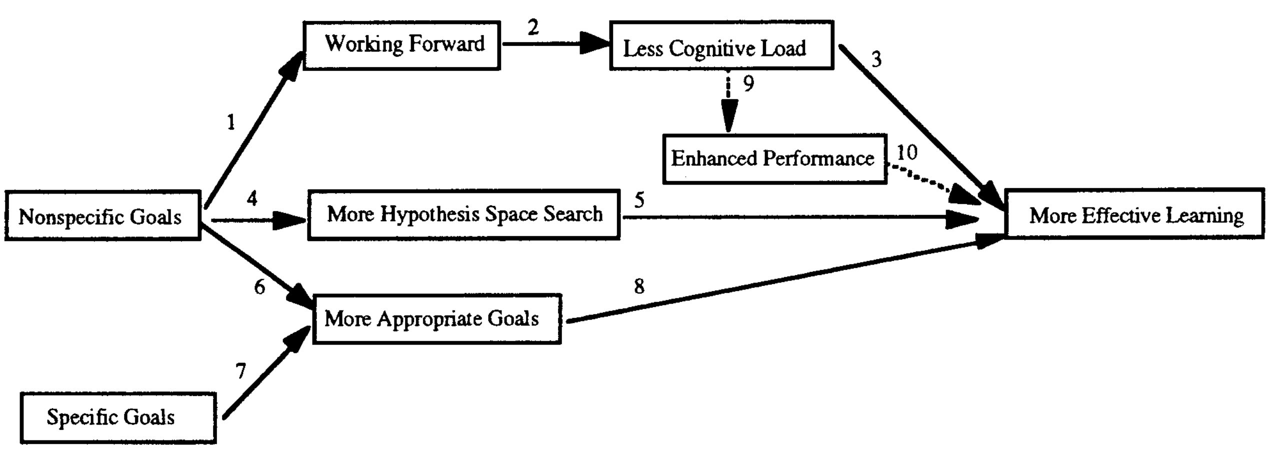

Figure 1: Alternative explanations for the relation between goal specificity and learning.

The proposition that cognitive load has a direct influence on learning (link 3, in Figure 1) is inconsistent with the ACT-R (Anderson, 1993) and Soar (Newell, 1993) theories of cognition. In ACT-R, learning results from successfully analogizing from past examples (Anderson, 1993). Thus, learning occurs when a correct representation of the current situation has been mapped to an analogous past situation. In Soar, learning occurs when new chunks are created to remedy an impasse during problem solving (Newell, 1990; Lehman, Laird, & Rosenbloom, 1996). Under both of these unified theories of cognition, learning mechanisms do not draw on cognitive resources and thus, cognitive load cannot have a direct impact on learning.

An alternative explanation for the relation between cognitive load and learning is that rather than a direct link, there is an indirect link between cognitive load and performance (link 9). Sweller has shown that under conditions of less cognitive load, participants in the worked example condition commit fewer errors during acquisition than participants who solved conventional problems (Sweller & Cooper, 1985). The enhanced performance during acquisition may account for the learning on the post-test (link 10). This explanation then is consistent with both theory and results.

Other researchers have proposed alternative theories for why nonspecific goals can lead to better learning. Further, there are also some experimental results in which the specific goals lead to more learning, suggesting, at best, the need for boundary conditions on the Cognitive Load theory.

Alternative Theoretical Accounts

The tasks and problems used by Sweller were quite simple. The mathematics problems required knowledge of only two or three theorems or principles. No one has tried to replicate these results for a bigger search space in domains similar to Sweller's. However, other researchers have explored goal specificity in different domains.

Vollmeyer, Burns, and Holyoak (1996), using a dynamic problem solving environment, found advantages of nonspecific goal problem solving over specific goal problem solving. In a biology-lab simulation, participants were asked to explore and learn the system. The nonspecific goal group was not given any goals to aim for during the learning rounds. However, the specific goal group was informed of the goal during the learning rounds. Learning or solving the biology-lab system requires relatively complex induction not only of the connections between four input and four output variables, but also the strengths of connections. As they learned the system, the participants completed a structure diagram to indicate how they believed the input variables affected the output variables. In the solution round, both groups were asked to bring the biology-lab to a specified state, which was the same as the specific goal group had during the learning rounds. In the transfer round, both groups were asked to bring the system to a new state not seen before by either group.

Vollmeyer et al. (1996) found that the nonspecific goal group produced better structure diagrams, but both groups performed equally well in the solution round. However, the nonspecific goal group committed fewer errors in the transfer round than those in the specific goal group. This result suggests that acquisition of the structure and rules of the system was fostered by using a nonspecific goal.

Stevenson and Geddes (1997) found analogous results with their dynamic control task. The task required participants to interact with a "computer person" named Clegg. The nonspecific goal group were asked to play with Clegg to find out the pattern that explained his behavior. The specific goal group were asked to learn how Clegg operated by trying to bring him to a specified emotional state (e.g., "make Clegg very friendly").

Stevenson and Geddes (1997) found that participants in the nonspecific goal condition outperformed participants in the specific goal condition. They were better able to predict Clegg's response on later test trials. They also provided better rules describing Clegg's behavior.

How do these researchers account for their results? Though the results are similar to Sweller's results, instead of using a cognitive load account, these researchers account for their data with a dual space theory of learning.

Dual Space Theory

Simon and Lea (1974) defined two spaces where search during problem solving can occur: instance space and rule space. They proposed that a search in instance space is characterized by testing a new state against a goal state. This behavior is an essential feature in problem solving. However, when a problem solver examines the rule space by generating a hypothesis (i.e., a possible "rule", as opposed to a specific instance) and experimenting in the instance space to test the validity of the hypothesis, this behavior is characteristic of learning. A problem solver engages in rule induction when he or she makes connections between the rule space and the instance space. Klahr and Dunbar (1988) have proposed an analogous dual space theory consisting of hypothesis space and experiment space.

Applying the dual space theory described above, Vollmeyer et al. (1996) suggested that nonspecific goal problems accelerate the induction of rules by encouraging problem solvers to search through hypothesis space (see link 4 of Figure 1). On the other hand, specific goal problems promote the search through instance or experiment space alone, which may be an effective method for finding a solution to a problem but is not conducive for learning general rules and principles. Thus, participants in specific goal conditions are less apt to induce rules about the problem structure.

Stevenson and Geddes (1997) similarly interpreted their results with a dual space theory. They suggested that the rule learning that was facilitated by nonspecific goals and the instance learning that was fostered by specific goals can be best accounted for by a dual space model. Specific goals appear to lead to the acquisition of more superficial relations ("without understanding" as educators would say) while nonspecific goal problems lead to the acquisition of deeper domain principles (learning "with understanding"). This misdirected attention prevents the learning of rules.

Goal Appropriateness Theory

Miller, Lehman, and Koedinger (under review) found that free exploration (nonspecific goal problem solving) was not the only condition that had a positive influence on learning. These authors used a microworld called Electric Field Hockey (EFH) to test the specificity of the goal on learning. The pedagogical goal of this interactive microworld was to help students develop a qualitative understanding of the physics of electrical interaction. The immediate goal of the EFH computer game was to maneuver "pucks" into nets analogous to the sport of hockey.

The authors found that their EFH-Soar model of learning in the standard goal condition was able to become skilled at the game. However, it did not learn the underlying physics principles. Miller et al. hypothesized that by changing the task so that participants would have to follow a specific path, they would have to employ underlying physics principles. To test their hypothesis, the authors used three conditions. Students in the nonspecific condition were asked to freely explore the EFH environment (they played with charges and their interaction--no hockey-like goal was presented on the screen). Students in the standard-goal condition were asked to play the EFH game and try to get the puck into the net. Students in the specific-path condition were asked to get the puck into the net by following a specified path outlined for them. The authors posited that the specific-path was more likely to require physics principles and thus be more difficult and demanding of cognitive load.

These authors found that students in the exploratory learning (nonspecific goal) condition learned the physics principles better than students in the conventional standard-goal condition. This was illustrated by their higher scores on a post-test consisting of questions created to assess their acquisition of the underlying physics principles. However, students in the specific-path condition also performed better on the post-test than students in the standard-goal condition. Note, the specific-path condition is more detailed and "specific" than the "standard goal" condition that is most like the condition used by Sweller and other researchers. In addition, students in the specific-path condition progressed more slowly through the levels of the microworld than students in the standard-goal condition. This result suggests that the specific-path condition was more difficult than the standard-goal condition.

Obviously, these results contradict those of Sweller's and cannot be accounted for by Cognitive Load theory. How can an even more specific goal, which presumably is more taxing on cognitive resources than a conventional goal, result in better learning?

Perhaps, it is not the specificity of the goal per se that affects learning. Rather, learning may be influenced by the appropriateness of the goal and subgoals that it elicits to the target activities that are the objectives of instruction. What matters is that the goal is pedagogically related to the task at hand. So, if either goal (specific or nonspecific) elicits pedagogically relevant knowledge, it will lead problem solvers to engage in a search of hypothesis space that will result in effective learning. Thus, contrary to Sweller, Miller et al. (under review) suggest that learning may be achieved either through solving specific or nonspecific goal problems (see Figure 1, links 6, 7, and 8).

Under a Goal Appropriateness account, Miller et al. suggest that students in the nonspecific goal condition learned the principles of physics because they set appropriate goals for themselves (link 6). Students in the specific-path condition had the appropriate goals set for them (link 7). However, students in the standard-goal condition had an inappropriate goal set for them. Their goal was to get the puck into the net. In other words, their goal was to play the game. In contrast, students in the nonspecific and specific goal conditions set or were given specific pedagogical goals. Thus, their goals were appropriate and lead to more effective learning of the principles of physics (link 8).

The findings from Charney, Reder, and Kusbit (1990) provide additional support for a goal-appropriate account. These researchers trained participants on a computer spreadsheet application. Participants either tried to learn the system through a problem-solving tutorial or through free exploration. In the problem solving tutorial (specific goal) condition, participants were given problems with specific goals to solve. Participants in the free exploration (nonspecific goal) condition were allowed to set up their own goals to learn the system.

Charney, Reder, and Kusbit (1990) found that participants learned the spreadsheet application better when they went through a problem-solving tutorial. The tutorial set many specific goals for the students to solve. It tested them on skills and commands relevant to learning the spreadsheet. In other words, it set appropriate goals for them to learn the application. On the other hand, students in the free exploration condition did not learn the application because they did not set appropriate goals to learn it. These students were computer novices and were not aware of the function and utility of a spreadsheet. Thus, as Charney et al. (1990) suggested, the students in the nonspecific goal condition could not set appropriate goals for themselves.

Note the results of Charney, Reder, and Kusbit (1990) did not replicate those of earlier discovery learning studies. For example, Carroll, Mack, Lewis, Grischkowsky, and Robertson (1985) found that participants who explored (called "guided exploration" by the authors) a word processor learned the program better than participants who followed a conventional tutorial. These participants were guided in their exploration of the word processor by function cards which gave them hints (rather than step-by-step specifications given by the manual) to set up goals to perform a targeted function. The participants who followed the self-study manual had difficulty recognizing and adopting the appropriate goals.

In addition, as noted by Charney et al., participants in the Carroll et al. study were not naive to the function of the word processing program. The participants were temporary office workers who were skilled typists and had relevant knowledge of the goals and strategies used to produce business letters. Consequently, these participants were able to set specific, understandable, and appropriate goals on their own (Carroll et al., 1985, pp. 296-297). Though the data appear to be contradictory, the Goal Appropriateness theory accounts for the results from both studies.

Discussion

We have reviewed evidence illustrating the benefits of solving problems with nonspecific goals. Sweller and his colleagues suggest that nonspecific goals lead to more effective learning because they promote a working forward strategy which bears little cognitive load (Sweller, 1988). Other researchers, however, suggest the nonspecific goals lead to more effective learning because they induce a search through hypothesis space (Vollmeyer et al., 1996; Stevenson & Geddes, 1997). A third set of researchers (Miller et al., under review; Charney et al., 1990) suggest that irrespective of the specificity of the goal, learning is most effective when problem goals are appropriate and pedagogically related to the task at hand.

Although these theories appear to account for the results of studies they came from, do they generalize to the results of others' experiments? And do they do so in a way that is consistent with the unified theories, ACT-R and Soar?

Can the cluster of results supporting the dual space theory (Vollmeyer et al., 1996; Stevenson & Geddes, 1997) be accounted for by the Cognitive Load or Goal Appropriateness theories? We focus on the Vollmeyer results to investigate this question.

First, does Cognitive Load theory account for the Vollmeyer results? A straightforward application of Cognitive Load theory to Vollmeyer's task is difficult since it is not clear how the nonspecific goal condition in this case would lead to a working forward strategy (link 1). However, Vollmeyer et al. reported that specific goal subjects, despite being instructed to use a rule-induction strategy (i.e., vary one thing at a time), eventually switched to a difference-reduction strategy (i.e., iteratively tweak multiple input variables to achieve the specified outputs). Thus, according to Cognitive Load theory these subjects should experience greater cognitive load (inverse of link 2) and therefore less effective learning (inverse of link 3).

As discussed above, the direct relation between cognitive load and learning (link 3) is inconsistent with ACT-R and Soar. Furthermore, the data do not even support the modification of Cognitive Load theory (links 9 and 10) since the nonspecific group, under less cognitive load, did not perform better during training.

Does Goal Appropriateness account for the Vollmeyer results? Since the nonspecific goal group was instructed to learn the rules of the system, they operated under an appropriate goal to perform well in the transfer round. They also performed well in the solution round because they had already learned the rules of the system. The specific goal group performed equally well in the solution round because they already had previous attempts at attaining the goal. Thus, their goal was only appropriate for doing well in the solution round.

Can the cluster of results supporting the Goal Appropriateness theory (Miller et al, under review; Charney et al., 1990) be accounted for by the Cognitive Load or Dual Space theories? We focus on the Miller results to investigate this question.

The finding that the specific-path condition from Miller et al., which is more "specific" and requires higher load, facilitated learning is inconsistent with Cognitive Load theory. The Dual Space account also cannot explain this effect. The specific-path condition, much like the specific condition used by Vollmeyer et al. (1996) and Stevenson and Geddes (1997), requires the problem solver to bring the system to a particular state. Accordingly, Dual Space theory predicts this to lead to a search of only instance space and consequently should not facilitate learning.

Unlike Cognitive Load theory, Goal Appropriateness theory suggests that cognitive load should not have a direct effect on learning. What is important is that the goals bring to bear knowledge required to interact successfully with the system. Thus, as Miller et al. suggested, solving goal-based problems will transfer to learning when goal-dependent knowledge and pedagogically relevant material agree. This knowledge-dependent account is consistent with the learning mechanisms of both the Soar (Newell, 1990) and ACT (Anderson, 1993) unified theory of cognition.

Note Cognitive Load theory predicts that the more cognitive resources that are available, the more likely it is for learning to take place. This suggests a linear relation between cognitive load and learning. Accordingly, cognitive load theory suggests that solving very simple problems, which requires low load, will result in learning and that solving very difficult problems, which requires high load, will not likely result in learning.

The relation between cognitive load and learning is more likely U-shaped than linear. One can imagine conditions of reduced cognitive load (e.g., solving very simple problems, having the experimenter solve the problems for the participant) where learning would not likely occur. For example, if students were only given worked out examples to study and were not required to solve any problems, would we expect any learning? Cognitive Load predicts learning to occur due to the low load required. However, Anderson and Singley (1993) showed that calculus students who selected and applied their own operators (high load condition) learned more than students who had a computer select and apply the operators for them (low load condition). On the other end, one can imagine situations of very high cognitive load (e.g., performing multiple tasks simultaneously, solving very difficult problems) where learning should also be inhibited.

Finally, we address whether Sweller's results can be accounted for by Dual Space or Goal Appropriateness theory? In Sweller's studies, participants in both the specific and nonspecific goal groups were asked to find values to mathematics problems. Unlike participants in Vollmeyer et al. (1996) and Stevenson and Geddes (1997), they were not explicitly asked to induce rules to explain the system. Along this line of reasoning, it appears that Sweller's participants only operated in instance space and thus, the superior learning of the nonspecific group can not be accounted for by the Dual Space theory. However, it possible that because nonspecific goal participants did not get caught up in the high load difference-reduction strategy, they had more capacity available to engage in rule induction. And it is this rule induction activity that leads to more effective learning. This account, combining elements of both Cognitive Load theory and Dual Space theory, has the advantage of being consistent with the learning mechanisms of ACT-R and Soar.

Can Goal Appropriateness theory account for Sweller's results? We think it can, but here again, we need to borrow elements of Dual Space theory. Rather than low working memory load in the nonspecific goal condition, freeing students to do rule-induction, perhaps nonspecific goals encourage rule space search whereas specific goals encourage instance space search, distracting participants from the more general learning that results from rule space search. Thus, according to Goal Appropriateness, the attainment of a correct representation of the problem, which is aided an appropriate goal, which can be achieved under low or high cognitive load, is the important prelude to learning.

Is Goal Appropriateness consistent with a dual space account of problem solving? A distinction must be made between goals that are appropriate for learning and goals that are appropriate for problem solving. Goals that are appropriate for inducing rules may lead to a search of the hypothesis space. Goals that are intended for solving a problem may lead to an exploration of instance space.

Note solving a problem with goals to induce rules is analogous to a science model of problem solving. On the other hand, goals intended to find a specific value is analogous to an engineering model of problem solving (Schauble, Klopfer, & Raghavan, 1991). The science model of problem solving, which is characterized by the search for rules and causal relations among variables (i.e., searching hypothesis space), has been shown to lead to more general learning and understanding. In contrast, an engineering model, which is characterized by maximizing or achieving a certain specified target (i.e., searching instance space), does not result in the learning of rules and principles. However, it can lead to good results for just those specific goals targeted. This further supports the conjecture that more appropriate goals facilitate learning.

In summary, it appears that none of the theories can account for all of the results. Neither Cognitive Load or Dual Space can account for the Miller et al. and Charney et al. results. Without incorporating elements of Dual Space, the Goal Appropriateness theory cannot account for either the Vollmeyer or Sweller results. We have provided arguments for how combining elements for Goal Appropriateness and Dual Space can account for all results.

While the Cognitive Load theory alone fairs poorly in accounting for all the results and, at least without modification, is inconsistent with ACT-R and Soar, we are not ready to reject all elements of it. Modifications of the theory to include a U-shaped relation with learning mediated by either performance enhancement or increased hypothesis search appear to strengthen the theory. Furthermore, combining Cognitive Load and Dual Space seems to provide a better explanation for Sweller's results than Dual Space and Goal Appropriateness.

References

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

Anderson, J. R. & Singley, M. K. (1993). The identical elements of theory of transfer. In J. R. Anderson (Ed), *Rules of the Mind.* Hillsdale, NJ: Erlbaum.

Carroll, J., Mack, R., Lewis, C., Grischkowsky, N., & Robertson, S. (1985). Exploring exploring a word processor. *Human-Computer Interaction, 1*, 283-307.

Charney, D. H., Reder, L. M., & Kusbit, G. W. (1990). Goal setting and procedure selection in acquiring computer skills: A Comparison of tutorials, problem solving, and learner exploration. *Cognition and Instruction, 7(4)*, 323-342.

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

Koedinger, K. R., & Anderson, J. R. (1990). Abstract planning and perceptual chunks: Elements of expertise in geometry. *Cognitive Science, 14*, 511-550.

Larkin, J., McDermott, J., Simon, D., & Simon, H. (1980). Models of competence in solving physics problems. *Cognitive Science, 4*, 317-348.

Lehman, J. F., Laird, J. E., & Rosenbloom, P. S. (1996). A gentle introduction to soar, an architecture for human cognition. In *Invitation to Cognitive Science, vol. 4*. Cambridge, MA: MIT Press.

Miller, C., Lehman, J., & Koedinger, K. (under review). *Goals and learning in microworlds.*

Newell, A. (1990). *Unified Theories of Cognition.* Cambridge, MA: Harvard University Press.

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

Schauble, L., Klopfer, L., & Raghavan, K. (1991). Students' transition from an engineering model to a science model of experimentation. *Journal of Research in Science Teaching, 28(9)*, 859-882.

Simon, H. A., & Lea, G. (1974). Problem solving and rule induction: A unified view. In L. W. Gregg (Ed.), *Knowledge and cognition.* Hillsdale, NJ: Erlbaum.

Simon, D., & Simon, H. (1978). Individual differences in solving physics problems. In R. Siegler (Ed.), *Children's Thinking: What develops?* Hillsdale, NJ: Erlbaum.

Stevenson, R., & Geddes, B. (1997). Effects of goal specificity and explanations on instance learning and rule learning. *Proceedings of the Nineteenth Annual conference of the Cognitive Science Society* pp. (709-714). Hillsdale, NJ: Erlbaum.

Sweller, J. (1988). Cognitive load during problem solving: Effects on learning. *Cognitive Science, 12*, 257-285.

Sweller, J., & Cooper, G. A. (1985). The use of worked examples as a substitute for problem solving in learning algebra. *Cognition and Instruction, 2*, 59-89.

Sweller, J., Mawer, R. F., & Ward, M. R., (1983). Development of expertise in mathematical problem solving. *Journal of Experimental Psychology: General, 112*, 463-474.

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