Goal Specificity and the Generality of Schema Acquisition

David L. Trumpower (dtrumpow@unm.edu)
Timothy E. Goldsmith (gold@unm.edu)
Maureen Below (mollybe37@msn.com)
University of New Mexico
Department of Psychology, Logan Hall
Albuquerque, NM 87131 USA

Abstract

Fourteen statistics novices were asked to solve three statistics word problems under standard (SGS) or reduced (RGS) goal specificity. Later, they were asked to solve both structurally identical and structurally different transfer problems, and their structural knowledge of the domain was assessed. Results indicate that participants in the RGS condition performed better on the structurally different transfer problems and had acquired structural knowledge more similar to that of a domain expert. These results extend previous work in showing that the schematic knowledge acquired under reduced goal specificity training is more general than previously realized. The goal specificity effect is discussed in terms of the attentional focus required to solve RGS and SGS problems.

Introduction

Most theorists agree that schemas form the basis for problem solving expertise. Schemas are typically described as knowledge structures that represent generalized concepts, and are comprised of facts and procedures as well as the interrelationships among those facts and procedures. With respect to problem solving, it is generally accepted that schemas allow: (1) problems to be classified according to the general principles required for their solution (Chi, Feltovich, & Glaser, 1981), (2) solution planning (Priest, & Lindsay, 1992), and (3) use of forward-chained solutions (Koedinger, & Anderson, 1990), all of which are hallmarks of expertise. Thus, an important issue for cognitive scientists and educators alike is to understand how schemas are learned.

Cognitive Load Theory (CLT) has been advanced to describe the relationship between problem solving and learning (Sweller, & Levine, 1982; Sweller, 1988). CLT posits that acquisition of schematic knowledge during problem solving is not automatic; rather, it requires a certain amount of cognitive resources. Therefore, if a problem solving task or strategy demands a great deal of cognitive resources then learning will be impaired relative to a task or strategy that carries a low cognitive load.

CLT has been used to explain the finding that reducing the specificity of goals enhances problem solving performance, otherwise know as the goal specificity effect. The goal specificity effect has been shown in maze learning (Sweller, & Levine, 1982), kinematics (Sweller, Mawer, & Ward, 1983), geometry (Ayres, 1993; Sweller, Mawer, & Ward, 1983), trigonometry (Owen, & Sweller, 1985; Sweller, 1988), and several more complex, dynamic tasks (Miller, Lehman, & Koedinger, 1999; Vollmeyer, Burns, & Holyoak, 1996). According to CLT, problems with standard goal specificity (SGS), in which problem solvers are given values for several variables and asked to solve for the value of a specific unknown variable, encourages use of a means-ends strategy. Under a means-ends strategy, problem solvers' attention is focused on reducing the difference between the current problem state and the goal. Moves are guided by the goal state, which requires solvers to keep in memory the goal, any subgoals, and the current problem state. Because this task is cognitively demanding, it detracts from the learning of relations that are relevant for schema acquisition. Reduced goal specificity (RGS) problems, in which problem solvers are asked to solve for the value of as many unknown variables as possible rather than the value of a *specific* unknown variable, eliminate the possibility of a means-ends strategy. Instead, they require a forward-working strategy where moves are generated solely by the current problem state. Because this strategy is less cognitively demanding (see Sweller, 1988), resources are available for learning the relations relevant to schema acquisition, namely, relations between the appropriate operators and problem states.

According to CLT, training with RGS problems is more likely to lead to schema acquisition than training with SGS problems, where schemas are defined as knowledge of problem states and their associated operators. However, this definition of a schema is limited in that it is only applicable to problems with similar structure as those encountered during training (i.e., problems that share, at least some of, the same problem states as the training problems). We will call this the *limited schema view*. Actually, it is difficult to distinguish this view from one that simply postulates the storage of exemplar solutions. If one remembers previous problem solutions, they then have knowledge of problem states and their associated moves/operators…the same information contained in limited schemas. Under this *exemplar view*, the goal specificity effect can be explained by the notion that RGS solutions are easier to remember than SGS solutions (since they require less cognitive load to perform, more resources are available to store them), and they are forward-working. A third alternative is that schemas are acquired under RGS training and that they are more general than previously believed. We will refer to this possibility as the *general schema view*.

Most of the previous studies investigating the goal specificity effect cannot distinguish among these views, because they have predominantly looked at transfer performance on problems that were structural identical to training problems. For example, Sweller, et al. (1983) showed that novices who practiced with RGS kinematics and geometry problems were more likely to work forward on structurally identical test problems than those who practiced with SGS problems. Although consistent with the idea that RGS participants had acquired schemas (either limited or more general), this result is also compatible with the exemplar view. Since novices tend to use means-ends analysis on standard problems, the solutions to SGS practice problems will be backward-chained, whereas since RGS problems eliminate the possibility of using a means-ends strategy, the solutions to RGS practice problems will be forward-chained. Applying these stored exemplar solutions to test problems would result in forward solutions for RGS participants and backward solutions for SGS participants. Schematic knowledge is not required to account for this finding.

If we assume that the greater cognitive load associated with SGS problems interferes with storage of exemplar solutions, then an exemplar view can also account for the findings that SGS training leads to more errors on isomorphic transfer test problems (Owen, & Sweller, 1985), fewer practice problems accurately recalled (Sweller, 1988), and other related findings.

Furthermore, none of the results mentioned above can distinguish between the limited and general schema views, because both limited and general schemas would apply equally well to problems that are structurally the same as the problems from which the schemas were generated. Structurally different transfer problems, though, would help make the distinction. Limited schemas, comprised of relations between previously encountered problem states and associated operators, would not apply to structurally different problems that have different problem states and different solutions. Exemplar solutions of training problems would not apply either. General schemas that are based on abstract principles, though, would apply to structurally different problems, so long as they could be solved with the same general principle. One finding that may favor the general schema view comes from Owen and Sweller (1985). They trained participants to solve trigonometry problems under either SGS or RGS conditions. Training problems gave values for one side and one angle in a right triangle, and participants were asked to solve for either a specific side of an adjacent triangle, or to solve for the values of as many sides as possible, using the trigonometric ratios sine, cosine, and tangent. Later, performance was tested on structurally identical transfer problems, for which RGS participants showed an advantage. They also tested performance on a diagram construction task in which participants were given values for two sides of a right triangle and were asked to draw the triangle, labeling the values for all three sides. Due to the fact that RGS participants fared better on this diagram construction task as well, the authors concluded that mathematical schema acquisition involves learning mathematical principles, where mathematical principles seem more akin to general than limited schemas. Unfortunately, Owen and Sweller (1985) did not control for the number of sides solved for during training. Because the RGS condition tended to solve significantly more sides during training, any differences upon testing could be attributed to amount of practice rather than goal specificity per se. In the present study, we will use transfer problems that are structurally different from training problems, while also controlling for amount of practice.

Another issue with CLT that remains largely untested is the description of the processes used to account for the effect of goal specificity on schema acquisition. Although Sweller (1988) constructed computational models of SGS and RGS problem solving to show that solving SGS problems do indeed require more cognitive resources, it is nonetheless possible that the functional difference between SGS and RSG problems is solely where attention is focused when solving such problems, and does not depend on the amount of resources available to encode problem information during that time. That is, SGS training might produce just as much learning as RGS training, but if attention is focused in the wrong places, then SGS training will result in erroneous learning. In order to examine this idea, we will employ a structural knowledge measure that is measured independent of problem solving performance, and that allows relatively specific questions about the process of schema acquisition to be tested.

Structural knowledge refers to knowledge of the interrelationships among domain concepts, and is well correlated with domain expertise (for a review, see Goldsmith, Johnson, & Acton, 1991). As such, it is likely to be at least a subset of the knowledge contained in schemas. Trumpower (2000) has shown how network representations of structural knowledge can be used to assess schema acquisition. Briefly, the process involves comparing participants' knowledge networks with those of domain experts. Figure 1 displays a knowledge network of the statistics concepts used in the present study, derived from two statistics experts.

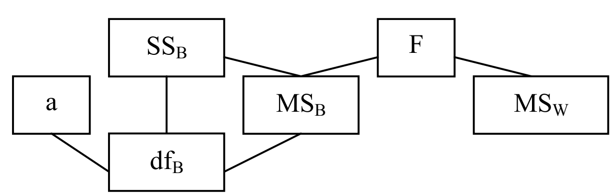

Figure 1: Expert knowledge network

According to an attentional focus explanation of the goal specificity effect, RGS training allows learning of the relationships between problem states (i.e., the subset of variables that are known at a given time) and appropriate operators (i.e., the equation that can be used at that same time to solve for an unknown). Thus, we might predict that RGS training will result in learning of the relations among concepts contained in the equations used to solve RGS problems, since the equations are the operators and contain the currently known variables. By inspecting the three equations needed to solve training problems in the current study (listed below in the Problem Domain & Materials section), we see that relations among concepts in those equations correspond almost perfectly with the pattern of links in the expert network shown above. Therefore, we predict that participants undergoing RGS training will acquire knowledge structures that look very similar to the expert network shown above.

For SGS training, an attentional focus explanation says that attention is directed toward the goal and reducing differences between current states and the goal, at the expense of noticing the local relationships described above. Therefore, we predict that SGS training will result in making associations between all problem states or known variables and the goal (e.g., links between a-SS_B, df_B-SS_B, MS_B-SS_B, MS_W-SS_B, F-SS_B), but a failure to notice the relevant relations between non-goal concepts (e.g., links between a-df_B, df_B -MS_B, F- MS_B, F- MS_W).

To summarize, the current study addresses three questions: (1) Does goal specificity have its effects primarily on storage of exemplar solutions or schema acquisition?, (2) If the effects are on schema acquisition, then how general are the acquired schemas?, and (3) Can the observed effects be better accounted for by the processes proposed in CLT or an attentional focus explanation? In order to examine these questions, we assessed problem solving performance on transfer problems that were structurally different than training problems, and used a measure of schematic knowledge that is independent of problem solving performance.

Method

Participants

Fourteen undergraduate students enrolled in an Introductory Psychology course at the University of New Mexico participated in this study for partial course credit. None of them had previously completed a college-level statistics course. Half of the participants were randomly assigned to receive training with standard goal specificity problems (SGS), while the other half received training with reduced goal specificity problems (RGS).

Problem Solving Domain & Materials

The problem solving domain used in the present study was one-way analysis of variance (ANOVA). All problems used were relatively simple word problems that could be solved with the following three equations: $df_B=a-1$, $MS_B=SS_B/df_B$, and $F=MS_B/MS_W$, where a = number of groups, df_B = between groups degrees of freedom, MS_B = between groups mean square, SS_B = between groups sum of squares, F = F-ratio, and MS_W = within groups mean square.

All training problems gave values for a, MS_W, and F. Those used in the SGS condition asked to solve for SS_B, while those used in the RGS condition asked to solve for as many unknown values as possible. Notice that in both conditions successful solutions required participants to first solve for df_B and MS_B (in either order), and then solve for SS_B.

Structurally identical transfer problems for both conditions were identical in structure to the training problems received in the SGS condition during training in that they gave values for a, MS_W, and F, and asked to solve for SS_B. Structurally different transfer problems were different in structure from the training problems in that they gave values for different variables, and asked to solve for a different variable. Thus, structurally different transfer problems still required use of the same three equations to solve, but they required that the equations be used in a different order and that the equations be manipulated in a different way than was done during training.

A relatedness rating task was also used in which participants were asked to rate the relatedness of all pairwise combinations of the six statistics terms contained in the equations listed above on a 5-point scale (1="Not at all related", 5="Very related").

Procedure

All participants were tested individually in the presence of an experimenter. Participants were first asked to solve three training problems. During this training period, they were given a Rolodex containing separate note cards containing each of the three equations necessary for solution of the problems, as well as a calculator to perform computations. Participants were allowed five minutes to solve each problem. Within this time, the experimenter would immediately notify the participant if they made a mistake, but would not tell them the nature of the mistake. If the problem was not solved within five minutes, the experimenter would guide them to the solution. After solving a problem, participants went on to the next problem and could not refer back to previous problems.

Upon completion of the third training problem, participants were asked to complete the relatedness rating task, which took approximately five minutes. The equations were not made available to participants during completion of this task.

Next, participants were asked to solve four transfer problems (2 structurally identical, 2 structurally different). Approximately half of the participants in each condition were given the two structurally identical transfer problems first, while the other half were given the two structurally different transfer problems first. Participants were again given the necessary equations, and problem solving proceeded as during training.

Results

Separate one-way ANOVAs were used to compare the SGS and RGS conditions on time to solve each of the training problems, and on time to solve structurally identical and structurally different transfer problems. Additionally, separate one-way ANOVAs were used to compare the number of various kinds of links found in the structural knowledge representations of participants in the SGS and RGS conditions. A .05 significance level was used for all tests.

Training

Participants in the RGS condition solved the first two training problems significantly faster than those in the SGS condition, $F(1,12)=5.03$, $p=.045$ and $F(1,12)=6.89$, $p=.022$, respectively for the first and second training problem. This is consistent with the idea that SGS problems require greater cognitive load, and should therefore require more time to solve. There was no significant difference between the SGS and RGS conditions on time to solve the final training problem, $F(1,12)=1.52$, $p>.10$, suggesting that participants in both conditions had acquired similarly efficient solution procedures by the end of training (see Table 1).

Table 1: Time (in seconds) to solve training, structurally identical transfer (S-I), and structurally different transfer (S-D) problems as a function of training condition.

Problem	SGS	RGS
	Mean (SD)	Mean (SD)
First training	300.00 (0.00)	238.14* (72.94)
Second training	219.00 (58.25)	145.00* (46.59)
Third training	139.00 (50.39)	106.29 (48.85)
S-I transfer	108.29 (31.30)	99.64 (55.44)
S-D transfer	254.93 (45.27)	164.93 (78.52)*

*$p<.05$

Structurally Identical Transfer

There was no difference between the SGS and RGS conditions on average time to solve structurally identical transfer problems, $F<1$ (see Table 1). Apparently, both conditions learned to solve problems of the structure that they were trained on equally well. Although CLT (both the limited schema and general schema views) had predicted better performance from the RGS condition, it is possible that the task was too easy to disrupt learning in the SGS condition. If so, then we would expect no difference on the structurally different transfer problems.

Structurally Different Transfer

Participants in the RGS condition solved the structurally different transfer problems significantly faster than those in the SGS condition, $F(1,12)=6.90$, $p=.022$ (see Table 1). This suggests that although both conditions learned to solve problems structured like the training problems equally well, those in the RGS condition gained qualitatively different knowledge that allowed superior transfer to structurally different problems. This is in contrast to both the limited schema and exemplar views. Schemas comprised of knowledge of problem states encountered during training and associated operators would not apply to the structurally different transfer problems, since these problems involved different problem states. Neither would exemplar solutions acquired during training apply, since the structurally different transfer problems required different solutions. Based on these results, it appears that the schematic knowledge acquired during RGS training is more general than previously thought.

Structural Knowledge

Participant's relatedness ratings were submitted to the Pathfinder scaling algorithm to generate a knowledge network for each (for a review of Pathfinder, see Schvaneveldt, 1990). These networks were then analyzed for the number of: (1) critical links with the training goal, (2) irrelevant links with the training goal, and (3) critical links with non-goal concepts (see Table 2).

There are two critical links with the training goal (SS_B) found in the expert network, one with each of the subgoals, (df_B and MS_B). There was no difference in the mean number of these links possessed by participants in the SGS and RGS conditions, $F<1$, as predicted by an attentional focus explanation.

Four other irrelevant links (i.e., those not found in the expert network) with the training goal are possible. As predicted by the attentional focus explanation, participants in the SGS condition possessed significantly more of these irrelevant links than participants in the RGS condition, $F(1,12)=7.59$, $p=.017$.

Four other critical links, not involving the training goal, are present in the expert network. Of these links, participants in the SGS condition possessed significantly fewer than participants in the RGS condition, $F(1,12)=7.36$, $p=.019$, again consistent with predictions made by the attentional focus explanation.

Taken together, these structural knowledge results are consistent with an attentional focus explanation. Under SGS training, attention is focused on the goal, resulting in both relevant and irrelevant associations being made with the goal, at the expense of other critical schematic associations. RGS training, on the other hand, focuses attention precisely where it is needed for schema acquisition, on the local relations described by the equations.

Table 2: Number of links as a function of training condition.

Link type	SGS	RGS
	Mean (SD)	Mean (SD)
Critical, with goal	1.14 (.69)	1.43 (.53)
Irrelevant, with goal	1.57 (.98)	.29 (.98)*
Critical, with non-goals	1.71 (.76)	3.00 (1.00)*

*$p<.05$

Discussion

The results of the current study support and extend previous studies of the goal specificity effect. Reducing the specificity of training goals led to problem solving advantages. However, the advantage was found on transfer problems that were structurally different than training problems. Thus, it is argued that the schematic knowledge that is more readily acquired under RGS than SGS training is more general than previously considered.

These results are consistent with Owen and Sweller's (1985) contention that schema acquisition involves learning abstract principles. It appears that these principles are not tied to problems of a specific form. With respect to the structural knowledge measure employed in the current study, results suggest that the acquired relational information is not unidirectional. Such findings are important for theories of expertise, since we expect expert-like schemas to be applicable to a wide range of novel problems, not just those encountered in the past. If schemas were limited, then experts would gain no advantage at solving novel problems. The very basis for schema theory is that experts possess not only more knowledge through experience, but also better structured knowledge.

The finding that RGS training leads to the acquisition of general knowledge that can be transferred to structurally different problems than encountered during training is pedagogically important as well. Despite being one of the foremost goals of educators, the difficulty in obtaining transfer to non-isomorphic problems has been well documented (e.g., Gick, & Holyoak, 1983).

The structural knowledge results obtained in the current study are consistent with an attentional focus explanation for the goal specificity effect. It should be pointed out that although the attentional processes invoked by this explanation are described in CLT, they are not dependent upon greater cognitive load being present in the SGS condition. Instead, the present results can be explained by SGS training focusing attention on pedagogically irrelevant relations, and RGS focusing attention towards pedagogically relevant ones. This explanation is similar to one advanced by Miller, et al. (1999). They had participants learn about electrical fields by interacting with a microworld called Electronic Field Hockey (EFH). Participants who practiced moving a puck around the EFH workspace in a no-goal condition performed better on a subsequent test of declarative and procedural knowledge of electrical fields than those who practiced by directing the puck around obstacles and into a specific goal. However, participants who practiced by trying to make the puck follow a well-specifed path denoted by a line leading around obstacles and into a goal, performed almost as well as those in the no-goal condition. Miller, et al. (1999) posit that eliminating the goal worked by requiring interaction with the pedagogically relevant aspects of EFH, just like the specific-path condition. In other words, the specific-path condition directs attention away from the ultimate goal toward a series of more immediate subgoals. By directing attention from more distant goals, it can be focused on local relations involved in solving current subgoals. Similarly, eliminating distant goals altogether allows attention to be focused on immediate local relations, which turn out to be the pedagogically relevant ones.

The results of Miller, et al. (1999) may also be explained by a cognitive load interpretation. If following a specific path shifts attention completely away from the ultimate goal, then the task becomes one of meeting a continuous series of smaller goals. If attention is directed at solving each of the immediate goals (i.e., staying on the path), and if each of these small goals can be solved without use of a means-ends strategy, then the specific-path condition would require no more cognitive resources than the no-goal condition. It may be argued that problem solvers solving no-goal problems do adopt a strategy of setting a series of small goals for themselves that can be solved in a forward-chained manner.

Thus, although neither the present results nor the results of Miller, et al. (1999) require an explanation based on cognitive load, they do not rule it out as a possibility. One way to resolve the issue concerning whether RGS training works due to reduced cognitive load or to a pedagogically relevant focus of attention would be through a dual task paradigm. Problem solvers could be asked to solve RGS problems either while concurrently performing another resource demanding task or not. If the concurrent task interferes with learning in a manner consistent with SGS training, then the cognitive load explanation would be justified.

Overall, this study indicates that eliminating specific goals during training can benefit schema acquisition, and that this advantage is more general than previously considered. Training on problems with non-specific goals allowed better transfer to structurally different problems. It is concluded that non-specific goals allow learning pedagogically relevant, local relations, as opposed to standard problems which interfere with such learning. It is suggested that problems with non-specific goals provide this advantage by focusing attention on relations necessary for schema acquisition.

References

Ayres, P.L. (1993). Why goal-free problems can facilitate learning. *Contemporary Educational Psychology, 18,* 376-381.

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.

Gick, M.L. & Holyoak, K.J. (1983). Schema induction and analogical transfer. *Cognitive Psychology, 15,* 1-38.

Goldsmith, T. E., Johnson, P. J., & Acton, W. H. (1991). Assessing structural knowledge. *Journal of Educational Psychology, 83,* 88-96.

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

Miller, C.S., Lehman, J.F., & Koedinger, K.R. (1999). Goals and learning in microworlds. *Cognitive Science, 23*, 305-336.

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

Priest, A.J. & Lindsay, R.O. (1992). New light on novice-expert differences in physics problem solving. *British Journal of Psychology, 83*, 389-405.

Schvaneveldt, R. W. (1990). *Pathfinder associative networks: Studies in knowledge organization*. Norwood, NJ: Ablex.

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

Sweller, J., & Levine, M. (1982). Effects of goal specificity on means-ends analysis and learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 8*, 463-474.

Sweller, J., Mawer, R.F., & Ward, M.R. (1983). Development of expertise in mathematical problem solving. *Journal of Experimental Psychology: General, 112*, 639-661.

Trumpower, D.L. (2000). Schema acquisition and solution strategy in statistics problem solving. *Proceedings of the Twenty-Second Annual Conference of the Cognitive Science Society* (pp. 1061). Maywah, NJ: Lawrence Erlbaum Associates.

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.