

Cognitive Load During Problem Solving: Effects on Learning

JOHN SWELLER

University of New South Wales

Considerable evidence indicates that domain specific knowledge in the form of schemas is the primary factor distinguishing experts from novices in problem-solving skill. Evidence that conventional problem-solving activity is not effective in schema acquisition is also accumulating. It is suggested that a major reason for the ineffectiveness of problem solving as a learning device, is that the cognitive processes required by the two activities overlap insufficiently, and that conventional problem solving in the form of means-ends analysis requires a relatively large amount of cognitive processing capacity which is consequently unavailable for schema acquisition. A computational model and experimental evidence provide support for this contention. Theoretical and practical implications are discussed.

Problem-solving skill is highly valued. For most of this century, many theorists and educational institutions have placed a heavy emphasis on this ability, especially in mathematics and science (see Dewey, 1910, 1916). Entire movements such as "discovery learning" (e.g., Bruner, 1961) were spawned, at least in part, by the perceived importance of fostering problem-solving skills. This emphasis on problem solving was not associated with a commensurate knowledge of its characteristics and consequences. In the last few years, this state of affairs has begun to change with our knowledge of relevant mechanisms increasing markedly. These mechanisms have implications for learning, as well as problem solving. The purpose of the present paper is to suggest that contrary to current practice and many cognitive theories, some forms of problem solving interfere with learning.

This research was supported by a grant from the Australian Research Grants Scheme. The computational model was constructed while the author was on leave at the Learning Research and Development Center, University of Pittsburgh. I wish to thank J. Greeno and H. Simon for discussing aspects of the model with me. I also wish to thank E. Rees for assistance with the PRISM language, and T. Cahn for research assistance. The cooperation the New South Wales Department of Education and of D. Brown, Principal, F. Navin, Deputy Principal, and also J. Kopp, Mathematics Master, South Sydney High School is gratefully acknowledged.

Correspondence and requests for reprints should be sent to J. Sweller, School of Education, University of New South Wales, P.O. Box 1, Kensington, NSW 2033, Australia.

These results can be closely integrated with those on memory of problem states. Experts are able to work forward immediately by choosing appropriate equations leading to the goal because they recognize each problem and each problem state from previous experience and know which moves are appropriate. The same cognitive structures which allow experts to accurately recall the configuration of a given problem state also allow immediate moves toward the goal from the givens. These cognitive structures will be called schemas where a schema is defined as a structure which allows problem solvers to recognize a problem state as belonging to a particular category of problem states that normally require particular moves. This means, in effect, that the problem solver knows that certain problem states can be grouped, at least in part, by their similarity and the similarity of the moves that can be made from those states. Novices, not possessing appropriate schemas, are not able to recognize and memorize problem configurations and are forced to use general problem-solving strategies such as means-ends analysis when faced with a problem.

Features Used in Categorizing Problems

Hinsley, Hayes, and Simon (1977) found that competent problem solvers could readily categorize algebra word problems with a high degree of inter-subject agreement. In related research, Chi, Glaser, and Rees (1982) found that expert physicists, when asked to categorize a series of physics problems, tended to group them on the basis of solution mode. Problems soluble by a basic principle of physics such as conservation of energy tended to be placed in the same category. In contrast, novices preferred to group problems according to surface structures such as the inclusion of shared objects in the problem statement. For example, problems mentioning an inclined plane tended to be placed in the same category.

The same basic principles leading to expert-novice distinctions in problem-solving strategies and memory of realistic problem-state configurations may also be instrumental in determining modes of problem categorization. Experts, possessing schemas allowing them to distinguish between problem states and their associated moves, may categorize problems according to those schemas. If an expert has a schema which suggests that conservation of energy should be used to solve a particular problem, then that problem is likely to be categorized with other problems to which the same schema can apply. Novices, not having sophisticated schemas of this type must resort to surface structures when classifying problems.

In summary, the expert-novice research suggests that domain specific knowledge, in the form of schemas, is a major factor distinguishing experts from novices in problem-solving skill. Differences in memory of problem states, strategies used and categories into which problems are placed can all be explained by assuming that experts have acquired schemas which play a crucial role in the way in which they approach and solve problems.

LEARNING

In the previous section it was suggested that schema acquisition constitutes a primary factor determining problem solving skill. The manner in which this skill is best acquired is an important question with both theoretical and practical ramifications. Surprisingly, little research has been carried out on this issue. It is commonly assumed by both theoreticians and those concerned with practical problem-solving issues that practice on a large number of conventional problems is the best way of gaining problem-solving skill. Given the domain-specific, knowledge-based nature of problem-solving skill discussed in the last section, there is reason to doubt this assumption. There are also theoretical reasons which will be discussed in the next section for supposing that conventional problem-solving is an inefficient way of acquiring schemas. In the last few years, my collaborators and I have obtained experimental evidence supporting the same conclusion. These results are summarized below.

Evidence of Interference Between Conventional Problem Solving and Schema Acquisition

The initial findings were obtained using puzzle problems. Mawer and Sweller (1982), Sweller (1983), and Sweller, Mawer, and Howe (1982) presented subjects with a variety of puzzle problems which could be solved either by means-ends analysis or by inducing a rule based on the problem structure. Over many experiments it was found that while subjects had little difficulty solving these problems, they tended not to induce the relevant rules. This aspect of the problem structure could only be readily induced if considerable additional information was implicitly or explicitly provided. It was concluded that conventional, goal-directed search heuristics such as means-ends analysis, while facilitating problem solution, could frequently prevent problem solvers from learning essential aspects of a problem's structure. Evidence supporting this conclusion has been obtained by Lewis and Anderson (1985).

Sweller and Levine (1982) obtained direct evidence for this suggestion using maze problems. For some subjects the position of the goal was known. A conventional means-ends strategy of attempting to reduce differences between a given problem state and the goal state could be employed. Other subjects were not shown the goal position. They had to find both the goal and the route to the goal. Under these circumstances, it is not possible to reduce differences between a given problem state and the goal state because the goal state is not known until it is attained. Subjects given a conventional goal, in most cases failed to induce essential structural features of the problems which under some circumstances prevented them from even solving relatively simple problems. In contrast, the use of nonspecific goals permitted rapid learning of essential structural characteristics. These results

provided further evidence of the negative effects of means-ends analysis on learning.

PROBLEM-SOLVING SEARCH VIA MEANS-ENDS ANALYSIS AND SCHEMA ACQUISITION: CONTRARY GOALS?

Why should some forms of problem-solving search such as means-ends analysis interfere with learning? There are two related mechanisms which may be particularly important when considering learning and problem solving: selective attention and limited cognitive processing capacity.

Selective Attention

Solving a problem and acquiring schemas may require largely unrelated cognitive processes. In order to solve a problem by means-ends analysis, a problem solver must attend to differences between a current problem state and the goal state. Previously used problem-solving operators and the relations between problem states and operators can be totally ignored by problem solvers using this strategy under most conditions. Previous states and operators need to be noted only to prevent retracing steps during solution.

We may contrast these mechanisms with those required by schema acquisition. In order to acquire a schema, a problem solver must learn to recognize a problem state as belonging to a particular category of problem states that require particular moves. As a consequence, we might expect attention to problem states previously arrived at and the moves associated with those states to be important components of schema acquisition.

Cognitive Processing Capacity

The cognitive load imposed on a person using a complex problem solving strategy such as means-ends analysis may be an even more important factor in interfering with learning during problem solving. Under most circumstances, means-ends analysis will result in fewer dead-ends being reached than any other general strategy which does not rely on prior domain-specific knowledge for its operation. One price paid for this efficiency may be a heavy use of limited cognitive-processing capacity. In order to use the strategy, a problem solver must simultaneously consider the current problem state, the goal state, the relation between the current problem state and the goal state, the relations between problem-solving operators and lastly, if subgoals have been used, a goal stack must be maintained. The cognitive-processing capacity needed to handle this information may be of such a magnitude as to leave little for schema acquisition, even if the problem is solved.

While selective attention and limited cognitive processing capacity mechanisms have been treated independently in the previous discussions, they are

related. Indeed, for practical purposes, under some conditions it may not be useful to distinguish between the two processes. Assume a problem solver whose entire cognitive processing capacity is devoted to goal attainment. It was suggested in the previous section that this leaves no capacity to be devoted to schema acquisition. Rather than using cognitive processing capacity terms, we could just as easily describe these circumstances in attentional terms. A problem solver whose cognitive processing capacity is entirely devoted to goal attainment is attending to this aspect of the problem to the exclusion of those features of the problem necessary for schema acquisition.

CATEGORIES OF FORWARD-WORKING STRATEGIES

It was suggested previously that backward working could provide an indicator of a means-end strategy. In contrast, a problem solver may work forward when using any one of several distinct strategies. First, the schema-driven approach used by experts will proceed forward from the givens because a schema encodes a series of problem states and their associated moves. All states encountered have schemas associated with them indicating appropriate forward moves. Second, forward working may occur during means-end analysis either because the problem solver chooses to attempt to reduce differences between a current problem state and the goal state by working from the current problem state, or because the problem does not contain a sufficiently well specified goal to allow backward working. As well as these two previously discussed strategies, there is a third forward-working strategy. Problem solvers may work forward without being controlled either by a schema or by a problem goal. They may simply explore the problem space in order to see what moves are possible.

The practicality of the third strategy is dependent on the problem structure. Many problems have extensive state spaces. An uncontrolled search of a space containing thousands of paths is clearly unlikely to be productive. Mathematics or mathematically based problems presented to students tend not to be of this type. Most contain very limited state spaces which can be explored fully in a few minutes by anyone familiar with the problem-solving operators. A kinematics or geometry problem is likely to contain no more than about a dozen (and probably far fewer) unknowns and even fewer equations or theorems which can serve as problem solving operators. Because mathematics problems are of primary interest in this paper, it is appropriate to place a heavy emphasis on a nonspecific goal, schema free approach both as an experimental tool and, more importantly, as a tool able to assist in theory building.

We might expect a nonspecific goal strategy to substantially decrease cognitive load. Using this strategy, a problem solver merely has to find an equation allowing the calculation of any unknown rather than assess differ-

ences between a current problem state and the goal. This, in turn, should allow schema acquisition to occur more readily. Subsequent sections provide both experimental and theoretical (via a computational model) evidence for these contentions.

CONSEQUENCES OF A NONSPECIFIC GOAL STRATEGY ON MATHEMATICAL PROBLEM SOLVING— EXPERIMENTAL EVIDENCE

Sweller, Mawer, and Ward (1983) presented problem solvers with simple physics or geometry problems which had been modified in order to eliminate the conventional, specific goal. This was done by replacing a conventional goal such as "What is the racing car's acceleration" by the statement "Calculate the value of as many variables as you can." It was suggested that this may be analogous to Sweller and Levine's (1982) replacement of a specific by a nonspecific goal using maze problems. In the case of physics and geometry problems, the same theoretical rationale can be used to hypothesize that reducing goal specificity will enhance schema acquisition: A nonspecific goal eliminates the possibility of using a means-ends strategy to solve the problems. The results of these experiments indicated that the development of problem-solving expertise was enhanced more rapidly using a reduced goal specificity procedure. While these experiments provided evidence that means-ends analysis interferes with learning, they do not indicate the mechanisms by which it might do so. The model described in the next section provides support for the suggestion that cognitive processing load is an important factor reducing learning during means-ends analysis.

RELATIVE COGNITIVE LOAD IMPOSED BY MEANS-ENDS ANALYSIS AND FORWARD WORKING— A COMPUTATIONAL MODEL

A direct measure of the cognitive load imposed by a particular strategy or procedure is not available currently. Any potential measure must be capable of simultaneously accounting for problem difficulty, subject knowledge, and strategy used. Computational models do this naturally. Programs which model problem solving using means-ends or alternatively, nonspecific goal strategies can be analyzed in order to obtain an indicator of the relative information-processing capacity required by the two strategies. This section describes such a procedure.

Separate forward- and backward-oriented models have been described by Larkin, McDermott, Simon, and Simon (1980). The backward-working system solves physics problems using a means-ends strategy. This model is

designed to model novice behavior. The other works forward and was designed to model expert behavior. The model to be described here is a single system designed to work backward or forward depending on whether a specific goal is included in the problem statement. Because it is designed to model backward or forward strategies, it bears some conceptual similarities to Larkin et al.'s models. One difference is due to the fact that it is designed only to model novice behavior when confronted with either goal-specific or goal nonspecific problems rather than novice-expert differences. For this reason it consists of a single program which must determine the type of problem faced.

The major differences between the current and previous models are due to the highly specific function of the current model. The primary purpose for constructing the program was to provide evidence for the suggestion that means-ends analysis imposes a greater cognitive load than a nonspecific goal procedure. Consequently, care had to be taken to ensure that the model's mechanisms (especially the means-ends mechanisms) were basic, with no possibly unnecessary features that might increase cognitive load. In this sense, the model is "minimal." It contains the minimal essentials of means-ends and nonspecific goal strategies. This facilitates a comparison of the capacity required by both strategies that reduces the risk that either strategy has been burdened by unnecessary factors. Thus, in contrast to previous work, the model is not intended as a detailed description of problem solving behavior. Such details could distort its function.

The model was constructed using PRISM, a production system language designed to model cognitive processes (see Langley & Neches, 1981). A production system is a set of inference rules that have conditions for applications and actions to be taken if the conditions are satisfied. The model permitted the conditions of a single set of productions to be matched with the elements of a single-working memory. Decisions concerning the order in which productions fired were determined by three factors. The initial decision was based on the relative strengths of the relevant productions. Each production was allocated an initial strength based on psychological assumptions. This could remain constant or alter during a run. If there were several productions with an equal strength which exceeded that of the remaining productions, an additional decision rule was used to break the tie. The "activation" of the statements in working memory that can match the condition side of each production was used for this purpose. The activation of statements can be considered analogous to the extent that those statements are known or familiar. For relevant productions, this activation was summed and the production with the highest activation fired. If, after this procedure, there was still a tie between productions, one of them was chosen at random to fire. In addition, a production which fired on one cycle could not match the same elements and fire again.

Measuring Cognitive Load Using a Production System

While a production system is not specifically designed to measure cognitive load, there are several aspects of production systems which could provide suitable measures. First, we might expect cognitive load to be correlated with the number of statements in working memory. We know that human short-term memory is severely limited and any problem that requires a large number of items to be stored in short-term memory may contribute to an excessive cognitive load. In so far as short-term memory corresponds to a production system's working memory, it is reasonable to suppose that an increased number of statements in working memory increases cognitive load.

The number of productions and the number of conditions that need to match statements in working memory should also provide measures of cognitive load. In order to make progress on a transformation problem by choosing a move, a production system must determine which of its various productions should fire, using the mechanisms described previously. The first and critical step, is to find those productions which match elements in working memory. It seems plausible to suggest that the more productions that need to be considered at each step in the problem and the more statements that need to be matched in order to decide between productions, the greater is the "cognitive load."

The analogy between a production system determining which production should fire and a person deciding what to do next, may be quite close. In both cases, the elements of the problem and the knowledge brought to bear on the problem must be coordinated. A production system with many productions each containing many statements, may be analogous to a person using a complex problem-solving strategy involving many choice points with each choice requiring a large amount of information.

It also should be noted that the general argument applies irrespective of the specific assumptions concerning the system's architecture. For example, a parallel system should require more routes or channels (communication bandwidth) to handle complex rather than simple search mechanisms. These routes or channels should no longer be available for learning. In other words, while the precise measurement will depend on the architecture of the system, the general principle that a more complex search mechanism will require increased capacity which may interfere with learning, should not.

PRODUCTION SYSTEM DETAILS

This section provides details of a production system designed to allow estimates of the relative cognitive load imposed while solving conventional problems by means-ends analysis, compared to nonspecific goal problems otherwise identical in structure. The program is essentially an equation

chaining system. As such, it can be used to solve problems which in essence, require the construction of a chain of previously given equations connecting the givens to the goal.

Geometry, trigonometry, and kinematics problems provide examples of this category. To solve a kinematics problem for example, a chain (or chains) consisting of equations of motion must be constructed. One end of the chain must contain the givens and allow the calculation of values for new variables. These new variables can be used in other equations which become intermediate links in the chain. The end of the chain must contain the goal which can be calculated using the previous links. With limitations to be discussed below, the system to be described will solve all problems of this type as well as structurally identical problems with nonspecified goals.

Working Memory

The system commences with the following information in working memory: (a) A list of the equations that might potentially be used in attempting to solve the problem; (b) A list indicating which of the variables found in any of the equations are known; (c) A similar list indicating which of the variables are unknown. A statement indicating the goal variable was included in the case of conventional problems with a specific goal.

These lists are a combination of the relevant information that subjects must extract from a problem statement and the problem-solving operators (equations) needed to solve the problem. We assume that problem solvers have this information immediately prior to attempting their first move. This only leaves the problem solving strategies that might be used to attain solution. The list of productions specifies these and by eliminating all else from this list, we have a clear delineation between the cognitive processes brought into play before and after the first move. Working memory can be used to measure cognitive processing capacity required before the first move while the production list can be used similarly for processing that occurs during and after the first move.

Description and Justification of Means-Ends Productions

Since it has been hypothesized that a means-ends strategy imposes a heavy cognitive load, the productions required to describe the strategy must be minimal in number in order to avoid artificially increasing this load. Table 1 contains a list of four such productions which are considered essential to a means-ends strategy. The elimination of any one of these productions will prevent equation chaining problems from being solved by a means-ends strategy. Furthermore, they are sufficient to allow solution of all equation chaining problems which do not require the processing of more than one equation at a time.

Problems requiring the use of more sophisticated algebraic procedures such as simultaneous equations can not be solved by these productions. It

TABLE 1
Set of Means-Ends and NonSpecific Goal Equation-Chaining Productions

Means-ends Productions	Conditions	Actions	Strength
1	If a problem has a specified goal and if an equation is known in which the goal is the only unknown	then the goal becomes known	1.2
2	If a problem has a specified goal and if an equation is known which contains the goal and one or more unknowns not previously set as subgoals	then the unknowns not previously set as subgoals and other than the goal become subgoals	1.0
3	If an equation is known which contains subgoals and one or more unknowns not previously set as subgoals	then the unknowns not previously set as subgoals become subgoals	1.0
4	If an equation is known in which a subgoal is the only unknown	then the subgoal becomes a known	1.1
Nonspecific goal Production			
5	If a problem does not have a specified goal and if an equation can be found with only one unknown	then the unknown becomes a known	1.0

should nevertheless be noted in this context, that a new set of equations can always be derived from the sets of simultaneous equations and this will allow any problem to be solved without the use of the simultaneous equations. In this sense, the four productions are sufficient to allow the solution of all equation chaining problems. Any failures can be rectified by deriving the relevant set of equations. For example, if a body is uniformly accelerated from rest, and if acceleration and distance travelled are known, then time travelled can be calculated using the equations; $s = vt$, $v = .5V$ and $V = at$. To solve this problem, simultaneous equations can be used to derive the equation $s = .5at^2$, assuming this is not known. The problem can then be solved in a single move using this equation. (Of course, the equation $s = .5at^2$ is normally taught. If all possible equations in a system are learned, then all possible problems can be solved in one move.)

The condition side of the first means-ends production in Table 1 tests whether the problem has a specific goal and whether an equation is known in which that goal is the only unknown. If these conditions are met, the action side of the production states that the goal is now known. The relative

strength of this production ensures that it will fire prior to any alternative productions whose conditions are also met. (The relative strengths of productions reflect ordinal relations only.) Nevertheless, because its conditions are restrictive—under most problem conditions there are more equations with several unknowns than equations in which the goal is the only unknown—this production is always the last to fire. Before it fires, the values of other unknowns normally need to be found using other productions (unless it is the only production to fire). Once the first production fires, the system comes to a halt. Without this production, a value for the goal cannot be found.

If the first production cannot fire, then the second production must fire. It does not require the goal to be the only unknown in an equation. If any equation is known containing the goal, then this production will fire. The action side specifies that all unknowns other than the goal will become subgoals. These subgoals are thus added to working memory. Despite its relatively low strength, normally, this production is the first to fire because in a soluble problem, its conditions can always be met immediately. Basically, it only requires an equation containing the goal variable. It might also be noted, that if this production fires repeatedly (which may not necessarily occur), the system is engaging in a breadth first search. It is attempting to find as many questions as possible containing the goal. This production is essential for the initial setting of subgoals.

Once subgoals have been added to working memory by the actions of the second production, the third and fourth productions can fire. The conditions of these include subgoals rather than goals and in the case of the third production, include statements preventing previous subgoals from being reset as subgoals.

The third production allows a search in depth rather than the search in breadth of the second production. If used repeatedly, this production can construct a chain of subgoals with a chain of equations. Nevertheless, it is not constrained to a search in depth. It can also conduct a search in breadth, finding as many equations as possible containing a particular subgoal and adding all unknowns as additional subgoals in working memory. Without this production, a chain of subgoals linking equations from the goal to the givens cannot be constructed.

The fourth production has a higher strength than the third but under most problem conditions will fire later because of its more restrictive conditions. In order to fire, this production must find an equation containing a subgoal as the only unknown. This contrasts with the third production which can fire after finding an equation with unknowns other than the subgoal. The fourth production can fire only after the second has fired because prior to this, working memory contains no subgoals. This is a working-forward production. Without it, the system could not calculate values which allow chaining from the givens to the goal.

Example of Means-Ends Operation

These four productions appear to be sufficient to solve all equation chaining problems provided suitable equations are known. Figure 1 indicates the flow of control. A kinematics problem discussed in detail will be used to provide a simple example of the system's operation. The problem states: A car that starts from rest and accelerates uniformly at 2 meters/s/s in a straight line has an average velocity of 17 meters/s. How far has it travelled? It can be solved using the three equations, $s = vt$, $v = .5V$ and $V = at$. As noted above, these equations can be used to solve all problems in which an object is uniformly accelerated from rest in a straight line.

In order to solve this problem, the system must have the three equations in working memory, a list of the variables in the equations which are known (a , v), a similar list of the unknowns (t , V) and a statement indicating the goal variable (s). This is assumed to be equivalent to a person who has read

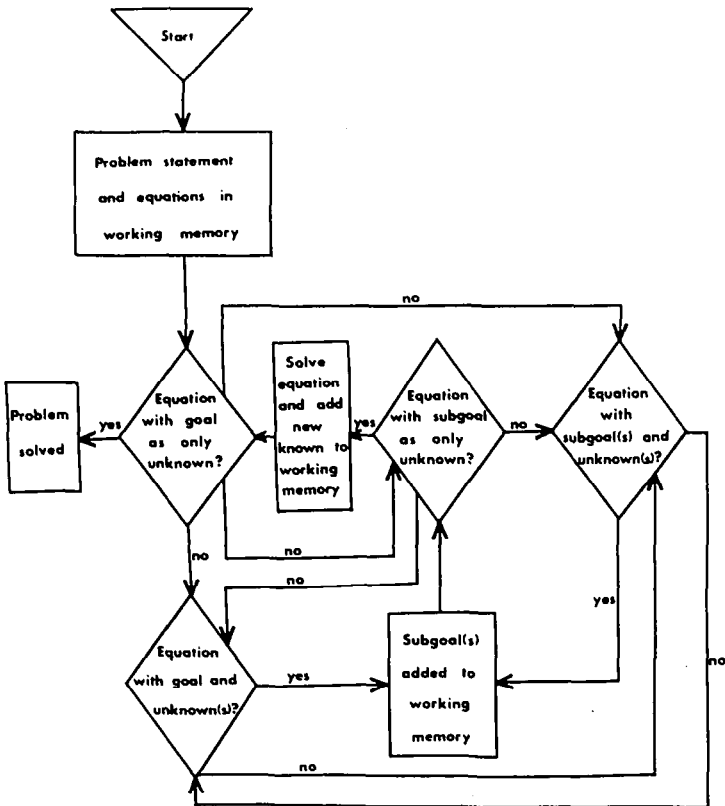


Figure 1. Flow of Control Under Means-Ends Production. (System halts either when the problem is solved or when no production can fire because no new subgoals can be generated and no equation can be solved.)

the problem statement, assumed a set of equations which might be relevant to solution, and determined with respect to each variable whether it is known, unknown or the goal. Translation and other processes leading to this representation of the problem are not reflected in the system. Only the end result is fed into working memory because only subsequent problem-solving processes are of concern when comparing conventional and goal-free problems. This is because translation processes are assumed to be identical in both cases and do not need to be compared.

Working Backward

The system will first attempt to solve this problem in a single step using the first production because this production has the highest strength. It will attempt to find an equation containing a , v , and s . Since such an equation is not available, the first production cannot fire. The fourth production has the next highest strength but this cannot fire because there are no subgoals in working memory. For the same reason, the third production cannot fire either. The second production can fire. An equation ($s = vt$) is known which contains the goal and unknown(s). Time (t) becomes a subgoal. The third and fourth productions now can be considered since subgoals are included in their conditions. The fourth production has a higher strength but cannot fire because no equation can be found in which t is the only unknown. The third production can fire using $V = at$. Final velocity is now a subgoal.

Working Forward

To this point, the system has been working backward from the goal, setting up subgoals. With final velocity as a subgoal, it can solve an equation and attempt to work forward. The conditions of the second, third, and fourth productions are all met at this juncture but the fourth production has the greatest strength. The equation $v = .5V$ contains the subgoal V and the known v . The action side of the production will add V as a known and delete it as a subgoal. The same production will now fire again, matched to different elements. The equation $V = at$ contains the subgoal t and the knowns a and V . Subgoal t is converted into a known. The conditions of the first production, which is the strongest, are now met. The equation $s = vt$ contains the goal s with v and t as knowns. Once this production fires, a value for s is obtained.

A Production to Solve Nonspecific Goal Problems

A single production is sufficient to solve problems in which the goal is not specified. The fifth production of Table 1 provides a description. It is designed to search for equations containing a single unknown and solve for that unknown with no reference to a goal. By firing repeatedly, all unknowns that can be derived from the givens of a problem statement will be found. Figure 2 diagrams the flow of control.




Figure 2. Flow of Control Under a NonSpecific Goal Production

Assume the last sentence of the previous kinematics problem is replaced by the statement "Calculate the value of as many variables as you can." The problem is now represented in working memory by the same equations and same knowns as the conventional problem. The only difference is that the previous goal (distance) is now listed as an unknown.

Because no goal or subgoal is listed in working memory, none of the means-ends productions can fire. Each time the nonspecific goal production fires, it essentially duplicates the forward-working actions of the means-ends productions but does so in a nondirected fashion. While the means-ends productions continually attempt to work forward due to the strength of the forward working productions (1 and 4), they are not able to do so until a suitable set of goals and subgoals have been added to working memory. The nonspecific goal production works forward automatically. On our kinematics example, it will fire three times, successively finding V using $v = .5V$, t using $V = at$ and s using $s = vt$. The same variables were found by the means-ends productions.

MEASURES OF COGNITIVE LOAD

Cognitive load can be measured in several ways. We will consider: (1) the number of statements in working memory; (2) the number of productions; (3) the number of cycles to solution; (4) the total number of conditions matched. Table 2 allows a comparison of these measures when the system solves a conventional and nonspecific goal version of the previously discussed kinematics problem.

Both the means-ends and the analogous nonspecific goal problems commence with an equal number of statements in working memory—14. The only difference is that the conventional problems contain a statement specifying a goal while the non-specific goal problems have the goal variable listed as another unknown. The equal number of statements reflects the assumption that the translation and general representation processes are

TABLE 2
Simulation Data Under Conventional and NonSpecific Goal Conditions When Solving a 3-Step Equation Chaining Problem

	Conventional	NonSpecific Goal
Average working memory	15.5	14
Peak working memory	16	14
Number of active productions	4	1
Number of cycles	5	3
Total number of conditions matched	29	17

Note. See text for an example problem.

similar in terms of cognitive load for conventional and nonspecific goal problems.

A means-ends strategy requires the addition of subgoals to working memory during solution. A nonspecific goal strategy has no net additions resulting in identical average and peak working memory loads. In so far as working memory corresponds to human short-term memory, this can be important due to the severe limitations of short-term memory. A small additional load could be critical. It must nevertheless be pointed out, that most problem solvers faced with the need to remember subgoals, are likely to use an external memory source such as pencil and paper. Differences in working memory may not be critical under these circumstances.

There are more substantial differences on all other measures of cognitive load. It might be noted that the ratio of cycles to conditions matched is approximately equal for both the conventional and nonspecific goal problems. This is reflected by the approximately equal number (7-9) of statements in each production.

The organization of the current system clearly requires more active productions, cycles, and conditions matched for a means-ends strategy than a nonspecific goal strategy. The major difference is in the number of productions and it is this difference that requires emphasis. If these productions reflect human cognitive processes, then they provide strong evidence for the contention that a means-ends strategy imposes a relatively high cognitive load. The plausibility of this claim rests heavily on the suggestion that the means-ends system is minimal. If it is minimal, then attempts to provide more realistic models of cognitive processes would result in expansions rather than contractions of the system.

The means-ends model of Larkin, McDermott, Simon, and Simon (1980) provides an example of a system that has many more productions than the current model. There are several classes of these productions omitted from the current system. For example, Larkin et al.'s model has productions which assign symbols to appropriate statements of the problem description rather than have the symbols and their status (known, unknown, goal) placed initially in working memory. As another example, productions are used to generate equations rather than assuming, as does the current model, that the relevant set of equations is known and in working memory. The inclusion of productions such as these is reasonable in a system designed to model as many aspects as possible of problem-solving performance. They are not nevertheless, essential and for this reason had to be excluded from a model pared to the bare minimum.

The nonspecific goal subsystem is of course, also minimal. It is nevertheless, plausible. Problem solvers faced with a nonspecific goal problem do simply search for equations that can be applied to a given problem state to enable solution of an unknown (see Sweller, Mawer, & Ward, 1983). The goal-free production mimics this activity.

Relations Between Cognitive Load and Number of Productions

Why should the number of productions (and elements that require matching) provide a measure of cognitive load? A qualitative analysis of the means-ends subsystem provides some evidence for the suggestion that it imposes a heavy cognitive load. A problem solver who is processing information in a manner similar to the means-ends subsystem, must at any given problem state decide whether there is an equation providing a single step solution (first production), whether the value for a desired subgoal can be calculated (fourth production), whether a chain of subgoals should be constructed (search in depth using the third production) or whether a series of unrelated subgoals should be established (search in breadth using the second and/or third productions). Each of these decisions must be determined by the relation between the current problem state and the goal, keeping in mind the available problem-solving operators (the equations). The four productions may accurately represent the difficulties faced by a problem solver using a means-ends strategy. It seems improbable that these complex decisions are automated and can run without a considerable strain on cognitive resources.

The issue may be examined in terms of a "human production system." At each choice point, he or she must attempt to match the known values with equations containing the goal to see if a solution is available; decide whether this is a futile exercise at this point; decide on the basis of available unknowns whether subgoals should be set up; decide on the basis of the knowns, unknowns, and equations which subgoal track should be followed; decide on the basis of the known values and available equations whether a value for a subgoal can be calculated; decide whether this is a futile exercise. It may not be surprising that under these circumstances the "system" frequently collapses (the problem solver gives up). Furthermore, none of these processes appears to be related to schema acquisition. Any learning processes must be imposed as additional mechanisms requiring additional cognitive capacity.

These processes may be contrasted with those imposed by a nonspecific goal problem. There is only one decision that needs to be made. Can a value be found for an unknown? It may be reasonable to suggest that this simple process poses little impediment to learning. It may not be surprising that the complex means-ends process can block learning.

A diagrammatic representation of the differences between the two strategies may be obtained by comparing Figure 1, which diagrams the flow of control under a means-ends strategy, with Figure 2, which diagrams the flow of control under a nonspecific goal strategy. The differences between the two diagrams may well reflect the differences in cognitive capacity required by the two processes.

It may be argued that in order to demonstrate that a heavy cognitive load during problem solving interferes with learning, it is important that learning

mechanisms be included in the model. In fact, such interference is inevitable assuming that: (a) the system (or person) has a fixed cognitive capacity; (b) both problem solving and learning require some of that capacity; (c) the problem solving and learning mechanisms differ. As indicated previously, these assumptions can explain the data. Under these assumptions, any increase in resources required during problem solving must inevitably decrease resources available for learning. Consequently, only the first step—evidence that required processing capacity may be relatively heavy during the means-ends analysis—is necessary. Consequences for learning follow inevitably.

In summary, a production system approach was used to provide some indication of the relative cognitive load of a means-ends compared to a non-specific goal strategy. Analysis of a system consisting of no more than the bare essentials needed to operate the strategies revealed that means-ends analysis required somewhat more information in working memory and a substantially more complex production system than a nonspecific goal strategy. This could be interpreted as suggesting that means-ends analysis requires more cognitive capacity than a nonspecific goal strategy¹

EXPERIMENTAL EVIDENCE

In previous sections, experimental results were discussed which suggested that means-ends analysis could impose a relatively heavy cognitive load. This evidence was indirect, consisting primarily of performance characteristics such as strategies used, categorization of solutions, speed of solution and errors on subsequent problems. Additional, stronger evidence for the contention was obtained by formal modelling techniques. There also is some relatively direct experimental evidence which is available.

If means-ends analysis imposes a heavy cognitive load, we might expect its use to simultaneously influence aspects of performance such as number of errors. In a mathematical task, novice problem solvers who do not have a substantial facility in the use of essential problem-solving operators—normally mathematical principles, equations, theorems, etc.—may be more likely to commit mathematical errors when using means-ends analysis than

¹ It should be noted that the implemented system required separate productions to handle equations containing two variables (e.g., $v = .5V$) and three variables (e.g., $s = vt$). Thus, each of the productions listed in Table 1 consisted of two productions in the implemented system. Furthermore, in the case of means-ends productions dealing with equations containing three variables, separate productions were required to generate subgoals where only one unknown existed as opposed to two unknowns. These variations resulted in ten means-ends productions and two goal-free productions. Additional productions, similar in structure to the existing ones, would need to be added to handle equations with more than three variables. The variations were necessitated by purely computational considerations and were not thought to have psychological significance. For this reason they have not been discussed in detail.

when using a nonspecific goal strategy. We might expect that if a strategy requires a large amount of cognitive processing capacity, less will be available for other, competing aspects of the task. Errorless use of, for example, equations, is one such aspect.

Results bearing directly on this issue have been obtained by Owen and Sweller (1985) in a series of studies. We used nonspecific goal trigonometry problems in which subjects were required to find the lengths of all the sides of a given diagram. This could be contrasted with the performance of subjects presented with a conventional problem consisting of the same diagram of which a specific side had to be found. Unlike previous experiments employing this paradigm (e.g. Sweller, Mawer, & Ward, 1983), the problems were not specifically structured to ensure that nonspecific goal subjects calculated the lengths of the same sides as those presented with the conventional problems. The groups were matched with respect to time spent on the problems rather than number of problems. This meant that while both groups spend a fixed amount of time solving problems, both the number of problems and number of sides calculated could vary.

Over several experiments, the major finding was that the conventional group made significantly more mathematical errors (e.g., misuse of the equation $\text{sine} = \text{opposite}/\text{hypotenuse}$) per side calculated as the nonspecific goal group. Four to six times as many mathematical errors were made by the conventional group. In fact, the total number of errors made by the conventional groups were consistently greater despite the fact that these groups consistently found fewer sides. As was the case in previous experiments, the advantage of nonspecific goal problems transferred to subsequent problems with fewer errors and faster performance in later problem solving.

The most obvious explanation for these results is in terms of the previously outlined model. Problem solvers organizing a problem according to means-ends principles, suffer from a cognitive overload which leaves little capacity for other aspects of the task. This overload can be manifested by an increase in the number of mathematical errors made.

The results of the Owen and Sweller (1985) experiments, in conjunction with the prior theoretical analysis, suggest that simultaneously solving a problem by means-ends analysis and attempting to acquire schemas associated with the problem, may be analogous to a dual task. Attempting to solve the problem may be considered the primary task. Acquiring knowledge of the problem structure and of elements which might facilitate subsequent solution attempts may be considered the secondary task.

A considerable volume of work has been carried out using a dual-task paradigm. Britton and his colleagues (Britton, Glynn, Meyer, & Penland, 1982; Britten, Holdredge, Curry, & Westbrook, 1979; Britton & Tesser, 1982) have used reaction times to a click as the secondary task with a variety of complex cognitive tasks as the primary task. They found that the second-

ary task could be used to indicate the cognitive capacity required by the primary task. Lansman and Hunt (1982) found that a secondary, reaction time task could be used to measure how much spare capacity was available while engaged on an easy primary task. This, in turn, could be used to predict performance on a subsequent, more difficult task. Fisk and Schneider (1984) obtained results indicating that increasing the extent to which subjects were required to attend to one task during controlled processing reduced the extent to which items were stored in long-term memory on a second task. Book and Garling (1980) and Lindberg and Garling (1982) found that knowledge concerning a traversed path was interfered with if subjects had to count backward while traversing the path.

All of these findings suggest strongly that a secondary task can be used as an indicator of the cognitive load imposed by a primary task. If, as suggested above, problem solving search via means-ends analysis and schema acquisition are independent tasks, then they may be considered as primary and secondary tasks respectively, within a dual task paradigm. Under these circumstances, if a strategy such as means-ends analysis is used to accomplish the primary task (attain the problem goal), then because the strategy imposes a heavy cognitive load, fewer resources may be available for the secondary task. Performance on aspects of the secondary task such as correct use of mathematical rules may be used to indicate the cognitive load imposed by the primary task. The Owen and Sweller (1985) experiments, in effect, used this procedure. Nevertheless, a direct use of the dual task paradigm may provide more evidence for the hypothesis that means-ends analysis imposes a heavy cognitive load.

An indication of the relative cognitive load imposed by means-ends and nonspecific goal strategies may be obtained by explicitly requiring subjects to engage in a secondary task while solving conventional or nonspecific goal problems. In the current experiment the secondary task was memory of the givens and the solution of a previously solved problem. This secondary task was chosen because it was thought that enhanced memory of these characteristics could provide some evidence of schema acquisition. If a schema allows subjects to classify a problem and indicates which moves are appropriate, then we might expect that enhanced memory of problem givens and solutions indicates enhanced schema acquisition. (Although it must, of course, be recognized that a schema requires more than just memory of givens and solutions. Nevertheless, these may be prerequisites for schema acquisition.)

Unlike previous experiments employing a nonspecific goal procedure (Owen & Sweller, 1985; Sweller, Mawer, & Ward, 1983) the current experiment was not designed to test whether a nonspecific goal procedure enhanced subsequent problem-solving performance. It was designed solely to test whether reducing cognitive load by reducing problem-solving search activity

using means-ends analysis could allow subjects to learn more of specific aspects of the problem. Problem solving search activity was reduced by presenting subjects with nonspecific goal trigonometry problems. Subjects were required to find the lengths of all but those sides not needed to solve equivalent conventional problems with conventional goals. The sides not required were indicated. A second group was presented with the equivalent conventional problems. Both groups were required to memorize and later reproduce the givens and solutions of the problems as a secondary task.

Two opposing hypotheses can be considered. First, as argued above, an increased cognitive load imposed by the conventional problems may decrease performance on the secondary task. Second, because the primary task requires subjects to deal with all of the elements of the secondary task, it can be hypothesized that increased cognitive load on the primary task, by strengthening those elements, should improve performance on the secondary task.

METHOD

Subjects

The subjects were 24 students from a Year 10 (age 15–16 years) class of a Sydney high school. All had been introduced previously to the sine, cosine, and tangent ratios.

Procedure

All subjects, tested individually, were presented a sheet explaining and giving examples of the use of the sine, cosine, and tangent ratios. When subjects were satisfied that the material was understood they were informed that they would be required to solve 6 problems. They were also told that after each problem was solved they would be required to precisely reproduce the original diagram and the correct solution of the problem preceding the one that had just been solved. There was no reproduction phase after the first problem (since there was no preceding problem to reproduce) and the fifth problem was the last requiring reproduction. Subjects were also told that their major task was to solve the problem. The problem statement and diagram were removed immediately after the last solution step had been taken. If on any problem the solution had not been obtained within 5 minutes, the experimenter provided the correct solution. There was no time limit on the reproduction phases. Each phase ended when subjects were satisfied that they could not improve their reproduction any further. Pencil and paper were used for both the problem solving and reproduction phases. Time and errors for each of the solution and reproduction phases were recorded.

A conventional and nonspecific goal group of 12 subjects each were used. Figure 3 provides an example of a conventional problem. In order to solve




Figure 3. Example Trigonometry Problem (If the goal of the problem is to find the length of CB, the solution is $CA = \text{sine } 35/4.4$; $CB = CA/\text{cosine } 49$.)

this problem the sine ratio needs to be used followed by the cosine ratio. The two trigonometric ratios needed to solve the remaining 5 problems were sine-sine, cosine-cosine, cosine-tangent, sine-tangent, and tangent-tangent, respectively. Each problem was identical in diagrammatic configuration to Figure 3 but the line segment labels (vertices) and the angles varied. The non-specific goal problems were identical except that subjects were told to find the lengths of as many sides as possible and the two sides not needed to solve the conventional problems were marked to indicate that they should not be solved for. No numerical values were required with subjects merely being asked to indicate the equations needed to solve the problems.

Results and Discussion

Table 3 indicates mean time to solve each problem for the two groups. (Non-solvers were allocated a time of 300 s.) There was no difference in total time to solve the 6 problems between groups, $F(1,22) = .17$. (The .05 level of significance is used throughout this article.) Nevertheless, it might be noted that on each of the 6 problems the goal-free group required marginally less time than the conventional group and the probability of this occurring by chance is .016. Trend analysis indicated a significant linear trend with later problems being solved more rapidly than earlier ones, $F(1,22) = 49.38$. There

TABLE 3
Mean Seconds to Problem Solution
(Mean mathematical errors are in brackets.)

Group	Problem					
	1	2	3	4	5	6
Conventional	274 (1.2)	239 (1.0)	170 (0.3)	188 (0.5)	193 (0.5)	138 (0.6)
NonSpecific Goal	272 (1.2)	227 (1.1)	156 (0.25)	170 (0.7)	191 (0.75)	126 (0.2)

was no group X problem interaction suggesting that both groups improved at approximately the same rate, $F(1,22) = .01$.

Table 3 also indicates the number of mathematical errors made where mathematical errors include algebraic errors or trigonometric errors such as defining the sine ratio as adjacent/hypotenuse. The results duplicate those for solution times with no difference between groups, $F(1,22) = 0$, a significant linear trend with fewer errors being made on later problems, $F(1,22) = 19.97$, and no group X problem interaction, $F(1,22) = .37$.

Table 4 indicates the number of subjects who were able to solve each problem within the allotted 5 minutes per problem. (Subjects who could not solve a problem were given the solution by the experimenter.) Most subjects could not solve the first problem but improved rapidly thereafter. Table 4 also indicates the mean number of sides calculated by each group on each problem.

Table 5 indicates mean reproduction times for the two groups on each of the 5 reproductions. There was no difference between the two groups in total time, $F(1,22) = .01$. A significant linear trend was obtained, $F(1,22) = 4.76$

TABLE 4
Number of Subjects Reaching Problem Solution Within 5 Minutes.
(Mean number of sides calculated in parentheses.)

Group	Problem					
	1	2	3	4	5	6
Conventional	4 (1.08)	5 (1.42)	10 (1.83)	9 (1.75)	9 (1.67)	11 (1.83)
NonSpecific Goal	2 (1.25)	8 (1.75)	11 (1.92)	11 (1.92)	8 (1.58)	10 (1.83)

TABLE 5
Mean Seconds for Reproduction

Group	Problem				
	1	2	3	4	5
Conventional	232	202	206	166	143
NonSpecific Goal	209	175	166	199	192

indicating increased speed in reproduction over problems. A nonsignificant group X problem interaction was obtained, $F(1,22) = 3.83$.

Errors on the reproduction task are the major focus of this experiment. Six categories of errors were used.

1. A segment-labelling error was scored if any of the line segments was incorrectly labelled—e.g., if line AB was labelled CD.
2. Angle-position errors occurred when an angle was incorrectly stated as a given or an unknown.
3. Angle-value errors occurred when a numerical value was incorrect. (If a new angle value appeared in a new position which had previously been an unknown, then both an angle-value and angle-position error was scored.)
4. Side-position errors occurred when an unknown side was labelled as known or vice-versa.
5. Side-value errors were scored when the given side was given the wrong length. (Side-position and side-value errors were scored when a value which had not appeared in the original was given to an originally unknown side.)
6. Solution errors were incorrect reproductions of the solution.

Each reproduction by each subject was given a score of 1 or 0 on each of these 6 criteria for each of the 5 problems. Any error on any of the criteria resulted in a score of 1 on that criterion. By adding across problems, a total score out of 5 could be obtained for each subject. Table 6 provides mean scores on each of the criteria for the two groups.

Because all scores fell in the limited range 0–5, severe floor or ceiling effects were obtained resulting in grossly skewed distributions. For this reason, nonparametric techniques were used to analyse this data. Mann-Whitney U tests were used to analyse differences between groups on either the segment-labelling errors, $U(12,12) = 55$, or the angle-value errors, $U(12,12) = 63$. The nonspecific goal group made significantly fewer angle-position errors, $U(12,12) = 29$, side-value errors, $U(12,12) = 27$, side-position errors, $U(12,12) = 41$, and solution errors, $U(12,12) = 17.5$.

The nature of the secondary task allowed 2 opposing hypotheses to be tested. Depending on the cognitive mechanisms that operate, a heavy cogni-

TABLE 6
Mean Reproduction Error Scores for Each Category

Group	Error Type					
	Segment-labels	Angle-value	Angle-position	Side-value	Side-position	Solution
Conventional	2.9	3.6	2.4	2.6	1.6	3.4
NonSpecific Goal	2.4	3.4	1.1	1.3	1.2	1.5

tive load during conventional problem solving relative to a nonspecific goal task, could either facilitate or inhibit memory of the problem structure. The results provide no evidence that an increased load during conventional problem solving assists problem solvers in assimilating information concerning the initial problem structure or the solution steps. Instead, the results conform to those obtained in more conventional dual task experiments with the cognitive load imposed by one task interfering with performance on the other. More excess capacity appears to be available after solving a nonspecific goal problem than a conventional problem.

The pattern of differences between the conventional and nonspecific goal problems also is of interest. A schema was defined above as a structure that permits problem solvers to categorize a problem as one which allows certain moves for solution. If a nonspecific goal problem enhances schema acquisition, we might expect subjects solving nonspecific goal problems to have superior recall of structural aspects of the problem such as the characteristics of the givens and the solution moves. In the current experiment, this should translate into superior performance on the angle-position, side-position, and solution measures since these are required to construct a schema. The nonspecific goal group was superior on all of these. The other measures—segment-labelling, angle-value and side-value—are presumably irrelevant to schema acquisition since a schema could be induced without reference to these details. Side-value was the only one of these measures providing a significant difference between groups. This may provide limited evidence that subjects followed instructions to concentrate on problem solution rather than the recall task. In the process, where cognitive capacity was available, they may have learned more of those characteristics needed to facilitate schema acquisition rather than irrelevant aspects of the problem. Thus the particular pattern of results may also be used to support the general hypothesis that cognitive load under conventional problem-solving conditions interferes with schema acquisition.

It might be noted, the results on the primary task (problem solution) do not replicate those obtained by Owen and Sweller (1985) or Sweller, Mawer, and Ward (1983) who repeatedly obtained improved performance under nonspecific goal as opposed to conventional conditions. Nevertheless, it should be noted that the previous studies did not use a secondary task. Assuming a fixed cognitive capacity for each subject, an advantage due to reduced cognitive load can be distributed over both tasks in a dual task exercise or concentrated over one or the other of the two tasks. To some extent, this distribution of cognitive resources is at the discretion of the subjects. In the present experiment, it appears that most subjects allocated excess capacity to the second task resulting in the usual performance difference occurring on that task rather than the primary one. In these terms, the results are in accord with previous findings. Nonspecific goal conditions have facilitated

performance. The only difference from previous experiments is that the major facilitation has been transferred from the primary to the secondary task. (In this context, it should be noted that the goal-modified group took less time to solve each of the 6 problems than the conventional group and this is unlikely to be a chance effect.)

THEORETICAL AND PRACTICAL IMPLICATIONS

Two general inferences can be drawn. First, it may be tentatively suggested that the use of computational models as measuring devices could have a more general applicability. Currently, there is no a priori method for determining the difficulty problem solvers will have in solving specific problems. Presenting the problem is the only viable technique. Inability to determine something as fundamental as problem difficulty may be a major impediment to progress. This gap in our basic technical repertoire is nevertheless understandable. A measure which simultaneously accounts for problem and problem solver characteristics is bound to be complex. By using minimal computational models it may be possible to simultaneously isolate and measure those aspects of a problem and a problem solver's strategy and knowledge that govern important aspects of performance.

The second conclusion, based partly on using a computational model as a measuring device, may be put more strongly. Conventional problem solving activity via means-ends analysis normally leads to problem-solution, not to schema acquisition. Both theoretical and practical implications flow from this conclusion.

The theoretical points made in the present paper suggest that cognitive effort expended during conventional problem solving leads to the problem goal, not to learning. Goal attainment and schema acquisition may be two largely unrelated and even incompatible processes. This may be relevant to all learning through problem-solving theories (e.g., see Anderson, 1982; Laird, Newell, & Rosenbloom, 1987).

The suggestions made in this article have clear applications, especially in an educational context. Most mathematics and mathematics-based curricula place a heavy emphasis on conventional problem solving as a learning device. Once basic principles have been explained and a limited number of worked examples demonstrated, students are normally required to solve substantial numbers of problems. Much time tends to be devoted to problem solving and as a consequence, considerable learning probably occurs during this period. The emphasis on problem solving is nevertheless, based more on tradition than on research findings. There seems to be no clear evidence that conventional problem solving is an efficient learning device and considerable evidence that it is not. If, as suggested here, conventional problems impose a heavy cognitive load which does not assist in learning, they may be

better replaced by nonspecific goal problems or worked examples (see Sweller & Cooper, 1985). The use of conventional problems should be reserved for tests and perhaps as a motivational device.

CONCLUSIONS

In summary we may conclude: (1) Both experimental evidence and theoretical analysis suggest that conventional problem solving through means-ends analysis may impose a heavy cognitive load; (2) The mechanisms required for problem solving and schema acquisition may be substantially distinct; (3) As a consequence, the cognitive effort required by conventional problem solving may not assist in schema acquisition; (4) Since schema acquisition is possibly the most important component of problem solving expertise, the development of expertise may be retarded by a heavy emphasis on problem solving; (5) Current theories and practice frequently assume problem solving is an effective means of learning and consequently may require modification.

■ Original Submission Date: ???.

REFERENCES

- Anderson, J. (1982). Acquisition of cognitive skill. *Psychological Review*, *89*, 369-406.
- Book, A., & Garling, T. (1980). Processing of information about location during locomotion. Effects of a concurrent task and locomotion patterns. *Scandinavian Journal of Psychology*, *21*, 185-192.
- Britton, B., Glynn, S., Meyer, B., & Penland, M. (1982). Effects of text structure on use of cognitive capacity during reading. *Journal of Educational Psychology*, *74*, 51-61.
- Britton, B., Holdredge, T., Curry, C., & Westbrook, R. (1979). Use of cognitive capacity in reading identical texts with different amounts of discourse level meaning. *Journal of Experimental Psychology: Human Learning and Memory*, *5*, 262-270.
- Britton, B., & Tesser, A. (1982). Effects of prior knowledge on use of cognitive capacity in three complex cognitive tasks. *Journal of Verbal Learning and Verbal Behaviour*, *21*, 421-436.
- Bruner, J. (1961). The act of discovery. *Harvard Educational Review*, *31*, 21-32.
- Chase, W., & Simon, H. (1973a). Perception in chess. *Cognitive Psychology*, *4*, 55-81.
- Chase, W., & Simon, H. (1973b). The mind's eye in chess. In W.G. Chase (Ed.), *Visual information processing*. NY: Academic.
- Chi, M., Glaser, R., & Rees, E. (1982). Expertise in problem solving. In R. Sternberg (Ed.), *Advances in the psychology of human intelligence*. NJ: Erlbaum.
- Cooper, G., & Sweller, J. (1987). Effects of schema acquisition and rule induction on mathematical problem-solving transfer. *Journal of Educational Psychology*, *79*, 347-362.
- De Groot, A. (1966). Perception and memory versus thought: Some old ideas and recent findings. In B. Kleinmuntz (Ed.), *Problem solving*. NY: Wiley.
- Dewey, J. (1910). *How we think*. Boston: Heath.
- Dewey, J. (1916). Method in science teaching. *Science Education*, *1*, 3-9.
- Egan, D., & Schwartz, B. (1979). Chunking in recall of symbolic drawings. *Memory and Cognition*, *7*, 149-158.

- Fisk, A., & Schneider, W. (1984). Memory as a function of attention, level of processing, and Automatization. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 10, 181-197.
- Greeno, J. (1978). Natures of problem solving abilities. In W.K. Estes (Ed.), *Handbook of learning and cognitive processes* (Vol. 5). Hillsdale, NJ: Erlbaum.
- Hinsley, D., Hayes, J., & Simon, H. (1977). From words to equations: Meaning and representation in algebra word problems. In P. Carpenter & M. Just (Eds.), *Cognitive processes in comprehension*. Hillsdale, NJ: Erlbaum.
- Jeffries, R., Turner, A., Polson, P., & Atwood, M. (1981). Processes involved in designing software. In J.R. Anderson (Ed.) *Cognitive skills and their acquisition*. Hillsdale, NJ: Erlbaum.
- Laird, J., Newell, A., & Rosenbloom, P. (1987). SOAR: An architecture for general intelligence. *Artificial Intelligence*, 33, 1-64.
- Langley, P., & Neches, R. (1981). PRISM user's manual. Technical report, Department of Psychology, Carnegie-Mellon University and Learning Research and Development Center, University of Pittsburgh.
- Lansman, M., & Hunt, E. (1982). Individual differences in secondary task performance. *Memory and Cognition*, 10, 10-24.
- Larkin, J., McDermott, J., Simon, D., & Simon, H. (1980). Models of competence in solving physics problems. *Cognitive Science*, 4, 317-348.
- Lewis, M., & Anderson, J. (1985). Discrimination of operator schemata in problem solving: Learning from examples. *Cognitive Psychology*, 17, 26-65.
- Lindberg, E., & Garling, T. (1982). Acquisition of locational information about reference points during locomotion: The role of central information processing. *Scandinavian Journal of Psychology*, 23, 207-218.
- Mawer, R., & Sweller, J. (1982). The effects of subgoal density and location on learning during problem solving. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 8, 252-259.
- Owen, E., & Sweller, J. (1985). What do students learn while solving mathematics problems? *Journal of Educational Psychology*, 77, 272-284.
- Simon, D., & Simon, H. (1978). Individual differences in solving physics problems. In R. Siegler, (Ed.), *Children's thinking: What develops?* NJ: Erlbaum.
- Sweller, J. (1983). Control mechanisms in problem solving. *Memory and Cognition*, 11, 32-40.
- Sweller, J., & Cooper, G. (1985). The use of worked examples as a substitute for problem solving in learning algebra. *Cognition and Instruction*, 2, 59-89.
- 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., & Howe, W. (1982). Consequences of history-cued and means-ends strategies in problem solving. *American Journal of Psychology*, 95, 455-483.
- Sweller, J., Mawer, R., & Ward, M. (1983). Development of expertise in mathematical problem solving. *Journal of Experimental Psychology: General*, 112, 639-661.
- Voss, J., Vesonder, G., & Spilich, G. (1980). Text generation and recall by high-knowledge and low-knowledge individuals. *Journal of Verbal Learning and Verbal Behaviour*, 19, 651-667.