

Resource-adaptive Selection of Strategies in Learning from Worked-Out Examples

Peter Gerjets, Katharina Scheiter & Werner H. Tack

Collaborative Research Center 378: Resource-Adaptive Cognitive Processes

University of the Saarland

D-66041 Saarbruecken/Germany

{pgerjet, katharis, tack}@cops.uni-sb.de

Abstract

Most tasks can be pursued by using different strategies (Logan, 1985; Reder & Schunn, 1998). In this paper we focus on strategies of learning from worked-out examples. Within a resource-oriented framework these different strategies can be classified according to their costs and benefits. These features may determine which strategy will be selected for accomplishing a task in situations with certain resource limitations. We investigate specific hypotheses about strategic adaptations to resource limitations (e.g., time pressure or lack of prior knowledge) within a hypertext-based learning environment. A comparison of the strategy selection of good and poor learners is used to assess the degree of subjects' resource adaptivity. Ideas for modeling resource-adaptive selection of strategies within the ACT-R architecture are discussed.

Resource-Adaptive Selection of Strategies

According to Reder and Schunn (1998) individual performance differences in learning and problem-solving tasks may not only depend on the *variability of cognitive parameters* (e.g. speed of processing, working-memory capacity) or on interindividual differences in the *availability of strategies* for solving the same task. Instead subjects may differ with regard to their *ability of shifting strategies* as a consequence of changes in task demands or other situational parameters. Therefore, the *adaptive selection of strategies* should be of major importance for success in learning and problem solving. Theoretically, *associative approaches* explain strategy selection as a reaction to cues related to certain strategies (cf. Reder & Schunn, 1998). On the contrary, *rational approaches* assume that subjects choose strategies according to their costs and benefits in terms of resource demands and expected utility (cf. Payne, Bettman & Johnson, 1993; Logan, 1985). In our paper we prefer a rational approach which is based on a wide conception of resources comprising all internal and external means that are useful or necessary for solving a specific task. We focus on internal resources like prior knowledge and external resources like learning time and external information. The *costs* of adopting a specific strategy increase with its resource demands. Besides differences in costs, strategies may additionally differ with respect to their *benefits* (e.g., effectiveness in solving the task at hand, success in solving subsequent tasks, acquisition of different kinds of knowledge).

To describe processes of strategy selection within a resource-based framework *two different types of resource adaptivity* have to be distinguished: (a) On the one hand evolution may have forced cognitive systems to generally employ *resource-adapted strategies*, i.e. strategies that do not lead to optimal task performance but that are compatible with the usual limitations of processing resources. According to this assumption resource-adapted behavior will be even displayed in situations with relatively high resource availability. A well known theoretical approach to resource-adapted behavior

that may be applied to strategy selection is the concept of *satisficing* (Simon, 1990). According to this concept bounded rational agents do not select the most effective strategy for solving a task but rather set a specific aspiration level (probably associated with the value of the respective goal) and select a strategy that exceeds it. (b) On the other hand cognitive systems may be *resource-adapting* in that the strategies employed to pursue a certain task are additionally constrained by the configuration of resources currently available for the agent. If resources like knowledge, time or external information are restricted, strategies should be adopted that are less demanding with respect to these resources. These strategic shifts may compensate for performance impairments expectable without such adaptations. Severe limitations of specific resources or certain combinations of resource limitations may prove impossible to compensate. Taken together it can be postulated that subjects generally choose satisficing strategies even when no strong resource limitations are present. Additionally they should adapt to specific limitations by choosing strategies that are more frugal with respect to these limited resources. Therefore, our main aim is to determine whether subjects working under certain resource limitations employ the same strategies as subjects without such limitations or whether they adapt to these limitations in a useful way.

Learning from Worked-Out Examples

In our empirical work we are specifically interested in strategies of learning from worked-out examples from the domain of probability word problems. Worked-out examples are instances of a certain problem type together with a detailed solution. They facilitate the learning of abstract procedures for later problem solving (Cummins, 1992; Catrambone, 1998) and the solving of novel problems by analogy (Gick & Holyoak, 1983; Reed, 1999). Prerequisites for the use of examples for knowledge acquisition and application are the generation of suitable example representations and the initiation of appropriate cognitive processes working on these representations. For our purposes strategies of learning from worked-out examples can be described on two dimensions: *rare versus frequent use* of examples and *brief versus extensive use* of examples. Van Lehn and Jones (1993) found that better learners preferred a rare use of examples and tried to solve training problems on their own. While good learners only inspected examples for getting specific information, poor learners referred back to examples as often as possible.

Beyond differences in the frequency of example use learners can use examples more or less extensively depending on the degree of example elaboration during learning. These elaborations may comprise the abstract deep structure of an example problem, the subgoal structure of its solution (Catrambone, 1998), or the similarities and differences of the example compared to other examples from the same problem type

(Cummins, 1992). Chi, Bassok, Lewis, Reimann & Glaser (1989) found that learners who elaborated examples during study substantially differed in their performance from learners that didn't elaborate on example problems. According to Chandler and Sweller (1991, p. 294) these results "indicate the importance to learning of an ability to properly process worked examples". Therefore, strategy selection in learning from worked-out examples may have a major influence on the quality of knowledge acquired.

Hypotheses

Based on the underlying theoretical framework we derived five experimental hypotheses about resource-limitations in learning from worked-out examples and about adaptive strategy selection: (a) Learning may be impaired if relevant resources like learning time, prior knowledge or external information are severely limited. (b) Different resource limitations may not act additively but interact with each other if more than one resource is limited. Therefore, different combinations of resource-limitations may result in different patterns of performance impairments. (c) Strategy shifts may help compensate for performance impairments associated with certain resource configurations, though some combinations of resource limitations may prove impossible to compensate by strategic choices. Hence, good and poor learners may differ in strategic variables under some but not under all resource configurations. (d) It can be expected that subjects select faster but less accurate processing strategies if learning time is limited. In the case of limited prior knowledge or external information, subjects may select less information-demanding strategies even if these strategies involve time-consuming inferences. (e) The dimensions rare - frequent use of examples and brief - extensive use of examples should be useful to characterize strategies for learning from worked-out examples and to describe relevant strategy shifts.

To investigate these hypotheses we conducted a series of three experiments in which subjects' had to work on a learning and problem-solving task from the domain of probability word problems. We developed a hypertext system to serve as experimental environment that allowed us to log subjects' strategic decisions in great detail. With regard to our hypotheses a question of central importance is whether possible strategic differences between experimental conditions can be interpreted as adaptive. To answer this question we employed the method of contrasting *strategic differences between experimental conditions* with *strategic differences between good and poor learners* within experimental conditions. This approach allows us to decide whether subjects who learn under a specific configuration of resource limitations change their behavior in a direction that can be identified as useful given this configuration of resources.

Experiment 1

Method

Participants The subjects were 46 students of the University of the Saarland (UdS), Germany who either participated for course credit or payment. Average age was 24.5 years.

Materials and procedure In the hypertext environment a short introduction to the domain of combinatorics was pre-

sented and subjects were instructed to solve a number of *probability word problems* following a *self-paced learning phase*. In the learning phase of the experiment subjects could retrieve *abstract explanations* of six solution principles from the domain of combinatorics (with their associated formula) by clicking on the respective links in the navigation bar. In the *test phase* the instructional information of the learning phase was no longer available. Three test problems were presented on the screen and one of the test problems had to be selected to begin with. In this experiment *no worked-out examples* were included because in the first step we wanted to study performance and strategy selection in learning with purely abstract information.

Design and dependent measures As independent variables time pressure and prior knowledge were manipulated by implementing three different learning conditions. In the *baseline condition with high resource availability* subjects possessed relatively high domain-specific prior knowledge and were instructed to take as much time as needed to understand the solution principles and then to begin with the test phase by clicking on the respective link. In the *condition with low prior knowledge* learning time was likewise unlimited, but subjects were rather unfamiliar with the domain of combinatorics. In the *condition with low learning time* we restricted the learning time of subjects with high prior knowledge to seven minutes (i.e., about two thirds of the mean learning duration in the condition without limitations). To induce time pressure subjects were informed that they would only be granted two thirds of the time usually needed for the learning phase. When the learning time (visible for the subjects on a digital clock) expired, the first page of the problem solving phase was automatically presented on the screen and subjects were instructed to begin working on the test problems. During problem solving there were no time limits. In the test phase the subjects had to mark the appropriate solution principle and the values of two variables for each of the three test problems in a multiple-choice form available in the hypertext environment. No calculations had to be made. One error was assigned for each wrong answer. Problem-solving time as well as total learning time, mean reading time per abstract page presented and frequency of retrieving abstract information pages were recorded by using logfiles. Following the test phase subjects had to pass a knowledge test with ten multiple-choice questions related to abstract concepts from the domain of combinatorics. One error was assigned for each wrong answer. Similar conceptual questions were posed as a pretest at the beginning of the experiment to control for domain-specific prior knowledge. Additionally, we registered subjects' last math grade as a general measure of mathematical ability which ranged from grade one (best) to grade six (worst).

Results and Discussion

First we investigated whether the three learning conditions differ with regard to performance and strategy measures. For this reason, we used the *baseline condition with high resource availability* as a point of reference and contrasted its data with the two other conditions (see table 1).

Table 1: Means and significance of differences

Learning without worked-out examples	A: Low prior knowledge	B: Base-line	C: Low learn. time	Significance of Difference
Problem-solving errors	52.3 %	32.5 %	42.1 %	A >> B << C
Knowledge-test errors	35.4 %	8.9 %	17.1 %	A >> B = C
Math grade	2.2	2.1	1.7	A = B = C
Pretest errors	65.8 %	32.2 %	30.2 %	A >> B = C
Frequency / abst. info.	22	24	15	A = B > C
Mean time / abst. info.	84 sec.	69 sec.	31 sec.	A = B >> C
Total learning time	823 sec.	722 sec.	397 sec.	A = B >> C
Problem solving time	782 sec.	726 sec.	759 sec.	A = B = C

Note: >>: $p \leq .05$; >: $p \leq .10$; =: $p > .10$ (p-values result from one-tailed t-tests)

A comparison with the *low-prior-knowledge condition* (A versus B) reveals strong differences in problem-solving errors and knowledge-test errors while there are no differences with regard to strategic measures. This may imply that subjects with low prior knowledge don't try to compensate for their performance impairments by increasing problem-solving time or learning time if only abstract information about the solution principles is available in the learning environment. Comparing the baseline condition with the *low-learning-time condition* (B versus C) yields similar differences in problem-solving errors while there are no differences with regard to knowledge-test errors. In addition, both conditions differ with respect to strategic measures. Compared to subjects in the baseline condition subjects under time pressure retrieve abstract information pages less frequently and spend less time on each abstract information page. This change in strategic behavior is not obligatory as subjects could as well have reacted to time pressure by only reducing the mean time reading abstract information but not the retrieval frequency (as they do in experiment 2).

In a second step we evaluated the adaptivity of strategy shifts in experiment 1 by comparing good and poor learners within the experimental conditions with regard to the strategy measures listed in table 1 (post-hoc median splits according to problem-solving performance). To rule out the hypothesis that differences between good and poor learners are caused by differences in prior knowledge or math grade we inserted these variables as covariates in the statistical comparison of good and poor learners. The respective analyses of covariance reveal that there are no differences with regard to strategy measures distinguishing between good and poor learners. This implies that subjects' strategic options (modifying the frequency or intensity of processing abstract information) are unsuitable for improving problem-solving in the conditions with purely abstract information. Accordingly, efficiency impairments caused by restrictions in either prior knowledge or learning time cannot be easily compensated by strategic shifts in this experiment. Therefore, no resource-adaptive processes of strategy selection could be evidenced here. We conducted experiment 2 to investigate whether strategies of information processing are of greater importance in example-based learning.

Experiment 2

Method

Participants and materials The subjects were 46 students of the UdS who either participated for course credit or

payment. Average age was 24.5 years. In experiment 2 the hypertext environment was supplemented by a single worked-out example per solution principle.

Design and dependent measures The same three learning conditions as in experiment 1 were implemented. Time pressure was induced analogously to experiment 1 by restricting learning time to nine minutes. The learning environment was augmented by a single worked-out example for each solution principle. These examples as well as the abstract information of the learning phase were no longer available in the test phase. Dependent measures were problem-solving errors, knowledge-test errors, domain-specific prior knowledge, last math grade, mean reading time per example provided, frequency of example retrieval (number of clicks), mean reading time per abstract information page, frequency of abstract information retrieval (number of clicks), total learning time, and problem-solving time.

Results and Discussion

Compared to the baseline condition with high resource availability subjects in the *low-prior-knowledge condition* show substantial performance impairments in problem solving and in the knowledge test (see A versus B in table 2). Furthermore, there are significant differences with regard to strategic measures between the two experimental groups. Subjects with low prior knowledge spend more time on learning and especially show an increased frequency of retrieving abstract information as well as an increased mean time reading these pages. There are, however, no differences concerning the use of examples between the two groups. Comparing the baseline condition with the *low-learning-time condition* (B versus C) yields similar differences in problem-solving errors and knowledge-test errors. With respect to strategic measures, subjects in the low-learning-time condition retrieve examples less frequently and spend less time reading examples and abstract information. Interestingly, subjects under time pressure retrieve abstract information more often than baseline subjects.

Table 2: Means and significance of differences

Learning with one worked-out example	A: Low prior knowledge	B: Base-line	C: Low learn. time	Significance of Difference
Problem-solving errors	55.4 %	32.0 %	54.5 %	A >> B << C
Knowledge-test errors	36.8 %	13.1 %	20.9 %	A >> B < C
Math grade	2.6	1.9	2.0	A > B = C
Pretest errors	66.1 %	33.3 %	32.3 %	A >> B = C
Frequency / example	8	7	3	A = B >> C
Mean time / example	56 sec.	44 sec.	14 sec.	A = B >> C
Frequency / abst. info.	22	13	16	A >> B < C
Mean time / abst. info.	62 sec.	45 sec.	37 sec.	A >> B > C
Total learning time	1047 sec.	809 sec.	516 sec.	A > B >> C
Problem solving time	600 sec.	606 sec.	571 sec.	A = B = C

Note: >>: $p \leq .05$; >: $p \leq .10$; =: $p > .10$ (p-values result from one-tailed t-tests)

A comparison of good and poor learners within the experimental conditions reveals the following strategic differences: In the *baseline condition* good learners spend more time reading examples than poor learners. In the *low-prior-knowledge condition* there are no strategic differences between good and poor learners. This implies that the performance in this con-

dition may not easily be improved by strategic shifts. Nevertheless, subjects with low prior knowledge try to improve their performance by learning longer (increased frequency and time reading abstract information). This shift, however, only increases costs in terms of time investment but doesn't yield any benefits in terms of performance. Therefore, subjects in this condition don't behave resource-adaptive.

In the *low-learning-time condition* good learners invest more time reading examples than do poor learners. In the light of this finding, it can be recommended that subjects under time pressure should save time by reducing time for abstract information processing without simultaneously confining the processing of examples. As the data in table 2 reveal, subjects under time pressure do not follow this recommendation towards resource-adaptive behavior. They only show a slight reduction in the mean reading time per abstract information page while there is a substantial decrease in the mean time reading examples. The respective interaction is significant and indicates that no resource-adaptive strategy shift took place. To conclude, performance impairments due to lacking prior knowledge cannot be compensated by selecting different strategies. Therefore, subjects' attempts to improve performance are in vain. On the other hand, performance impairments due to time pressure may be compensated by focussing on example information. Unfortunately, subjects do not shift their strategies in this direction. Therefore, no resource-adaptive strategy selection could be found in experiment 2.

We finally compared all six conditions from experiment 1 and 2. Contrasting the two conditions with low prior knowledge doesn't reveal any decrease in problem-solving errors due to the provision of examples. However, subjects in the one-example condition need less time for problem solving which indicates a slight increase in overall efficiency. A similar pattern of results can be found for the two baseline conditions. Unexpectedly, subjects in the low-learning-time condition deteriorate significantly with regard to problem-solving errors when provided with one example. Their problem-solving time is decreased analogously to the two other resource conditions. The respective interaction between time pressure (with/ without) and example availability (with/ without) with regard to problem-solving errors is significant.

To sum up, in our experimental setting learning with examples doesn't seem to be more effective than learning with only abstract information. At least the mere provision of instructional examples is obviously not sufficient to improve learning. Rather, the availability of examples must be accompanied by an extensive example-processing. As the differences between good and poor learners in the baseline condition and in the low-learning-time condition reveal this is crucial to performance. Furthermore, we found first support for the assumption that different kinds of resources may interact with regard to their effects on learning and problem solving. The augmentation of abstract information with one worked-out example slightly improves problem solving (i.e., reduces problem-solving time) if prior knowledge is restricted while it can have detrimental effects on problem-solving errors in the case of time limitation. In order to test whether these effects can also be observed when providing more than one example we conducted a third experiment.

Experiment 3

Method

Participants and materials The subjects were 43 students of the UdS who either participated for course credit or payment. Average age was 24.7 years. In experiment 3 the hypertext environment was supplemented by three worked-out examples of varying complexity to illustrate the application of each solution principle to different problem situations.

Design and dependent measures The same three conditions as in experiment 1 and 2 were used in this experiment. Time pressure was induced by allowing 13 minutes for learning in the time-limited condition. Dependent measures were the same as in experiment 2.

Results and Discussion

Compared to the baseline condition with high resource availability subjects in the *low-prior-knowledge condition* again show an increase in both types of error rates (see table 3, A versus B). With regard to strategic measures, subjects with low prior knowledge spend less time reading examples but simultaneously show an increase in time reading abstract information. Their time for problem solving is slightly decreased. Surprisingly, the comparison between the baseline condition and the *low-learning-time condition* (B versus C) shows that time pressure does not lead to impairments in problem solving like it did in experiment 1 and 2. There are, however, differences in knowledge-test errors as expected. Concerning strategic measures, subjects under time pressure spend less mean time reading examples and retrieve examples less frequently.

Table 3: Means and significance of differences

Learning with three worked-out examples	A: Low prior knowledge	B: Baseline	C: Low learn. time	Significance of Difference
Problem-solving errors	50.3 %	32.5 %	28.9 %	A >> B = C
Knowledge-test errors	33.7 %	13.6 %	22.0 %	A >> B < C
Math grade	2.8	2.0	2.5	A >> B = C
Pretest errors	59.6 %	29.4 %	35.6 %	A >> B = C
Frequency / example	13	17	7	A = B >> C
Mean time / example	25 sec.	32 sec.	9 sec.	A < B >> C
Frequency / abst. info.	28	23	21	A = B = C
Mean time / abst. info.	71 sec.	49 sec.	54 sec.	A > B = C
Total learning time	1179 sec.	1153 sec.	751 sec.	A = B >> C
Problem solving time	522 sec.	640 sec.	753 sec.	A < B = C

Note: >>: $p \leq .05$; >: $p \leq .10$; =: $p > .10$ (p-values result from one-tailed t-tests)

Comparing good and poor learners within the *baseline condition* shows that good learners spend more time on learning (especially on abstract information pages) and more time on problem solving. In the *low-prior-knowledge condition* good learners' frequency of retrieving examples and of retrieving abstract information is increased as well as their mean time reading example pages. Hence, it would be resource-adaptive in this condition to study abstract information and example information more intensively and in a well-balanced way. However, subjects with low prior knowledge even show a reduced mean time reading examples compared to subjects in the baseline condition. Furthermore, there is a significant

cross-interaction between prior-knowledge (with/ without) and retrieval frequency of different instructional material (examples/ abstract information). This interaction shows that low-prior-knowledge subjects focus on the retrieval of abstract information instead of handling examples and abstract information in a well-balanced way. In the *low-learning-time condition* good learners show an increased frequency of retrieving abstract information and examples. Thus a useful recommendation to subjects working under time constraint could be to retrieve example information and abstract information in a well-balanced way. A significant cross-interaction between time pressure (with/ without) and retrieval frequency of different instructional material (examples/ abstract information) reveals that subjects under time pressure focus on the retrieval of abstract information instead of handling examples and abstract information in a well-balanced way. Their behavior can thus not be classified as resource-adaptive. However, this is the only condition in which time pressure does not lead to significant performance impairments. This unexpected finding can be explained by considering that subjects more or less ignored the examples provided and therefore could spend the same amount of time in processing abstract information as subjects without time pressure and without instructional examples (i.e., baseline condition in experiment 1). Accordingly, their performance is comparable to that condition.

Contrasting the results from experiment 2 and 3 reveals that subjects with three examples learning in the baseline condition and in the low-prior-knowledge condition do not perform any better than the respective subjects in the one-example conditions. As explained before, improvements under time pressure are presumably not attributable to the provision of three examples but rather to the fact that subjects ignore the examples to save time for processing abstract information. The augmentation of instructional resources to three examples therefore does not prove as beneficial as could be expected when considering theories of learning by analogy (Gick & Holyoak, 1983) or theories of learning from worked-out examples (Cummins, 1992; Quilici & Mayer, 1996). At least the *mere provision* of three examples is obviously not sufficient to improve learning. Rather, the provision of multiple examples must be accompanied by a balanced processing of example information and of abstract information in order to acquire the relevant knowledge for problem solving. As the differences between good and poor learners in each of the three-example conditions reveal this is crucial to performance. Contrary to subjects learning with one example who profit most from studying the example intensively subjects learning with three examples should equally focus on abstract information. This finding fits theoretical assumptions about schema abstraction and the acquisition of transferable knowledge according to which it is necessary to compare different examples with respect to relevant abstract properties to induce theoretical concepts that may be applicable to analogous problems (Cummins, 1992).

General Discussion

Contrary to our *first hypothesis (a)* we found that limitations of relevant resources are not always associated with performance impairments and accordingly that the provision of rele-

vant resources is not always associated with performance improvements. E.g., the provision of additional instructional information doesn't always improve problem-solving. It can even lead to impairments if subjects are overwhelmed by information selection and integration. This interpretation is in line with the fact that subjects with low learning time suffer from the provision of one example and that they resign from the processing of examples when provided with three examples. Furthermore, as postulated in our *second hypothesis (b)* effects of resource limitations are not always additive, but may even be cross-interacting. For example, the augmentation of instructional resources by worked-out examples is slightly beneficial for subjects with low prior knowledge (decreased problem-solving time), while it can even have harmful effects for subjects with low learning time (increased problem-solving errors). Contrary to our *third hypothesis (c)* no cases of resource-adaptive strategy shifts could be identified. There are no patterns of differences between experimental conditions that can be classified as adaptive with respect to differences between good and poor learners within these experimental conditions. Our *fourth hypothesis (d)* stating that subjects with limited learning time should select faster but less accurate example processing strategies was confirmed in experiment 2 and 3. Contrary to our expectations, subjects with low prior knowledge do not adopt more time-consuming strategies of example processing. Finally, as predicted in the *fifth hypothesis (e)*, the dimensions brief versus extensive use (time per example provided) and rare versus frequent use (frequency of example retrieval) are important dimensions for describing strategies of learning from worked-out examples. This can be inferred from the differences between good and poor learners and between experimental conditions with respect to these variables.

In conclusion, our experiments show that strategic options to improve one's learning performance become the more numerous the more instructional material is provided. At the same time it could be demonstrated that one has to make use of these strategic options, i.e., adopt adequate strategies in order to benefit from this additional information.

Cognitive Modeling Approach

In the next step we intend to develop a more detailed model of resource limitations and their influences on processes of strategy selection. Within a cognitive science framework high-level processes of executive control like strategy selection in learning and problem-solving may be best modeled by means of cognitive architectures that are designed as comprehensive theories of human cognitive abilities. As a theoretical basis for the cognitive modeling of strategy selection in learning from worked-out examples we will refer to the ACT-R architecture (Anderson & Lebiere, 1998) that is based on a rational analysis approach compatible with our framework of resource-adaptive strategy selection. If one defines strategies for performing tasks as sets of procedures or operations that may be adopted in order to implement a certain goal, strategies can be easily represented in ACT-R by sets of productions that are sufficient to solve a task successfully. Based on this representation, two mechanisms of ac-

tion control can be distinguished in ACT-R that are useful in modeling strategy selection.

On the one hand, processing in ACT-R is controlled by the *currently active goal*. Productions referring to other than the current goal cannot be selected for execution. Strategy selection by setting strategy-specific subgoals can be interpreted as a choice process that is based on discrete symbolic knowledge and may be useful to model more deliberate aspects of strategy selection. Accordingly, goal setting implies that the accomplishment of the current task is interrupted for a period of meta-level decision making.

On the other hand, control in ACT-R is determined by the *mechanism of conflict resolution* that selects one of the conflicting productions that are compatible with the current goal for the next processing step. Strategy selection based on conflict resolution may be described as a subsymbolic process embedded within the fundamental mechanisms of the architecture. Conflict resolution is assumed to be an automatic process that is not consciously accessible and accordingly is initiated without changes in the current goal of information processing. ACT-R's mechanism for conflict resolution is based on an estimation of the expected gain E of the conflicting productions. For every feasible production i the value of E is determined by the formula $E = P G - C$ with P being the expected probability of goal achievement when using i , G being the goal value, and C being the expected costs of goal achievement when using i . Within this framework the resource limitations studied in our experiments can be modeled as follows.

Time pressure In ACT-R the goal value G is operationalized by the maximum amount of time that may be invested for goal achievement. Costs of goal achievement C are likewise measured by the time needed for goal accomplishment. Based on these conventions the mechanism of conflict resolution inherently produces a speed-accuracy trade-off depending on the available time. Time pressure will result in a decrease of G , leading to a lower weight of success probability and a higher weight of processing costs in production selection. As a result, less effective but at the same time less costly strategies will be selected for task accomplishment. Thus ACT-R enables the modeling of subjects' adaptation to limitations in time resources by automatic strategy shifts.

Lack of prior knowledge Limitations in domain-specific knowledge can be represented as gaps in declarative knowledge, i.e., the appropriate conceptual apparatus to encode the instructional material. In ACT-R these limitations can be best represented by missing chunk-types (representing concepts). Thus gaps in prior knowledge cannot be compensated automatically. Rather a deliberative setting of a specific learning goal may be necessary to first initiate activities to acquire the required conceptual knowledge.

Limited external information If external information necessary for the execution of the production with the highest expected gain in the conflict set is lacking this production is automatically abandoned in ACT-R and a production with less expected gain is selected that matches the currently available information. Thus a task can be handled successfully as long as there are productions available whose information demands are satisfied by the current external informa-

tion. Augmenting external information beyond these minimal requirements will improve performance if this information can be encoded correctly and if there are productions available that properly use this information within the time available. To model the interaction between time limitations and example availability with regard to problem-solving errors we assume that subjects under time pressure may lack the necessary time to process example information properly. This may explain why the provision of one worked-out example is harmful for low-learning-time subjects while subjects with sufficient learning time don't show any efficiency impairments.

Acknowledgements

We thank Simon Albers, Tina Schorr, Julia Schuh, Markus Werkle and the German Research Foundation for their support.

References

- Anderson, J. R. & Lebiere, C. (1998). *The atomic components of thought*. Hillsdale, NJ: Erlbaum.
- Catrambone, R. (1998). The subgoal learning model: Creating better examples so that students can solve novel problems. *Journal of Experimental Psychology: General*, 127, 355-376.
- Chandler, P. & Sweller, J. (1991). Cognitive load theory and the format of instruction. *Cognition and Instruction*, 8, 293-332.
- Chi, M. T. H., Bassok, M., Lewis, M., Reimann, P. & Glaser, R. (1989). Self-explanations: How students study and use examples in learning to solve problems. *Cognitive Science*, 13, 145-182.
- Cummins, D. D. (1992). Role of analogical reasoning in the induction of problem categories. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 18, 1103-1124.
- Gick, M. L. & Holyoak, K. J. (1983). Schema induction and analogical transfer. *Cognitive Psychology*, 15, 1-38.
- Logan, G. D. (1985). Executive control of thought and action. *Acta Psychologica*, 60, 193-210.
- Payne, J., Bettman, J. R. & Johnson, E. J. (1993). *The adaptive decision maker*. Cambridge: Cambridge University Press.
- Quilici, J. L. & Mayer, R. E. (1996). Role of examples in how students learn to categorize statistics word problems. *Journal of Educational Psychology*, 88, 144-161.
- Reder, L. M. & Schunn, D. (1998). Bringing together the psychometric and strategy worlds: Predicting adaptivity in a dynamic task. In D. Gopher & A. Koriati (Eds), *Cognitive regulation of performance (Attention and performance XVII)*. Cambridge, MA: MIT Press.
- Reed, S. K. (1999). *Word problems*. Mahwah, NJ: Erlbaum.
- Simon, H. A. (1990). Invariants of human behavior. *Annual Review of Psychology*, 41, 1-19.
- VanLehn, K. & Jones, R. M. (1993). Better learners use analogical problem solving sparingly. In P. E. Utgoff (Ed.), *Machine Learning: Proceedings of the Tenth Annual Conference*. San Mateo, CA: Morgan Kaufmann.