Accelerated Future Learning via Explicit Instruction of a Problem Solving Strategy

Min CHI and Kurt VANLEHN
Learning Research and Development Center
University of Pittsburgh, PA, 15260

Abstract. Explicit instruction in a problem-solving strategy accelerated learning not only in the domain where it was taught but also in a second domain where it was not taught. We present data from a study in which students learned two unrelated deductive domains: probability and physics. During the probability instruction, the Strategy group was trained with an Intelligent Tutoring System (ITS) that explicitly taught a domain-independent backward chaining problem-solving strategy while the No-strategy groups trained with another ITS without any explicit strategy instruction. During the subsequent physics instruction, both groups were trained with the same ITS, which did not explicitly teach any strategy. The Strategy group gained significantly more than the No-strategy group in both domains. Moreover, their gains were evident both on multiple-principle problems, where the strategy should make problem solving more efficient, and on single-principle ones, where the strategy should make no difference. This suggests that the strategy increased the students' learning of domain principles and concepts, because that is all they had to know in order to solve the single-principle problems.

Keywords. Inter-domain transfer, Intelligent Tutoring Systems, Problem-solving Strategy

Introduction

One of the most important goals of learning is transfer, so transfer should also be at the core of AI in Education. This paper focuses on inter-domain transfer. That is, if students first learn one task domain, X, does that improve their learning of a second task domain, Y? For the remainder of this paper we will use the term "transfer" for inter-domain transfer only.

We define three types of transfer based upon how the benefits are measured in the second domain Y; *Direct Application*, *Savings*, and *Accelerated Future Learning*. For the purposes of clarity, we will use the term X students for students who studied domain X first and non-X students for those who did not.

Direct Application occurs if studying domain X increases students' competence to solving problems in domain Y. The X students receive instruction in domain X and then are tested in domain Y with no additional instruction. If they score higher on the Y-test than the non-X students, then Direct Application has occurred. For instance, Bassok and Holyoak [1] investigated transfer across two isomorphic domains, arithmetic-progression word problems in algebra and physics problems involving motion in a straight line with constant acceleration. They found that students who learned the algebra problems more likely spontaneously recognized that physics problems can be solved by the same equation. However, students who learned the physics topics exhibited little detectable transfer to the isomorphic algebra ones.

In both *Savings* and *Accelerated Future Learning*, the second domain Y is taught, not just tested. *Savings* occurs when training on X gives students a head start in learning Y even though both groups learned Y at the same rate. In Singley and Anderson's [6] study, the X and Y domains were the word processors WordPerfect and Word respectively, which share many operations and concepts. Students who had mastered WordPerfect, the X students, had a higher initial competence than novices with no prior word-processor experience, the non-X ones. Although both groups learned how to use Word at roughly the same rate, the X students had fewer elements to learn and thus took less time to master Word than the non-X ones. Both Direct Application and Savings occur when the domains share identical elements [6][7].

Accelerated Future Learning (AFL) occurs when both groups start with the same initial competence in domain Y, but the X group learns domain Y at a faster *rate* than the non-X group. Slotta and Chi [7] taught half their students, the X students, a domain-independent schema for understanding certain kinds of scientific phenomena. Then both X and non-X students studied Y, a physics text about electricity. Although the X training included no mention of electricity, the X students gained more than the non-X students when studying Y.

In some studies of transfer, the independent variable is not the presence or absence of instruction on domain X, but the different teaching method used. Schwartz et al.[5] compared inquiry methods to didactic methods for teaching domain X, and assessed the students learning in domain Y. Students who learned X via inquiry learned Y more effectively than students who learned X didactically. Unlike the other studies cited here, X and Y were different problems in the same task domain; however their study illustrates a type of AFL that is more similar to the type we observed in this study.

In this study, both X and Y are deductive domains. Like Schwartz et al.[5], we compared two different methods for teaching X. Half of the students studied X with explicit problem-solving strategy instruction while the other half studied it without such explicit instruction. Then both groups received the same instruction in the Y domain. Our primary research question is: for deductive task domains, can explicit instruction on problem solving strategies cause Accelerated Future Learning?

1. Problem solving strategies in Deductive domains

A task domain is *deductive* if solving a problem in it means producing a proof or derivation consisting of one or more inference steps each of which is the result of applying a general domain principle, operator, or rule. Deductive task domains are common parts of mathematical and scientific courses.

Broadly speaking, a *problem-solving strategy* is a policy or procedure that determines how a problem could be solved. In deductive task domains, two common strategies are *forward chaining* and *backward chaining*. Forward Chaining starts with a set of known propositions, applies a principle to some subset of them which produces at least one new known proposition. This process repeats until the problem is solved. The reasoning proceeds forwards from the known propositions toward the goal. Backward chaining is *goal-directed*. It works backward from a goal state to the known propositions by repeatedly applying deductive rules that infer the current goal from some set of subgoals. The goal is progressively broken down into subgoals until the known propositions are reached, which results in a partial solution plan. This plan is followed until more planning is necessary, and then the backward chaining is repeated.

Neither forward nor backward chaining are typically observed in a pure form in natural human problem solving [2][3] and neither one is generally taught to human problem solvers. Yet their success in guiding computer problem solvers suggests that explicitly teaching them might improve students' problem solving. Although there have been several tests of this hypothesis, they have been complicated by many problems including: confounded designs [8][9], null effects [10], teaching suboptimal strategies [4], and floor effects [11]. Despite all the work, we still lack an unequivocal answer to this simple question: given that some problem-solving strategies work so well for computer problem solving, should we teach them to human students? Moreover, would the strategy be domain-independent in that it would apply to others domains even though it was learned in just the initial one?

2. Methods

In this study, the students were divided into two groups. The Strategy group was explicitly taught the Target Variable Strategy, a domain-independent Backward Chaining problem-solving strategy [11]; the No-strategy group received no explicit instruction and was free to solve problems in any fashion. Two deductive task domains, probability and physics, were taught. Each domain contained ten major principles. In the first domain, probability, the Strategy students were required to study and use the Target Variable Strategy while the No-strategy were not required to use any specific strategy. In the second domain, physics, all students were taught in the same way and were not required to use any specific strategy.

2.1. Participants:

Data was collected over a period of four months during the fall of 2005 and early spring of 2006. We recruited 91 undergraduate students. They were required to have a basic understanding of high-school algebra, but not to have taken college-level statistics or physics courses. They were randomly assigned to the two groups. Each took from two to three weeks to complete the study over multiple sessions. Due to the winter break and length of the experiment, only 44 completed the experiment. Two students were dropped one scored perfectly on the probability pre-test and the other lacked time consistency. Of the remaining 42 participants (59.5% female), 20 were assigned to the Strategy group and 22 to the No-strategy group.

2.2. Three ITSs

Three ITSs were used in our study: Pyrenees-probability, Andes-probability, and Andes-physics. Apart from the declarative knowledge, Andes-probability and Andes-physics were identical, and we will use the term Andes to refer to them both. All three employed a multi-paned user interface that consisted of a problem statement window, a variable window, an equation window, and a dialog window. In Pyrenees, students interacted with the tutor entirely via a text-based menu and all interactions occur in the dialog window. Andes is multi-modal and the students interacted with it in all four windows using the mouse and keyboard entry.

The salient difference between Pyrenees and Andes is that Pyrenees explicitly taught the Target Variable Strategy and required students to follow it. Andes provided no explicit strategic instruction nor did it require students to follow any particular strategy. Students using Andes could input any entry, and Andes colored it green if it was correct and red if incorrect. Students could enter an equation that was the algebraic combination of several principle applications on Andes but not on Pyrenees.

Besides providing immediate feedback, both Pyrenees and Andes provided help when asked. When an entry was incorrect, students could either fix it on their own or ask for *what's-wrong help*. When they did not know what to do next, they could ask for *next-step help*. Both Pyrenees and Andes gave similar *what's-wrong help* based upon the domain knowledge but the *next-step help* differed. Because Pyrenees required students to follow the Target Variable Strategy, it knew exactly what step the student should be doing next so it gave specific hints. In Andes, on the other hand, students could enter correct equations in any order, and an equation was considered correct if it was true, regardless of whether it was useful for solving the problem. So Andes did not attempt to figure out the student's problem-solving plans or intentions. Instead, it picked a step that it would most like to do next, and hinted that step. Both *next-step help* and *what's-wrong help* were provided via a sequence of increasingly specific hint. The last hint in the sequence, the *bottom-out hint*, told the student exactly what to do.

In this study, the Strategy students learned probability on Pyrenees and then physics on Andes-physics; while the No-strategy students learned both probability and physics on Andes.

2.3. Procedure

This study had 4 main segments: background survey, probability instruction, Andes interface training, and physics instruction (see Table 1, left column). The background survey included high school GPA, Math SAT, Verbal SAT score, and experience with algebra, probability, and physics, and other information.

Table 1: Procedure

Segment	Strategy group	No-strategy group
Background Survey	Background survey	
Probability instruction	Probability pre-training	
	Probability pre-test	
	Pyrenees video	Andes-Probability video
	Probability Training on Pyrenees	Probability Training on Andes-Prob
	Probability post-test	
Andes interface training	Andes-Probability video	
	Solve a problem with Andes-Prob.	
Physics instruction	Physics pre-training	
	Physics pre-test	
	Andes-Physics video	
	Physics Training on Andes-Physics	
	Physics Post-test	

Both the probability and physics instruction consisted of the same five phases: 1) pre-training 2) pre-test, 3) watching the video, 4) training on the ITS, and 5) post-test. During phase 1, pre-training, the students studied the domain principles. For each principle, they read a general description, reviewed some examples, and solved some single and multiple-principle problems. After solving a problem, the student's answer was marked in green if it was correct and red if incorrect, and the expert's solution was also displayed. If the students failed to solve a single-principle problem then they were asked to solve an isomorphic one; this process repeated until they either failed three times or succeed once. The students had only one chance to solve each multiple-principle problem and were not asked to solve an isomorphic problem if their answer is incorrect. During phase 2, students took the pre-test. They were not given feedback on their answer nor were they allowed to go back to earlier questions, (this was also true of the post-tests). During phase 3, they watched a video demo of a problem solving in the corresponding ITS. During probability instruction, the Strategy students also read a text description of the Target Variable Strategy.

In phase 4, both groups solved the same probability problems or physics problems in the same order on the corresponding ITS. Each main domain principle was applied at least twice and the textbook was available for students to review. During probability instruction, the Strategy students were also able to review the textual description of the Target Variable Strategy. Finally, in phase 5 students took the post-test. In each post-test, five of the multiple-principle problems were isomorphic to training problems in phase 4. The other post-test problems (5 for probability; 8 for physics) were novel, non-isomorphic multiple-principle problems. Most of the multiple-principle problems had dead-end search paths so that the Target Variable Strategy could show an advantage in search efficiency.

Only the Strategy students took the third segment, Andes interface training. Its purpose was to familiarize them with Andes without introducing any new domain knowledge. After watching the Andes-Probability video that was given to the No-strategy students during probability instruction, the Strategy students solved a problem in Andes-probability. The problem was one of the twelve problems that they had previously solved in Pyrenees. Our pilot studies showed that solving one problem was enough for most students to become familiar with the Andes interface.

The two key procedural differences between the conditions were: 1) during the probability instruction, the Strategy trained on Pyrenees while the No-strategy trained on Andes, and 2) the Strategy students spent additional time studying the Andes interface before studying physics.

2.4. Scoring Criteria for pre-trainings, pre-tests, and post-tests

Test problems required students to derive an answer by writing and solving one or more equations. We used three scoring rubrics: *binary*, *partial credit*, and *one-point-per-principle*. Under the *binary* rubric, a solution was worth 1 point if it was completely correct or 0 if not. Under the *partial credit* rubric, each problem score was the proportion of correct principle applications evident in the solution -- a student who correctly applied 4 of 5 possible principles would get a score of 0.8. Under the One-*point-per-principle* rubric gave a point for each correct principle application. Solutions were scored by a single grader in a double-blind manner.

3. Results

Several measures showed that the incoming students' competence was balanced across conditions. There was no significant difference between groups on (1) the background survey, (2) the probability pre-test with respect to their scores on three types of problems, single-principle, multiple-principle, and overall, across all 3 scoring rubrics, or (3) the probability pre-training on all three types of problems. Thus, despite attrition, the conditions remained balanced in terms of incoming competence.

Results also showed that the two conditions did not differ on the four training times: (1) probability pre-training; (2) probability training on Pyrenees or Andes-probability, (3) physics pre-training, and (4) physics training on Andes-physics. This is fortuitous, as it implies that any difference in post-test scores is due to the effectiveness of the instruction rather than time-on-task variations.

The main outcome (dependent) variables are the students' scores on the probability post-test, the physics pre-training, pre-test, and post-test. We discuss each in turn.

On the probability post-test using the binary scoring rubric, the Strategy students scored significantly higher than the No-strategy ones: $t(40)=3.765$; $p=0.001$. The effect size was 1.17, and is denoted "d" subsequently. Moreover, they also scored higher than the No-strategy students on single-principle problems, $t(40)=3.960$; $p<0.001$; $d=1.24$, and multiple-principle ones, $t(40)=2.829$; $p=0.007$; $d=0.87$. The same pattern was found with under the partial credit and one-point-per-principle scoring rubrics.

During physics pre-training under the binary scoring rubric, the Strategy students solved more single-principle problems correctly in the *first* try than the No-strategy ones: $t(40)=2.072$, $p=0.045$; $d=0.64$. No significant difference was found between the two groups on the total number of simple- and multiple-principle problems solved correctly. This finding could be due to an unlucky random assignment: the Strategy students may have had higher prior physics knowledge than the NS. Although we cannot rule this interpretation out, there was no difference across the two conditions on the background questionnaire items relating to the students' physics experience and grades. Therefore, it is more likely that the Strategy students did, even at this early state, learn physics faster than the No-strategy students. This explanation is consistent with the results presented next.

On the physics pre-test, the Strategy students scored higher than the No-strategy under all three scoring rubrics. Under the binary scoring rubric, $t(40)=2.217$, $p=0.032$, $d=0.69$. On the single-principle problems, the two conditions did not differ significantly regardless of the scoring rubric. This was probably due to a ceiling effect. On the multiple-principle problems, the Strategy students scored higher than the No-strategy ones under the partial credit rubric, $t(40)=2.913$, $p=0.0058$, $d=0.90$ and one-point-per-principle rubric $t(40)=.800$, $p=0.008$, $d=0.86$, but not on the binary rubric $t(40)=1.148$, $p=0.147$. This could be due to the binary rubric's lack of sensitivity. In any case, the overall physics pretest scores indicate that the Strategy students learned more effectively during the physics pre-training than the No-strategy students.

On the physics post-test, the Strategy students scored much higher than the No-strategy students on the single-principle and multiple-principle problems, and overall under all three scoring rubrics. Under the binary rubric, $t(40)=4.130$, $p<0.0002$, $d=1.28$; on overall problems, $t(40)=3.211$, $p=0.003$, $d=1.00$ on single-principle problems and $t(40)=3.395$, $p<0.001$, $d=1.23$ on multiple-principle ones.

The results presented so far are consistent with two kinds of transfer: *Savings*: the Strategy students had a head start over the No-strategy ones on learning physics; or *Acceleration of Future Learning*: the Strategy students may have learned physics faster than the No-strategy ones. In order to differentiate between a head-start and a faster learning rate, we ran an ANCOVA on the physics post-test scores using the physics pre-test scores as a covariant. This yielded a post-test score for each student, adjusted for the difference in his/her physics pre-test score. Under binary scoring rubric, the Strategy students had higher adjusted post-test scores ($M=0.705$, $SD=0.21$) than the No-strategy students ($M=0.478$, $SD=0.22$). This difference was large and reliable $F(39)=11.079$, $p=0.002$, $d=1.05$. A similar pattern held for the other two rubrics. This suggests that learning the strategy in probability did accelerate the students' learning of physics. This is intuitively satisfying, as we chose task domains that had very little overlap, making a Savings-style transfer unlikely.

The results are summarized in Table 2. Numbers in the cells are effect sizes. An *ms* or *ns* indicates that the difference between the two conditions was marginally significant or not-significant respectively. On the measures labeled ns*; the non-significance was probably due to a ceiling effect. In sum, the Strategy and No-strategy students tied on the measure of prior knowledge (the probability pre-training and pre-test), and the Strategy students scored higher than the No-strategy ones at all the subsequent assessment points: probability post-test, physics pre-test and post-test.

Table 2: Effect sizes: Strategy compared to No-strategy

scoring rubrics			Single-	Multiple-	Overall
binary	Probability	Post-test	1.24	0.87	1.17
	Physics	Pre-test	*ms*	*Ns*	0.69
		Post-test	1.00	1.23	1.28
		Adjusted Post-test	*ns**	1.02	1.05
partial credit	Probability	Post-test	1.24	1.04	1.27
	Physics	Pre-test	*ns*	0.90	0.93
		Post-test	0.92	1.17	1.21
		Adjusted Post-test	*ns**	0.78	0.81
one-point-per-principle	Probability	Post-test	1.24	0.99	1.07
	Physics	Pre-test	*ms*	0.86	0.91
		Post-test	0.67	1.17	1.18
		Adjusted Post-test	*ns**	0.73	0.75

4. Discussion

Consistent with our previous study of physics [11], the Target Variable Strategy improved students' learning significantly in the initial domain, probability. Because the Target Variable Strategy increased learning in two deductive domains, it may be an effective strategy in other deductive domains as well. More importantly, we found that teaching students the Target Variable Strategy in probability significantly accelerated their learning in a new domain, physics, even though it was not taught there. Let us consider three hypotheses about *why* teaching the Target Variable Strategy improved learning in both domains.

The first hypothesis is that teaching students the Target Variable Strategy improved their search efficiency. The multiple-principle problems were constructed so that using the strategy would reduce the average number of steps required to solve the problems, and hence reduce both time and the likelihood of error. If the search-efficiency hypothesis was true, the Strategy students should perform better than the No-strategy ones on the multiple-principle problems; while on the single-principle problems, where search is not required, the two groups should perform. However, this latter prediction did not pan out (see Table 3). The Strategy students outscored the No-strategy ones on single-principle post-test problems in both probability and physics.

The second hypothesis is that it improved students' schema learning. A schema is defined as a structure that allows problem solvers to recognize a problem as belonging to a particular category of problems that normally require particular moves. If teaching students the Target Variable Strategy improved their schema acquisition, then the Strategy students should have performed better than the No-strategy ones on the problems that are isomorphic to training problems; on the novel, non-isomorphic test problems, however, the two groups should perform equally. This latter prediction also turned out to be false. On the non-isomorphic multiple-principle problems (5 in probability post-test and 8 in physics post-test), the Strategy students performed significant better than the No-strategy ones: $t(40)= 2.27$, $p= 0.029$ on the probability post-test and $t(40)= 3.803$, $p< 0.0005$ on the physics post-test respectively.

The third hypothesis is that teaching the strategy improves students' learning of the domain concepts and principles. Although this may seem implausible, the Target Variable Strategy requires students to apply one principle at a time (by selecting the desired principle from a menu), whereas Andes did not force them to identify the principles nor to apply them one at a time. Therefore, if students made a mistake, they would get principle-specific feedback in Pyrenees but not in Andes. Such focus on principles during probability instruction may have convinced the Strategy students to continue to focus on principles during physics instruction. If this hypothesis is true, we expect that on all types of problems, the Strategy group should perform better than the No-strategy one. This turned out to be the case. Therefore, it suggests that the main effect of teaching the strategy was to get students to focus on the domain principles in both domains.

In sum, we have shown that teaching students an explicit problem-solving strategy improved their performance in the initial domain and more importantly, it seems to cause Accelerated Future Learning in a new domain. Because the improvement occurred with all types of problems, it seems likely that the strategy instruction accelerated the learning of domain principles, as opposed to accelerating the learning of problem schemas or improving the search efficiency. It is rare to observe inter-domain transfer at all, especially one with such a large effect size. Acceleration of Future Learning is similarly uncommon as is getting the students to focus not on the whole problem solution but on each principle individually. Given that this hypothesized shift in learning orientation was obtained entirely with text and ITS, it may be possible to bring it to scale quickly, with large benefits in many deductive domains.

5. References

[1] Bassok M, & Holyoak K 1989. Interdomain Transfer between Isomorphic Topics in Algebra and physics. Journal of Experimental Psychology: Learning, Memory, and Cognition, 15, 1 153-166.

[2] Newell, A. and Simon, H. (1972). Human Problem Solving. Englewood Cliffs, NJ: Prentice Hall.

[3] Priest, A. G., and R. O. Lindsay (1992). "New Light On Novice-Expert Differences in Physics Problem Solving," British Journal of Psychology, 83, 389-405.

[4] Scheines, R & Sieg, W 1994.Computer environments for proof construction. Interactive Learning Environments, 4(2), 159-169.

[5] Schwartz, D., Bransford, J. et al. (2005). Efficiency and innovation in transfer. In J. Mestre (Ed.), Transfer of learning from a modern multidisciplinary perspective. CT: Information Age Publishing.

[6] Singley, K. & Anderson, R. 1989. The transfer of cognitive skill. Cambridge, MA: Harvard Press.

[7] Slotta, J. D. & Chi, M.T.H. (2006). Students Understand Challenging Topics in Science Through Ontology Training. Cognition & Instruction, 2006, Vol. 24 Issue 2, p261-289.

[8] Sweller 1988 Cognitive load during problem solving Effect on Learning. Cognitive Science12 257-285

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

[10] Trafton, J, & Reiser, B. (1991). Providing natural representations to facilitate novices' understanding in a new domain: Forward and backward reasoning in programming. *The Proceedings of the 13th Annual Conference of the Cognitive Science Society* (pp. 923-927). Hillsdale, NJ: Erlbaum.

[11] VanLehn, K., Bhembe, D., Chi, M. et al. 2004. Implicit versus explicit learning of strategies in a non-procedural cognitive skill. 7th Conference on Intelligent Tutoring Systems, Maceio, Brazil