Preparing Unprepared Students For Future Learning

Mark Abdelshiheed, Mehak Maniktala, Song Ju, Ayush Jain, Tiffany Barnes, and Min Chi
Department of Computer Science
North Carolina State University
Raleigh, NC 27695
{mnabdels, mmanikt, sju2, ajain37, tmbarnes, mchi}@ncsu.edu

Abstract

Based on *strategy-awareness* (knowing *which* problem-solving strategy to use) and *time-awareness* (knowing *when* to use it), students are categorized into *Rote* (neither type of awareness), *Dabbler* (strategy-aware only) or *Selective* (both types of awareness). It was shown that Selective is often significantly more prepared for future learning than Rote and Dabbler (Abdelshiheed et al., 2020). In this work, we explore the impact of *explicit strategy instruction* on Rote and Dabbler students across two domains: logic and probability. During the logic instruction, our logic tutor handles both Forward-Chaining (FC) and Backward-Chaining (BC) strategies, with FC being the default; the Experimental condition is taught *how to use BC via worked examples* and *when to use it via prompts*. Six weeks later, all students are trained on a probability tutor that supports BC only. Our results show that Experimental significantly outperforms Control in both domains, and Experimental Rote catches up with Selective.

Keywords: Preparation For Future Learning; Metacognitive Skills Instruction; Backward Chaining

Introduction

It is commonly observed that certain learners are better prepared for future learning in that they are less sensitive to learning environments and can always learn; while others are more sensitive to variations in learning environments and may fail to learn (Ju, Zhou, Abdelshiheed, Barnes, & Chi, 2021; Sanz Ausin et al., 2019; Zhou & Chi, 2017; Chi & VanLehn, 2010; Cronbach & Snow, 1977). It is not fully understood why such differences exist and one of many hypotheses is that certain learners lack specific skills such as general problem-solving strategies or metacognitive skills (Abdelshiheed et al., 2020; Pintrich, 2002; Hartman, 2001).

In this work, we focus on two types of metacognitive skills related to problem-solving strategies: *strategy-awareness* and *time-awareness*, that is, knowing *which strategy* to use and *when* to use it. Based on these two types, we classify students into three groups: "***Rote***" referring to those who never switch problem-solving strategies and simply follow the default strategy, "***Dabbler***" for those who seemingly know different problem-solving strategies but not proficient enough to know *when* to use each, and finally, "***Selective***" refers to students who know both *which* strategy to use and *when* to use it. In our prior work, we analyzed data from 495 college students who studied logic and then probability using two intelligent tutoring systems (ITSs) and we found that Selective outperformed both Rote and Dabbler across both tutors (Abdelshiheed et al., 2020). These findings were consistent with prior research, in that skilled problem solvers tend to understand how, when, and why to use each strategy (de Boer et al., 2018; Pintrich, 2002; Alexander et al., 1998). In this work, we explore whether *explicit instruction on **how** and **when** to use a problem-solving strategy* in one domain can make Rote and Dabbler catch up with Selective using two intelligent tutoring systems (ITSs) across two deductive domains: logic and probability.

In deductive task domains, solving a problem often requires deriving an argument or proof consisting of one or more inferential steps, and each step is the result of applying a domain principle, operator, or rule. Two common problem-solving strategies in deductive domains are forward chaining (FC) and backward chaining (BC) (Russell & Norvig, 2010). Prior work has shown that students often use a mixture of FC and BC during their problem-solving (Priest & Lindsay, 1992; Simon & Simon, 1978). In this work, our logic tutor handles both FC and BC strategies but defaults to FC, while our probability tutor supports BC only. During the logic instruction, the Experimental condition is explicitly taught BC strategy through *worked examples* (strategy-awareness) and the tutor would *prompt* them to switch to it when it is proper to do so (time-awareness), while the Control condition is trained on the same tutor without any explicit strategy instruction. After six weeks, both conditions are trained on a probability tutor. Our results show that the Experimental condition significantly outperforms Control in both domains. More importantly, the Rote students benefit the most from our intervention, as they significantly outperform their peers on the logic tutor and continue to perform well in the probability tutor. By analyzing their strategy switch behavior, we find that they indeed catch up with their Selective peers, who know in advance which strategy to use when.

Related Work

Generally speaking, metacognition indicates one's ability to regulate, understand and monitor their cognitive skills (Chambres et al., 2002; Roberts & Erdos, 1993). Many studies have investigated the impact of teaching metacognitive skills on student learning (Zepeda et al., 2015; Chi & VanLehn, 2010) and explored different ways of teaching them effectively (de Boer et al., 2018; Cardelle-Elawar, 1992). Based on Winne and Azevedo (2014), mastering strategy selection alone is a cognitive skill, but when incorporated with awareness about when such strategy should be used, it can be considered as a metacognitive skill. Thus in this work, we consider two types of awareness: *strategy-awareness* and *time-awareness*, which are respectively, which strategy to use and when to use it (de Boer et al., 2018).

Strategy Instruction

Researchers have explored a variety of instruction types for teaching strategies (de Boer et al., 2018; Likourezos & Kalyuga, 2017; Zepeda et al., 2015; Spörer et al., 2009). For instance, Spörer et al. (2009) analyzed the role of explicit instruction of multiple reading strategies on the comprehension of third- to sixth-graders. They found that students who were instructed explicitly outperformed their peers, who were taught traditionally by the instructors' text interactions, on a transfer posttest task and follow-up test.

In this work, we focus on two problem-solving strategies: FC and BC. In the former, students progress from givens to the goal, while in the latter, they reduce the goal to givens. Prior research has highlighted the impact of teaching *FC* (Shrestha et al., 2013), *BC* (Chi & VanLehn, 2010), or both strategies (Priest & Lindsay, 1992; Larkin et al., 1980; Simon & Simon, 1978). For example, Priest and Lindsay (1992) compared how experts and novices solve physics problems. They found that both groups used a similar mixture of forward and backward chaining. However, *only* the experts knew when and why to use each strategy, as they significantly produced more complete plans and stages compared to their novice peers. As far as we know, prior work taught FC or BC by asking students to solve problems with/without feedback or support. In this work, we use *worked examples* to teach students BC in the logic tutor, where students are allowed to choose either FC or BC to solve each problem.

The Worked-Example Effect

A worked example (WE) is a provided step-by-step solution that solves a problem or completes a task. In cognitive load theory, the *worked-example effect* refers to the observed learning outcome from teaching with WEs (Sweller & Cooper, 1985). Substantial work has leveraged WEs to enhance students' problem-solving skills (Renkl & Atkinson, 2010; Van Gog et al., 2004; Paas et al., 1994), prepare them for direct instruction (Likourezos & Kalyuga, 2017), seal knowledge-gap experience (Glogger-Frey et al., 2015), reduce cognitive load (Gerjets et al., 2004), and deepen the conceptual understanding (Schwonke et al., 2009).

Renkl and Atkinson (2010) addressed when to present WEs to students. They found that WEs contribute positively to the learning performance when offered in the *early* stages of skill acquisition. However, students will likely stop paying attention to them in later stages. The authors argued that WEs should elicit some steps from students, and possess self-explanatory prompts to strengthen their impacts. Furthermore, Schwonke et al. (2009) investigated the impact of faded worked examples (FWEs), which is a mixture of solution steps and incomplete steps to be solved by students. They found that students who were given FWEs were able to learn more efficiently and show deeper conceptual knowledge of geometry principles, compared to their peers who had supportive tutoring in the form of corrective feedback and self-explanation prompts.

In short, given that WEs have shown great promise of improving student content learning, in this work, we use WEs in the early stages of the logic tutor to teach the *BC* strategy.

The Importance of "When to use Which Strategy"

Prior work suggests that students often lack metacognitive knowledge because they do not understand *when or why* to use a specific strategy (de Boer et al., 2018; Pintrich, 2002; Hartman, 2001), and thus in order to raise their metacognitive awareness, they must master such skills (Winne & Azevedo, 2014; Alexander et al., 1998). For instance, de Boer et al. (2018) investigated the long-term effects of metacognitive strategy instruction on students' academic performance. They found that students who were given interventions that include when, why, how, and which strategy to use, showed superior planning skills. Moreover, their learning performance was the highest, as they did not only outperform their peers on a posttest task, but also a far follow-up test. de Boer et al. argued that in multi-strategy domains, it is insufficient to only learn *what* each strategy is. Rather, it is equally important to learn *when* to use each.

On the other hand, compared with a large amount of prior work on explicit strategy instruction, little research has been done on how to teach students to become time-aware. Some related work has shed light on this subject. For example, it has been shown that young adults (< 30 years) are more flexible in and capable of, revising their initial strategy choice for a better strategy (Taillan et al., 2015; Ardiale & Lemaire, 2012). More specifically, Taillan et al. (2015) compared the strategy switch behavior between young and older adults in a computational task and found that the former were more receptive to trying new strategies. Moreover, Taillan et al. argued that the likelihood of switching a strategy is much higher when all strategies have similar difficulty. Inspired by these findings, we use *prompts* to recommend students to switch their problem-solving strategy when it is proper to do so.

Methods

Two Tutors and Our Interventions

Logic Tutor and Instructional Interventions Our logic tutor teaches students propositional logic proofs. It consists of five ordered levels with an *incremental degree of difficulty* and each level consists of four problems. A student can solve any problem by either a **FC** or **BC** strategy. Figure 1a shows that for *FC*, one must derive the conclusion at the bottom from givens at the top; while Figure 1b shows that for *BC*, students need to derive a contradiction from givens and the *negation* of the conclusion. Problems are presented by *default* in FC, but students can switch to BC by clicking the

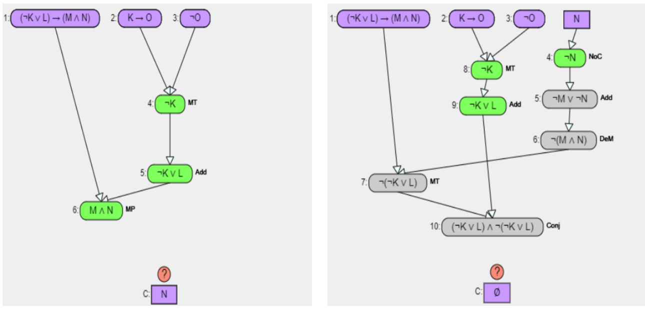

(a) Forward Chaining (b) Backward Chaining

Figure 1: Logic Tutor Problem-Solving Strategies

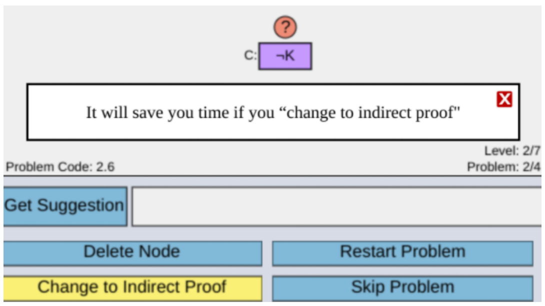

Figure 2: Strategy Switch Prompt

yellow button in Figure 2.

For the purpose of this work, we modified our logic tutor by implementing two WEs on explicit BC strategy instruction and by offering *advisory* prompts (the text in the black box in Figure 2) to switch to the BC strategy when it is proper to do so. Figure 3 shows that the two BC WEs are presented as the first problem at Levels 1 and 2; all problems in green are 'proper' problems to be solved by BC. We used the strategy switch behavior from our prior study (Abdelshiheed et al., 2020) to guide us *which 'proper'* problems to display the prompts in, by picking the most frequently switched problems, and *when* to display them, by learning a probability distribution of the duration lengths that students take before switching. More specifically, 55% of the time the tutor would wait for 1.5 minutes, 35% for 3 mins, and only 10% for 6 mins. For the remaining problems (colored in white in Figure 3), the tutor behaves the same as the original tutor.

Our goal in this work is to investigate whether explicit BC strategy instruction using WEs combined with prompts would make Rote and Dabbler catch up with Selective. We expect that the former two groups would benefit from our instructional interventions which are designed to scaffold the metacognitive skills that they lack. On the other hand, for the Selective group, we expect that providing them with additional scaffolding may interfere with their existing skills. Therefore, *only* the Experimental Rote and Dabbler groups will get the treatment shown in Figure 3, while the Control Rote and Dabbler groups and the Selective group will get no treatment.

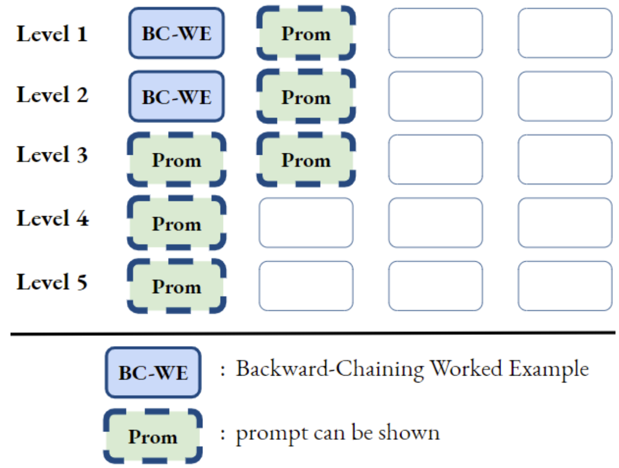

Figure 3: Training on the Modified Logic Tutor

Probability tutor: The interface of the tutor can be shown in Figure 4. It is a web- and text-based tutor that teaches students how to solve mathematical probability problems using 10 major principles, such as the Complement Theorem and Bayes' Rule. It consists of 12 problems. Each problem can *only* be solved by BC in that it requires students to derive an answer by *writing and solving equations* until the target is completely reduced to the givens.

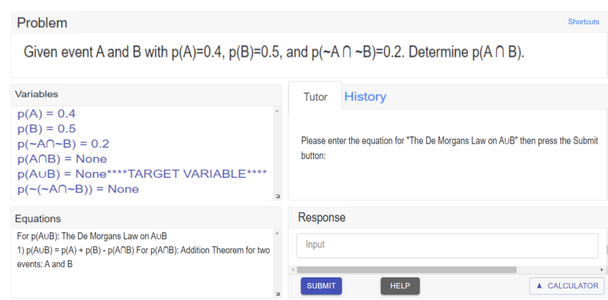

Figure 4: Probability Tutor Interface

Participants

Our participants are Computer Science undergraduate students at North Carolina State University. The tutors are assigned as a regular homework assignment and their completion is required for full credit. Students are told that the assignment would be graded based on their demonstrated effort, not performance. A total of 128 students completed both tutors.

One of the main challenges in this work is how to determine which of the three groups a student belongs to before their training on the logic tutor starts. Previously, the three metacognitive groups were determined based on a post-hoc analysis after students completed the entire training (Abdelshiheed et al., 2020). For the purpose of this work, we trained a random forest classifier (RFC) using prior

Table 1: Overview of the Study Procedure

Logic	Pretest (2 problems)
	Training (20 problems):
	Experimental $\Longrightarrow$ Intervention
	Control $\Longrightarrow$ Original
	Selective $\Longrightarrow$ Original
	Posttest (6 problems, including 2 isomorphic)
Six weeks later	
Prob.	Textbook
	Pretest (14 problems)
	Training (12 problems)
	Posttest (20 problems, including 14 isomorphic)

data for the task of the early prediction of the metacognitive group. More specifically, based on *how* students solve the logic pretest problems, our RFC will assign them to one of the three groups before training on the logic tutor. As a result, we had 60 Rote, 42 Dabbler and 26 Selective students.

Then the Rote and Dabbler students are randomly assigned to two conditions[1]: $N = 61$ for Experimental (35 Rote$_{Exp}$ + 26 Dabbler$_{Exp}$) and $N = 41$ for Control (25 Rote$_{Ctrl}$ + 16 Dabbler$_{Ctrl}$). We found no significant difference in the distribution of Rote and Dabbler across the two conditions: $\chi^2(1, N = 102) = 0.13, p = .72$. Finally, for all students who are classified as Selective, no intervention is provided and they use the original logic tutor and are referred to as the *Selective* group later on. The accuracy of our RFC is further confirmed by the students' behaviors using the Control and the Selective students since they received no intervention. Our results show that the RFC achieved 95.5% accuracy, 94.9% macro-recall, and 96.1% macro-precision which is comparable to its performance on the training data: 95.9% accuracy, 95.1% macro-recall, and 96.2% macro-precision.

Procedure

Table 1 summarizes our procedure. During the logic instruction, all students go through the same standard pretest-training-posttest procedure. The first two posttest problems are isomorphic to the two pretest problems. The *only* difference occurs during the training on the logic tutor in that the Experimental condition (Rote$_{Exp}$ and Dabbler$_{Exp}$) is given WEs with prompts (highlighted in Figure 3), while the Control condition (Rote$_{Ctrl}$ and Dabbler$_{Ctrl}$) and the Selective group receive no such intervention.

Six weeks later, students would be all trained on the same probability tutor following the same standard procedure: textbook, pretest, training on ITS, and posttest. In the textbook, students study the domain principles; In pre- and post-test, students solve 14 and 20 open-ended problems that require them to derive an answer by writing and solving one or more equations. 14 of the posttest problems are isomorphic to the pretest problems. For the training section, all students would go over the same problems in the same order. Note that for both tutors, the posttest is *much harder* than the pretest.

Grading criteria

In logic, a problem score is based on students' time, accuracy, and solution length. Students' *pre-* and *post-test* scores are calculated by averaging their pre- and post-test problem scores. In probability, students' answers are graded in a double-blind manner by experienced graders using a partial-credit rubric, where grades are based *only* on accuracy. The *pre-* and *post-test* scores are the average grades in their respective sections. For comparison purposes, all test scores are normalized to the range of $[0, 100]$.

Results

Learning Performance:

Experimental vs Control Table 2 compares the two conditions across the two tutors and it shows the mean(SD) of students' pre- and post-test scores, isomorphic posttest, training time (in hours), and the learning performance in terms of the normalized learning gain (NLG) defined as: $(NLG = \frac{post-pre}{\sqrt{100-pre}})$, where 100 is the maximum test score. The last column in Table 2 shows the one-way ANOVA statistical comparisons between the two conditions with their corresponding effective size η. As shown in Table 2, while no significant difference was found between the two conditions in the logic pretest: $F(1, 100) = 0.8, p = .38$, and probability pretest: $F(1, 100) = 2.7, p = .11$, Experimental significantly outperformed Control in all other aspects, except for the training time on the probability tutor.

Next, we will investigate whether both Rote and Dabbler

Table 2: Comparing the Two Conditions across Tutors

	Experimental ($N = 61$)	Control ($N = 41$)	1-way ANOVA
Logic Tutor			
Pre	62.8 (19.6)	59.4 (18.4)	$p = .38$
Iso Post	75.9 (2.5)	62.4 (4)	$F(1, 100) = 15.3, p <$ **.001**, $\eta = .13$
Iso NLG	0.19 (.03)	0.03 (.06)	$F(1, 100) = 8.2, p =$ **.005**, $\eta = .08$
Post	77.4 (3.6)	65.4 (5.2)	$F(1, 100) = 38.9, p <$ **.001**, $\eta = .28$
NLG	0.2 (.05)	0.05 (.06)	$F(1, 100) = 10.6, p =$ **.002**, $\eta = .1$
Time	4.32 (4.9)	7.23 (8)	$F(1, 100) = 5.2, p =$ **.02**, $\eta = .05$
Probability Tutor			
Pre	69.6 (18.6)	74.9 (15.1)	$p = .13$
Iso Post	92.9 (2.8)	84.4 (4.1)	$F(1, 100) = 18.5, p <$ **.001**, $\eta = .16$
Iso NLG	0.4 (.05)	0.11 (.08)	$F(1, 100) = 24.3, p <$ **.001**, $\eta = .2$
Post	88.9 (5)	72.2 (6.1)	$F(1, 100) = 58.2, p <$ **.001**, $\eta = .37$
NLG	0.33 (.06)	-0.18 (.21)	$F(1, 100) = 51.4, p <$ **.001**, $\eta = .34$
Time	1.86 (.5)	1.83 (.49)	$p = .76$

[1]The difference in size is due to the fact that we gave priority to have a sufficient number of Experimental students to perform a meaningful analysis of our intervention.

Table 3: Comparing The Groups' Mean(SD) Scores

Logic Tutor					
	Experimental		Control		
	$Rote_{Exp}$ ($N=35$)	$Dabbler_{Exp}$ ($N=26$)	$Rote_{Ctrl}$ ($N=25$)	$Dabbler_{Ctrl}$ ($N=16$)	Selective ($N=26$)
Pre	61.8 (23)	64.2 (14)	60.1 (20)	58.3 (16)	62.3 (21)
Iso Post	78.9 (1.9)	71.9 (1.6)	64.2 (3.8)	59.4 (4.2)	73.1 (5.3)
Iso NLG	0.25 (.04)	0.1 (.04)	0.05 (.05)	-0.02 (.06)	0.08 (.05)
Post	80.3 (1.7)	73.4 (1.5)	64.3 (3.5)	67.2 (2.9)	72.3 (5.5)
NLG	0.25 (.03)	0.13 (.03)	0.02 (.04)	0.09 (.07)	0.11 (.06)
Probability Tutor					
Pre	67 (20)	73.1 (16)	73.2 (15)	77.7 (15)	70.6 (19)
Iso Post	92.5 (3.4)	93.5 (3.3)	82.5 (3.9)	87.4 (5.8)	91.7 (6.2)
Iso NLG	0.43 (.06)	0.37 (.12)	0.09 (.21)	0.14 (.23)	0.37 (.16)
Post	88 (3.1)	90.2 (3.1)	71.3 (3.5)	73.5 (5.5)	85.8 (5.7)
NLG	0.35 (.05)	0.3 (.08)	-0.16 (.23)	-0.21 (.21)	0.24 (.15)

benefited from our intervention by comparing them across the two conditions. The first five columns in Table 3 compare the two Experimental groups (Rote$_{Exp}$ and Dabbler$_{Exp}$) against the two Control groups (Rote$_{Ctrl}$ and Dabbler$_{Ctrl}$). A two-way ANOVA using condition {Experimental, Control} and metacognitive group {Rote, Dabbler} as factors shows no significant difference among the four groups in the logic pretest: $F(1,98) = 0.28, p = .6$. A two-way ANCOVA using the logic pretest as a covariate, with condition and metacognitive group as factors, finds a significant interaction effect on the logic posttest: $F(1,97) = 17.3, p < .0001, \eta = 0.06$. There is also a main effect of condition: $F(1,97) = 66.7, p < .0001, \eta = 0.23$, as the two Experimental groups significantly outperform the two Control groups. Subsequent contrast analyses show that while no significant difference is found between the two Control groups, a significant difference is found between the two Experimental groups: Rote$_{Exp}$ > Dabbler$_{Exp}$ ($t(59) = 2.9, p < .01, d = 4.3$). A two-way ANOVA using the same two factors on NLGs shows similar results with the only exception that Dabbler$_{Exp}$ does not outperform the two Control groups: Rote$_{Ctrl}$ and Dabbler$_{Ctrl}$.

In probability, a two-way ANOVA using condition {Experimental, Control} and metacognitive group {Rote, Dabbler} as factors shows no significant difference among the four groups in the probability pretest: $F(1,98) = 0.05, p = .82$. Additionally, a two-way ANCOVA using the probability pretest as a covariate, with condition and metacognitive group as factors, shows no interaction effect in the posttest but a significant main effect of condition: $F(1,97) = 81.9, p < .0001, \eta = 0.43$ in that the two Experimental groups significantly outperform the two Control groups. Subsequent contrast analyses show that no significant difference is found neither between the two Experimental groups nor between the two Control groups. The same findings are found in probability NLGs.

In summary, our results show that our explicit strategy instruction using worked example combined with timely prompts is indeed effective in that Experimental not only outperforms the Control in logic where the intervention occurs but also the former continues to outperform the latter in probability six weeks later when there is no such intervention. More specifically, our results show that Rote students benefit more from our interventions in that $Rote_{Exp}$ has significantly higher posttest and NLG scores compared to the two control groups in logic, while $Dabbler_{Exp}$ did not learn significantly better than the two Control groups. For probability, both $Rote_{Exp}$ and $Dabbler_{Exp}$ significantly outperform the two Control groups in terms of posttest and NLG scores, and there is no significant difference between the two Experimental groups on the same measures.

Comparisons with the Selective Group It is essential to note that Selective is the "desired" group but only about 20% of the students are "naturally" classified as such in this work. To save space, the last column in Table 3 shows the performance of the Selective group across all measures. In this section, we will explore whether our instructional intervention is indeed effective from two aspects: 1) whether it would make the Rote and Dabbler students in the Experimental condition catch up with the Selective group; and 2) whether, without such intervention, the students in the Control condition would perform worse than the Selective group.

As for the first aspect, our overall results show that the Experimental condition performs as well as or better than the Selective group in that no significant difference is found between the two on all measures in both logic and probability. The only exception is that the former significantly outperforms the latter in the logic posttest: $t(36.3) = 2.1, p = .04$. Next, we compare the two Experimental groups, Rote$_{Exp}$ and Dabbler$_{Exp}$, against the Selective group. Table 3 shows that no significant difference is found among the three groups on both logic and probability pretests. While no significant difference is found between Selective and Dabbler$_{Exp}$ in logic, Rote$_{Exp}$ significantly outperforms Selective in the logic posttest: $t(59) = 3.2, p = .002, d = 2$ and the logic NLG: $t(59) = 2.3, p = .02, d = 2.9$. In probability, no significant difference is found among Rote$_{Exp}$, Dabbler$_{Exp}$, and Selective across all measures.

As for the second aspect, almost as expected, the Selective group outperforms the Control condition in both domains, except for the logic NLG; In the logic posttest: $t(53.9) = 2.4, p = .02$, in the probability posttest: $t(53.4) = 4.8, p < .001$ and $t(63.7) = 4.4, p < .001$ for the probability NLG. Similarly, when comparing the Control groups against the Selective group, Table 3 shows that: while no significant difference is found among the three groups on both logic and probability pretests, Selective significantly outperforms the two Control groups in both logic posttest: ($t(49) = 3, p < .01, d = 1.7$ for Rote$_{Ctrl}$ and $t(40) = 3.9, p < .001, d = 1.1$ for Dabbler$_{Ctrl}$) and probability posttest: ($t(49) = 4.7, p < .0001, d = 3.1$ for Rote$_{Ctrl}$ and $t(40) = 3.5, p < .001, d = 2.2$ for Dabbler$_{Ctrl}$).

To summarize, our results show that with our instructional

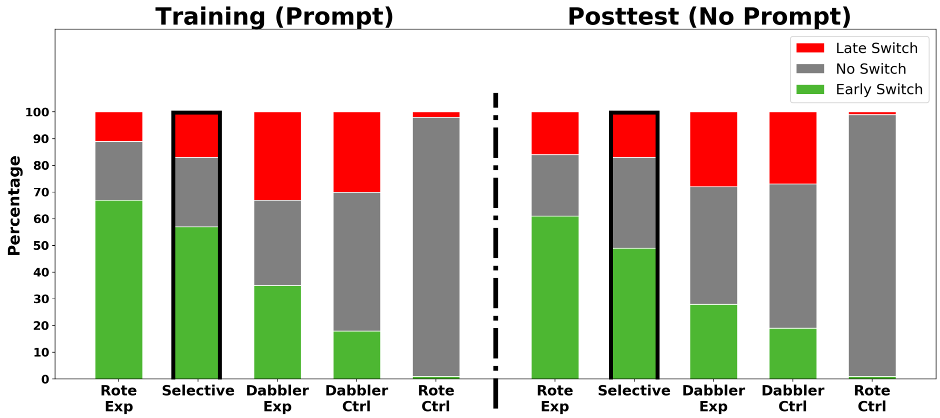

Figure 5: Strategy Switch Behavior in The Logic Tutor

intervention, Experimental indeed catches up with the Selective as the former performs at least as well as Selective on both tutors. On the other hand, without the interventions, Selective outperformed Control on the two tutors. More specifically, our results show that our interventions are especially beneficial to Rote students in that Rote$_{Exp}$ surpassed all other groups, even Selective, in logic and continued to perform well in probability. For Dabbler, Dabbler$_{Exp}$ outperformed the two control groups in the logic posttest, but not in logic NLG; In probability, Dabbler$_{Exp}$ outperformed the two control groups.

Strategy Switch Behavior

Figure 5 shows the strategy switch behaviors of the five groups (from FC into BC). We compared their behaviors during the logic tutor *Training* where only Rote$_{Exp}$ and Dabbler$_{Exp}$ were offered prompts, and *Posttest* where no student got prompts. Here, by following the same definition as described in our prior work (Abdelshiheed et al., 2020), we display the percentage of *No Switches* (sticking to the default strategy), *Early Switches* (switching within the first 30 actions), and *Late Switches* (switching after the first 30 actions). In Figure 5, the five groups are ordered by the percentage of *Early Switches*, the most desired behavior, and Selective is highlighted in bold as the gold standard.

A one-way ANOVA found that the switch behaviors differed significantly among the five groups: $F(4,123) = 71.2, p < .0001, \eta = 0.7$ for the training and $F(4,123) = 62.6, p < .0001, \eta = 0.67$ for the posttest. More importantly, the behaviors of each group were very similar in the training and posttest. Subsequent contrast analyses showed that while no significant difference was observed between Rote$_{Exp}$ and Selective on their switch behaviors, both groups switched early significantly more than the other three groups: Dabbler$_{Exp}$, Dabbler$_{Ctrl}$ and Rote$_{Ctrl}$. For example, Selective switched early significantly more than Dabbler$_{Exp}$: $t(50) = 5.4, p < .0001, d = 1.4$ for the training and $t(50) = 4.9, p < .0001, d = 1.3$ for the posttest.

In short, analyzing strategy switch behaviors confirms that Rote$_{Exp}$ indeed caught up with Selective, as the former showed very similar behaviors to the latter during the training when the prompts were available and more importantly, during the posttest when prompts were not present. On the other hand, much to our surprise, the strategy switch behaviors of Dabbler$_{Exp}$ stayed similar to their Control peers, Dabbler$_{Ctrl}$.

Conclusions & Discussions

Ever since the theory of Preparation for Future Learning (PFL) has been proposed (Bransford & Schwartz, 1999), substantial work has investigated how to assess PFL (Abdelshiheed et al., 2020; Mylopoulos et al., 2016; Belenky & Nokes-Malach, 2012) and factors that facilitate it such as metacognitive skills (Zepeda et al., 2015; Veenman et al., 2004), but relatively little research has been done on the design of interventions to prepare students for future learning (Zepeda et al., 2015). In this work, we investigate how explicit instruction on *how* and *when* to use a problem-solving strategy would impact student learning across two domains: logic and probability, and more importantly, whether such instruction would indeed eliminate the gap among different learners. Overall, our results show that teaching students the BC strategy using worked examples with prompts can significantly improve their learning not only in the domain they were taught but also in a new domain six weeks later. Specifically, such instructions seem to be more beneficial to students who were neither strategy-aware nor time-aware. Such findings suggest that explicit instructions on *how* and *when* can prepare students, especially those with little metacognitive skills, for future learning. Despite these findings, it is important to note that there is at least one caveat in our analyses: our probability tutor only supports the BC strategy. A more convincing testbed would be to use any ITS that supports different types of strategies so we can investigate whether Rote$_{Exp}$ would continue to switch.

Acknowledgments

This research was supported by the NSF Grants: Generalizing Data-Driven Technologies to Improve Individualized STEM

Instruction by Intelligent Tutors (2013502), Integrated Data-driven Technologies for Individualized Instruction in STEM Learning Environments (1726550) and CAREER: Improving Adaptive Decision Making in Interactive Learning Environments (1651909).

References

Abdelshiheed, M., et al. (2020). Metacognition and motivation: The role of time-awareness in preparation for future learning. In *Proceedings of the 42nd annual conference of the cognitive science society.*

Alexander, P. A., et al. (1998). A perspective on strategy research: Progress and prospects. *Educ. Psychol. Rev.*, *10*(2), 129–154.

Ardiale, E., & Lemaire, P. (2012). Within-item strategy switching: An age comparative study in adults. *Psychology and aging*, *27*(4), 1138.

Belenky, D., & Nokes-Malach, T. J. (2012). Motivation and transfer: The role of mastery-approach goals in preparation for future learning. *J. Learn. Sci.*, *21*(3), 399–432.

Bransford, J. D., & Schwartz, D. L. (1999). Rethinking transfer: A simple proposal with multiple implications. *Review of research in education*, *24*(1), 61–100.

Cardelle-Elawar, M. (1992). Effects of teaching metacognitive skills to students with low mathematics ability. *Teaching and teacher education*, *8*(2), 109–121.

Chambres, P., et al. (2002). *Metacognition: Process, function, and use*. Kluwer Academic Publishers.

Chi, M., & VanLehn, K. (2010). Meta-cognitive strategy instruction in intelligent tutoring systems: How, when, and why. *Educational Technology & Society*, *13*(1), 25–39.

Cronbach, L. J., & Snow, R. E. (1977). *Aptitudes and instructional methods: A handbook for research on interactions.*

de Boer, H., et al. (2018). Long-term effects of metacognitive strategy instruction on student academic performance: A meta-analysis. *Educational Research Review*, *24*, 98–115.

Gerjets, P., et al. (2004). Designing instructional examples to reduce intrinsic cognitive load: Molar versus modular presentation of solution procedures. *Instr. Sci.*, *32*, 33–58.

Glogger-Frey, I., et al. (2015). Inventing a solution and studying a worked solution prepare differently for learning from direct instruction. *Learning and Instruction*, *39*, 72–87.

Hartman, H. J. (2001). Developing students' metacognitive knowledge and skills. *Metacognition in learning and instruction*, 33–68.

Ju, S., Zhou, G., Abdelshiheed, M., Barnes, T., & Chi, M. (2021). Evaluating critical reinforcement learning framework in the field. In *International conference on artificial intelligence in education.*

Larkin, J., et al. (1980). Expert and novice performance in solving physics problems. *Science*, *208*, 1335–1342.

Likourezos, V., & Kalyuga, S. (2017). Instruction-first and problem-solving-first approaches: alternative pathways to learning complex tasks. *Instr. Sci.*, *45*, 195–219.

Mylopoulos, M., et al. (2016). Preparation for future learning: a missing competency in health professions education? *Medical education*, *50*(1), 115–123.

Paas, F. G., et al. (1994). Variability of worked examples and transfer of geometrical problem-solving skills: A cognitive-load approach. *J. Educ. Psychol.*, *86*(1), 122.

Pintrich, P. R. (2002). The role of metacognitive knowledge in learning, teaching, and assessing. *Theory into practice*, *41*(4), 219–225.

Priest, A., & Lindsay, R. (1992). New light on novice—expert differences in physics problem solving. *British journal of Psychology*, *83*(3), 389–405.

Renkl, A., & Atkinson, R. K. (2010). Learning from worked-out examples and problem solving.

Roberts, M. J., & Erdos, G. (1993). Strategy selection and metacognition. *Educational Psychology*, *13*, 259–266.

Russell, S. J., & Norvig, P. (2010). Artificial intelligence: a modern approach.

Sanz Ausin, M., et al. (2019). Leveraging deep reinforcement learning for pedagogical policy induction in an intelligent tutoring system. In *Proceedings of the 12th international conference on educational data mining.*

Schwonke, R., et al. (2009). The worked-example effect: Not an artefact of lousy control conditions. *Computers in human behavior*, *25*(2), 258–266.

Shrestha, A., et al. (2013). Using point-of-view video modeling and forward chaining to teach a functional self-help skill to a child with autism. *J. Behav. Educ.*, *22*, 157–167.

Simon, D. P., & Simon, H. A. (1978). Individual differences in solving physics problems.

Spörer, N., et al. (2009). Improving students' reading comprehension skills: Effects of strategy instruction and reciprocal teaching. *Learning and instruction*, *19*(3), 272–286.

Sweller, J., & Cooper, G. A. (1985). The use of worked examples as a substitute for problem solving in learning algebra. *Cognition and instruction*, *2*(1), 59–89.

Taillan, J., et al. (2015). Relationships between strategy switching and strategy switch costs in young and older adults: A study in arithmetic problem solving. *Experimental aging research*, *41*(2), 136–156.

Van Gog, T., et al. (2004). Process-oriented worked examples: Improving transfer performance through enhanced understanding. *Instr. Sci.*, *32*(1), 83–98.

Veenman, M. V., et al. (2004). The relation between intellectual and metacognitive skills from a developmental perspective. *Learning and instruction*, *14*(1), 89–109.

Winne, P. H., & Azevedo, R. (2014). Metacognition. In R. K. Sawyer (Ed.), *The cambridge handbook of the learning sciences* (2nd ed.). Cambridge University Press.

Zepeda, C. D., et al. (2015). Direct instruction of metacognition benefits adolescent science learning, transfer, and motivation: An in vivo study. *J. Educ. Psychol.*, *107*(4), 954.

Zhou, G., & Chi, M. (2017). The impact of decision agency & granularity on aptitude treatment interaction in tutoring. In *Cogsci.*