UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Stoic Behavior in Hint Seeking when Learning using an Intelligent Tutoring System

Permalink

Journal

ISSN

Authors

Miwa, Kazuhisa
Terai, Hitoshi
Kanzaki, Nana
et al.

Publication Date

Peer reviewed

Stoic Behavior in Hint Seeking
when Learning using an Intelligent Tutoring System

Kazuhisa Miwa (miwa@is.nagoya-u.ac.jp)
Hitoshi Terai (terai@is.nagoya-u.ac.jp)
Nana Kanzaki (kanzaki@cog.human.nagoya-u.ac.jp)
Graduate School of Information Science, Nagoya University

Ryuichi Nakaike (nakaike@educ.kyoto-u.ac.jp)
Graduate School of Education, Kyoto University

Abstract

Stoic behavior is defined as a behavior in which students tend not to seek help with a challenge. We investigated two types of stoic behavior: keeping-off behavior, in which students restrain themselves from requesting help, i.e., keep levels of help support at a minimum, and self-fading behavior, in which students voluntarily lower levels of support on their own volition. Three experiments were conducted. Overall, results showed that the participants actually exhibited stoic behavior when learning in an actual classroom setting. Self-fading was more difficult than the keeping-off behavior. The participants who maintained levels of support at a minimum through exhibiting active keeping-off behavior achieved greater learning gains, suggesting that stoic behavior resulted in positive impacts on learning. However, our experiment did not detect this effect for self-fading behavior. These experimental results were discussed with the assistance dilemma problem, generally occurring in instruction by intelligent tutoring systems.

Keywords: Intelligent tutoring system; Help seeking; Assistance dilemma.

Introduction

Recent intelligent tutoring systems include highly interactive features. Such systems give participants various types of feedback such as verification, correct response notification, try again encouragement, error flagging, and elaboration messages (Shute, 2008). In this context, the assistance dilemma has been recognized. Koedinger & Aleven (2007) asked a crucial question: How should learning environments balance assistance giving and withholding to achieve optimal learning? (Koedinger & Aleven, 2007) High assistance sometimes provides successful scaffolding and improves learning, however, it may also elicit superficial responses without consideration from the students. In contrast, low assistance sometimes encourages students to make great efforts to learn, while at other times it results in enormous errors and interferes with effective learning. To resolve this issue, the levels of support (LOS) must be adaptively controlled by tutoring systems.

Equipping tutoring systems with intelligent functions for help control and optimization of feedback information for participants are important issues. However, students must seek help intelligently. From this perspective, in this study, we investigate students' active help-seeking behaviors rather than passive help-receiving behaviors, which were managed by intelligent tutoring systems. Razzaq & Heffernan (2011) (Razzaq & Heffernan, 2010) confirmed that active hint-seeking, in which on-demand hints were given, was more effective for learning than passive hint-seeking, in which participants were given hints proactively when facing errors.

Help-seeking is a representative metacognitive activity in learning behavior. Help-seeking is valuable, not only for maximizing learning effects but also for acquiring a domain-independent meta-learning strategy. Some trials have instructed students using such metacognitive abilities. A domain-independent agent, called Help Tutor, for teaching better help-seeking skills by tracing students actions, was developed (Roll et al., 2006). Such help-seeking support was successful in improving students' declarative help-seeking knowledge, but did not improve their overall learning (Roll, Aleven, McLaren, & Koedinger, 2007). In more recent trials, the Help Tutor improved students' help-seeking behavior, and the improved help-seeking skills are transferred to learning new domain-level content during the month following the intervention (Roll, Aleven, McLaren, & Koedinger, 2011). To instruct such metacognitive activities, we must learn more about the nature of students' help-seeking behavior.

Students themselves have to manage their help-seeking behavior to maximize learning effects. However, previous studies have demonstrated that students' help-seeking behavior does not follow rational principles (Wood & Wood, 1999). Hint abuse is a representative irrational behavior that appears in hint-seeking whereby students tend to seek the most specific hints to find answers rather than acquiring understanding (Aleven & Koedinger, 2000).

In this study, we focus on stoic behavior in hint seeking. Stoic behavior is defined as behavior in which students tend not to seek help for their challenge. We will investigate two types of behavior: the keeping-off and self-fading behaviors. Keeping-off behavior is defined as behavior in which students try to solve problems by themselves without a system's assistance. Students restrain themselves from receiving help from a tutoring system even when permitted to do so. This is regarded as a type of behavior with the purpose of avoiding the hint abuse. On the other hand, we define self-fading behavior as one in which students voluntarily decrease an LOS by their own volition. This behavior is recognized, along with scaffolding, as a central concept for effective learning. To enable students to effectively learn, scaffolding should be eliminated gradually as learning progresses. In learning by examples, fading methods have also been used as an effective principle

for controlling the flow of learning (Atkinson, Renkl, & Merrill, 2003).

In this study, we investigate the following two research questions: Do students exhibit stoic behavior in hint seeking when learning with an intelligent tutoring system in classroom settings? and does such stoic behavior in hint seeking promote learning gains? We examined these research questions through three empirical studies.

Learning System and Task

We investigated participants' help-seeking behavior using a relatively complex learning task in which participants learned natural deduction (ND). Natural deduction is a kind of proof calculus in which logical reasoning is expressed by inference rules closely related to a natural way of reasoning. The following is an example solution process of an example problem: inducing a proposition $\neg Q \rightarrow \neg P$ from a premise $P \rightarrow Q$.

(1)	$P \rightarrow Q$	Premised
(2)	$\neg Q$	Assumption
(3)	P	Assumption
(4)	$P \rightarrow Q$	Reiteration of (1)
(5)	Q	$\rightarrow$ Elimination from (3) and (4)
(6)	$\neg Q$	Reiteration of (2)
(7)	$\neg P$	$\neg$ Introduction from (3), (5), and (6)
(8)	$\neg Q \rightarrow \neg P$	$\rightarrow$ Introduction from (2) and (7)

Participants learned inference rules and strategies for applying the rules. Participants in this study learned eight basic rules and four strategies, which are the fundamental basis of ND; the majority of problems can be solved using these rules and strategies.

Our tutoring system, which was developed for teaching ND to university undergraduates, has two important features.

First, it does not have a database that contains a set of ND problems and their solutions. Our system solves each problem on demand. It includes a production system model, which consists of the working memory, whose layout is consistent with the structure of ND problems, and production rules, which correspond to the inference rules and strategies for solving ND problems.

As a second feature, our system was established based on a server-client framework. Miwa et al. developed a web-based production system architecture called DoCoPro that enables such a system design to be established (Miwa, Morita, Nakaike, & Terai, 2013). A problem solver constructed on a server performed the complex inferences in ND. Client computers connected to the server performed easy processing for the interface. Using this server-client framework, our system can operate in any educational environment where various types of computers, e.g., high performance, poorly performing, and on different types of operating systems. Participant learning processes are saved as log data on the server.

Figure 1 shows a screenshot of the tutoring system. The system provides the participants with lists of the inference rules and strategies. They select one of the rules or strategies from a list, and the system automatically runs the rules and presents partial or complete results of inference. The system
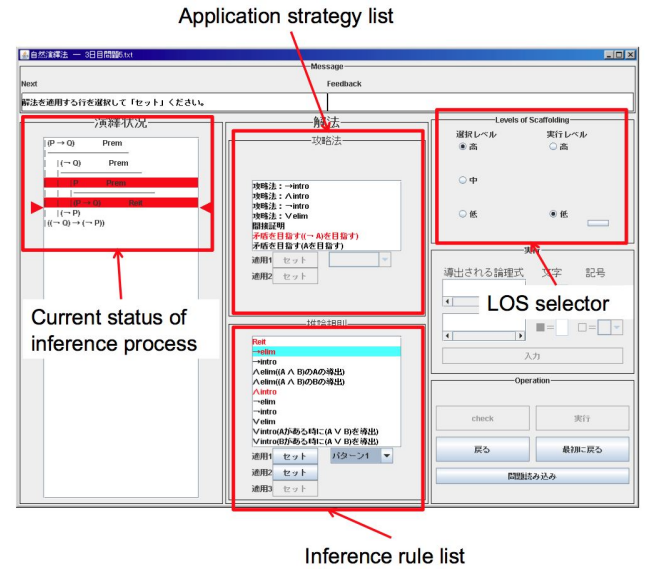

Figure 1: Example screenshot of the tutor terminal.

scaffolds the students by providing helpful information about the selection of the rules and strategies.

The LOS can be controlled from two viewpoints: rule selection and application.

LOS for rule selection: Level 3 (high): The system presents applicable candidates (rules and strategies) and the propositions to which the rules should be applied. For example, in the middle window in Figure 1, the system proposes that three highlighted inference rules and one strategy could be applied. When "$\rightarrow$ Elimination" is selected, P and $P \rightarrow Q$ are highlighted in the left window, indicating that the selected rule should be applied to these propositions. Level 2 (middle): The system presents only applicable candidates (rules and strategies). When this level is selected, students are required to find the propositions to which the selected rule should be applied without receiving support from the system. Level 1 (low): The system presents only a set of inference rules and strategies (no support is provided).

LOS for rule application: Level 2 (high): The system infers a proposition and automatically presents it in the left window. Shortly after students select an inference rule and the propositions to which the rule will be applied, the system displays the current status of deduction. Level 1 (low): The system infers a proposition, but presents only partial information of the inferred result. Students are required to complete the inference process by filling terms in blank spaces of a template.

Experiment 1

In Experiments 1 and 2, the initial setting of LOS at the beginning of solving each problem was lowest. Once a new problem was set, the LOS was initialized to Level 1. The participants were required to determine whether to raise an LOS from the initial setting while solving each example problem. Therefore, Experiments 1 and 2 investigated the participants'

keeping-off behavior in help-seeking. We will investigate the self-fading behavior in Experiment 3.

Experiment 1 was a preliminary experiment. Experiment 1 was performed in a laboratory setting; Experiments 2 and 3 were performed in a real classroom setting.

Participants and Procedure

Thirty-three participants joined Experiment 1. In the initial phase of the experiment, the participants learned the basics of ND through handout materials and an instructional video. They learned eight inference rules and four strategies without the tutoring system. After the participants were instructed on how to use the tutoring system, they learned ND by solving six example problems with our tutoring system through the 80-min learning phase. Two of the problems were difficult and required a second-order subproof, and two were easy and either did not require a subproof or required only a first-order subproof. The data recorded in this phase were analyzed.

Results

We focused on two kinds of help control behavior. One is a relatively simple behavior. In our experiments, participants were allowed to solve problems at their own pace. Some participants quit solving a problem, moved to other problems, and then revisited the initial problem and attempted to complete it. Our first hypothesis is that participants in the second and following attempts, compared to the first attempt, would not raise it for their challenge. The other is a more sophisticated behavior: we expected that participants would adopt an LOS according to the degree of difficulty of each problem. Our hypothesis is that participants solving easy problems would select a lower LOS than when solving difficult problems, despite the fact that they were permitted to receive help if they wished.

We compared the average LOS of the first attempt with that of the second-and-following attempts. In certain cases, the participants attempted to solve a problem more than twice. In such cases, we used the average score of the second-and-following attempts. Figure 2 shows the result of analysis about the LOS control between the first and second-and-following attempts. A t-test revealed a significant difference between the first and second-and-following attempts in rule selection ($t(116) = 6.10$, $p < 0.01$), but not in rule application ($t(115) < 1$, n.s.).

Figure 3 shows the result of analysis about more sophisticated behavior, i.e., the LOS control when solving easy and difficult problems. A t-test revealed a significant difference between the easy and difficult problems in both rule selection and application ($t(31) = 2.59$, $p < 0.05$; $t(31) = 4.58$, $p < 0.01$).

The above results indicated that the participants kept an LOS at low in the second and following attempts, relative to the first attempt, but only in rule selection, and did not raise it when solving easy problems compared to when solving difficult problems. These results supported our hypotheses about the participants' keeping-off behavior in help-seeking. This
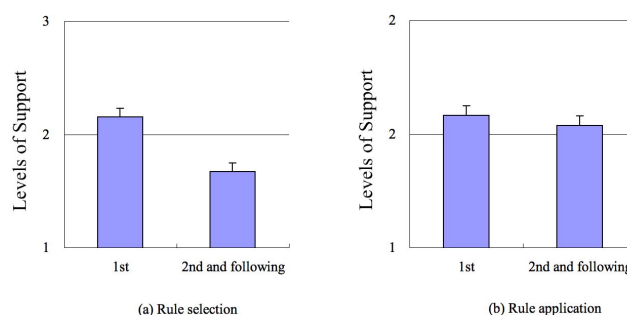

(a) Rule selection (b) Rule application

Figure 2: Levels of support versus number of attempts in Experiment 1.

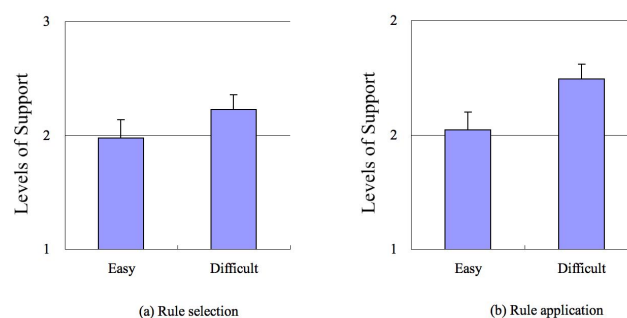(a) Rule selection (b) Rule application

Figure 3: Levels of support versus problem difficulty in Experiment 1.

stoic behavior was observed greatly in rule selection than application.

Experiment 2

In Experiment 2, we performed both pre- and post-tests, before and after the learning phase, to examine the relationship between help-seeking behavior and the learning effects. We also focused on whether the participants' stoic help-seeking behavior depends on their problem solving ability. In Experiment 1, we confirmed that stoic behavior was greatly observed in rule selection; therefore, in Experiments 2 and 3, the LOS in rule application was fixed at Level 1, and only the LOS in rule selection was investigated.

Participants and Procedure

Forty-nine participants from a cognitive science class joined the 2011 experiment. Three lessons were assigned to learn ND. In the first lesson, an instructor lectured on the basics of formal inference systems and ND as an example of such systems.

In the second lesson, the participants initially solved six problems while learning how to use the tutoring system. First, the instructor presented an example flow of problem solving. Then, participants followed the flow and reached the solution using the system. After the initial training, the participants were given two new problems to solve. Finally, the
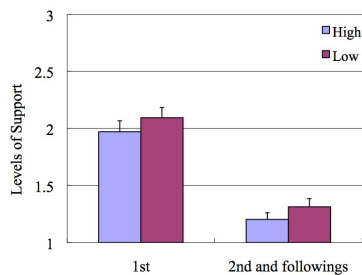
Figure 4: Levels of support versus number of attempts in Experiment 2.

Figure 5: Levels of support versus problem difficulty in Experiment 2.

Figure 6: Levels of support versus learning gains in Experiment 2.

participants were given a paper test in which they solved a test problem; we used this test as a pretest in the following analysis.

The log data from the third lesson were analyzed. The participants solved eight problems at their own paces and selected an LOS. Three of the eight problems were easily solved by applying the basic rules learned in the first lesson. However, three problems were relatively difficult, and their solutions required more complex rules and solution strategies, such as subgoal settings. The learning session lasted for an hour. After the learning session, a post-test was performed.

Results

To investigate whether the participants' help-seeking behavior is dependent on their problem solving ability, we divided the participants into two groups on the basis of their pre-test scores, and formed lower- and higher-score groups.

Figure 4 shows the results of analysis on simple help management behavior, i.e., the LOS control between the first and second-and-following attempts. A two (attempt: first and second-and-following) x two (ability: high and low) ANOVA revealed that the main effect of the attempt factor reached significance ($F(1, 158) = 136.93$, $p < 0.01$), but the main effect of the ability factor did not ($F(1, 158) = 1.54$, n.s.). There was no interaction between the two factors ($F(1, 158) < 1$, n.s.).

Figure 5 shows the results of analysis on more sophisticated behavior, i.e., the LOS control when solving the easy and difficult problems. A two (problem: easy and difficult) x two (ability: high and low) ANOVA revealed that the main effect of the problem factor reached significance ($F(1, 44) = 33.02$, $p < 0.01$), but the main effect of ability factor did not ($F(1, 44) < 1$, n.s.). There was no interaction between the two factors ($F(1, 44) = 2.96$, n.s.).

The above results duplicated the participants' stoic behavior captured in Experiment 1. Additionally, the same tendency was observed in both low- and high-score groups, meaning that such help-seeking behavior does not depend on the participants' problem solving ability.

Next, we focus on the analysis of the relation of help-seeking behavior and learning effects. We hypothesize that a lower LOS may provide greater learning effects and a higher
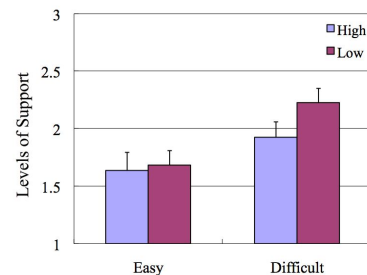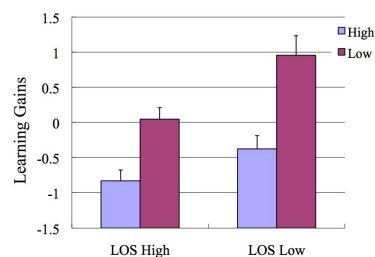
LOS may obstruct effective learning. We divided the participants into two groups on the basis of their average LOS during problem solving in the learning phase. The problem used in the pre-test was different from those used in the post-test; therefore, we cannot directly compare the scores of the two tests. Accordingly, we transferred the test scores to the z-scores in each of the two tests and calculated the gains of the z-score from the pre- to post-tests.

Figure 6 shows the results of the analysis. A two (LOS in learning phase: high and low) x two (ability: high and low) ANOVA revealed that both the main effects of the LOS factor and the ability factor reached significance ($F(1, 45) = 8.28$, $p < 0.01$; $F(1, 45) = 26.98$, $p < 0.01$). There was no interaction between the two factors ($F(1, 45) = 1.44$, n.s.).

The result shows that the participants who learned with a lower LOS in the learning phase gained greater learning effects. This means that stoic behavior, especially the keeping-off behavior in this case, promoted learning.

Experiment 3

In Experiment 2, we focused on the keeping-off behavior. Experiment 3 investigated the self-fading behavior in help-seeking,

Participants and Procedure

Twenty-eight participants from a cognitive science class joined our 2012 experiment. Three lessons were assigned for learning ND and the learning content and procedures were almost identical to Experiment 2. The crucial differ-
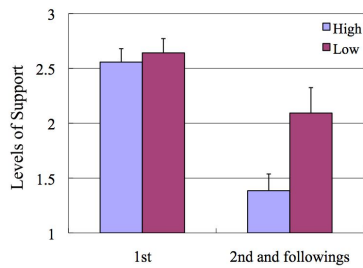
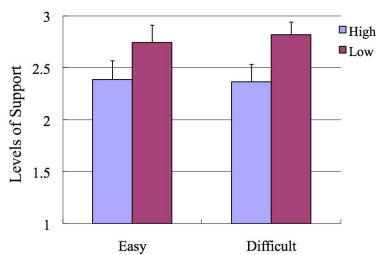Figure 7: Levels of support versus number of attempts in Experiment 3.

Figure 8: Levels of support versus problem difficulty in Experiment 3.

ence was that the initial setting of LOS at the start to solve each problem was the highest (Level 3) in Experiment 3. The participants were required to determine whether to lower an LOS from the initial setting, while solving example problems. Therefore, Experiment 3 investigated the participants' self-fading behavior in help-seeking.

Results

Figure 7 shows the results of analysis on the LOS control between the first and second-and-following attempts. A two (attempt: first and second-and-following) x two (ability: high and low) ANOVA revealed that the main effects of both the attempt and ability factors reached significance ($F(1, 54) = 34.43$, $p < 0.01$; $F(1, 54) = 4.98$, $p < 0.05$). There was significant interaction between the two factors ($F(1, 54) = 4.47$, $p < 0.05$). The simple main effect of the ability factor at the first attempt was not significant ($F(1, 108) < 1$, n.s.), but the effect at the second-and-following attempts was significant ($F(1, 108) = 9.41$, $p < 0.01$).

Figure 8 shows the result of analysis about the LOS control when solving the easy and difficult problems. A two (problem: easy and difficult) x two (ability: high and low) ANOVA revealed that there was neither a main effect of the problem factor nor a main effect of the ability factor ($F(1, 25) < 1$, n.s.; $F(1, 25) = 2.67$, n.s.). There was no interaction between the two factors ($F(1, 25) < 1$, n.s.).

In Experiment 3, we confirmed the stoic behavior only in the LOS control between the first and second-and-following attempts, but not in the LOS control when solving easy and difficult problems. In the former case, the higher abil-
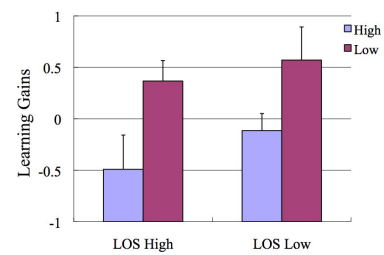

Figure 9: Levels of support versus learning gains in Experiment 3.

ity participants greatly lowered the LOS in the second-and-following attempts compared with the lower ability participants.

Next, we focus on the analysis of the relationship between the self-fading behavior and learning effects. The same analysis as in Experiment 2 was performed. Figure 9 presents the results of the analysis. A two (LOS in learning phase: high and low) x two (ability: high and low) ANOVA revealed that there was neither a main effect of the LOS factor nor interaction between the two factors ($F(1, 24) < 1$, n.s.; $F(1, 24) < 1$, n.s.). However, the main effect of the ability factor reached significance ($F(1, 24) = 5.09$, $p < 0.05$). Learning effects by the stoic help-seeking behavior were not confirmed in Experiment 3.

Discussion and conclusions

The first research question we posed was: Do students exhibit stoic behavior in hint seeking? We examined two types of stoic behavior: Experiment 2 investigated the keeping-off behavior, and Experiment 3 investigated the self-fading behavior. We hypothesized that participants would lower an LOS with the development of learning (i. e., from the first to the second-and-following attempts). This hypothesis was fully supported in both Experiments 2 and 3. More specifically, with regard to self-fading, high ability participants more actively lowered the LOS in the second-and-following attempts compared with the lower ability participants. The second hypothesis was whether participants would adaptively manage their help-seeking behavior based on the degree of problem difficulty. We expected that they lower an LOS or would not raise it when solving easy problems, compared to when solving difficult problems. This hypothesis was supported only in Experiment 2, indicating that the participants kept an LOS at low (Experiment 2), but that they did not reduce an LOS from high to low (Experiment 3) when solving easy problems, which suggests that the keeping-off behavior was confirmed, though the self-fading behavior was not.

These results imply that the self-fading behavior, as an adaptive behavior in help-seeking, was more difficult for the participants than the keeping-off behavior. The latter behavior comes from a strategy to set the LOS at low by stopping action (i.e., stopping raising an LOS). However, the former

behavior comes from a strategy to set the LOS at low by performing an action (i.e., beginning to reduce an LOS). The latter is relatively passive, while the former is an intentional and active behavior. These results suggest that an active type of stoic behavior was more difficult for users.

The second research question was: Does the stoic behavior in hint seeking promote learning gains? This relates to a trade-off of selecting either the problem-solving or the learning goal. Participants learn while solving instance problems given by a tutoring system. Attaining the problem-solving goal means solving such instance problems as accurately and rapidly as possible. However, the learning goal requires another attainment that is usually more essential. The primary objective is not to solve instance problems, but to learn by solving instances. Dweck classified two types of goals: learning and performance (Dweck, 1986; Ames, 1992). Highly motivated children tend to set learning goals in an effort to increase their competence levels for understanding or mastering something new rather than simply solving problems. Our previous study confirmed that high learning supports promote the problem solving goal setting, and refrain the learning goal setting (Miwa, Terai, & Nakaike, 2012). In the high LOS situation, participants may solve training problems accurately and rapidly in the learning phase, but tend to learn least from the training.

The assistance dilemma hypothesizes an optimum point of learning effects. Koedinger et al. (2008) demonstrated a reverse U-shape learning curve as a function of cognitive load (Koedinger, Pavlik, Mclaren, & Aleven, 2008). This means that extremely lower and higher cognitive loads result in negative impacts on learning. The levels of help support are correlated with learners' cognitive load while learning. Much help reduces their cognitive load for problem solving in the learning phase where students simply respond to help indications from a tutoring system, e.g., a direct instruction about what to do next, without deeper consideration. From this viewpoint, our experimental results are considered to capture the right side of the reverse U-shape. We compared learning effects when the participants learned with a low and a high LOS. In the right half, the reversed U-shape predicts that a lower LOS provides more learning effects; Experiment 2 supported this prediction. However, we also expected that the effects of learning decrease gradually, in the left side of the reverse U-shape, as the LOS is reduced. In another experiment (Miwa et al., 2012), we confirmed this prediction using the same tutoring system, in which we set up two experimental conditions. In the system condition, the participants learned ND using our tutoring system. They were permitted to control the LOS. In the control condition, i.e., the paper-and-pencil condition, participants learned ND without a tutoring system, instead they learned ND using only a textbook. The latter was the no support condition. Results showed that learning effects in the system condition were greater than in the paper-and-pencil condition. In this experiment, no support relates to the leftmost side of the reversed U-shape.

References

Aleven, V., & Koedinger, K. R. (2000). Limitations of student control: Do student know when they need help? In *Proceedings of the 5th international conference on intelligent tutoring systems, its 2000* (p. 292-303).

Ames, C. (1992). Classrooms: Goals, in structure, and student motivation. *Journal of Educational psychology*, *84*, 261-271.

Atkinson, R. K., Renkl, A., & Merrill, M. M. (2003). Transitioning from studying examples to solving problems: Effects of self-explanation prompts and fading worked-out steps. *Journal of Educational Psychology*, *95*, 774-783.

Dweck, C. S. (1986). Motivational processes affecting learning. *American Psychologist*, *41*, 1040-1048.

Koedinger, K. R., & Aleven, V. (2007). Exploring the assistance dilemma in experiments with cognitive tutors. *Educational Psychology Review*, *19*, 239-264.

Koedinger, K. R., Pavlik, P., Mclaren, B., & Aleven, V. (2008). Is it better to give than to receive? the assistance dilemma as a fundamental unsolved problems in the cognitive science of learning and instruction. In *Proceedings of the 30th annual conference of the cognitive science society* (p. 2155-2160).

Miwa, K., Morita, J., Nakaike, R., & Terai, H. (2013). Learning through intermediate problems in creating cognitive models. *Interactive Learning Environments. (in press)*.

Miwa, K., Terai, H., & Nakaike, R. (2012). Tradeoff between problem-solving and learning goals: Two experiments for demonstrating assistance dilemma. In *Proceedings of the 34th annual conference of the cognitive science society* (p. 2008-2013).

Razzaq, L., & Heffernan, N. (2010). Hints: Is it better to give or wait to be asked? In *Proceedings of the 10th international conference on intelligent tutoring systems, its 2010* (p. 349-358).

Roll, I., Aleven, V., McLaren, B., & Koedinger, K. R. (2007). Can help seeking be tutored? searching for the secret sauce of metacognitive tutoring. In *Proceedings of the 13th international conference on artificial intelligence in education, aied 2007* (p. 203-210).

Roll, I., Aleven, V., McLaren, B., & Koedinger, K. R. (2011). Improving students' help-seeking skills using metacognitive feedback in an intelligent tutoring system. *Learning and Instruction*, *21*, 267-280.

Roll, I., Aleven, V., McLaren, B., Ryu, E., Baker, R., & Koedinger, K. R. (2006). The help tutor: Does metacognitive feedback improve students' help-seeking actions, skills, and learning? In *Proceedings of the 8th international conference on intelligent tutoring systems, its 2006* (p. 360-369).

Shute, V. J. (2008). Focus on formative feedback. *Review of Educational Research*, *78*, 153-189.

Wood, H., & Wood, D. (1999). Help seeking, learning and contingent tutoring. *Computers and Education*, *33*, 153-169.