Learning and Instruction 21 (2011) 267–280

Learning and Instruction

www.elsevier.com/locate/learninstruc

Improving students' help-seeking skills using metacognitive feedback in an intelligent tutoring system

Ido Roll [a,*], Vincent Aleven [b], Bruce M. McLaren [b], Kenneth R. Koedinger [b]

[a] *Carl Wieman Science Education Initiative, University of British Columbia, 6224 Agricultural Road, Vancouver, British Columbia V6T 1R9, Canada*
[b] *Human Computer Interaction Institute, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh PA 15213, USA*

Abstract

The present research investigated whether immediate metacognitive feedback on students' help-seeking errors can help students acquire better help-seeking skills. The Help Tutor, an intelligent tutor agent for help seeking, was integrated into a commercial tutoring system for geometry, the Geometry Cognitive Tutor. Study 1, with 58 students, found that the real-time assessment of students' help-seeking behavior correlated with other independent measures of help seeking, and that the Help Tutor improved students' help-seeking behavior while learning Geometry with the Geometry Cognitive Tutor. Study 2, with 67 students, evaluated more elaborated support that included, in addition to the Help Tutor, also help-seeking instruction and support for self-assessment. The study replicated the effect found in Study 1. It was also found that the improved help-seeking skills transferred to learning new domain-level content during the month following the intervention, while the help-seeking support was no longer in effect. Implications for metacognitive tutoring are discussed.

Keywords: Help seeking; Self-regulated learning; Cognitive tutors; Intelligent tutoring systems; Metacognition

1. Introduction

1.1. Help seeking in tutoring systems

Knowing when and how to seek help during learning is a key self-regulatory skill (Nelson-Le Gall, 1981; Newman, 1994; Pintrich, 2000). Research in classrooms suggests that adaptive help-seeking behavior helps students learn more effectively (Arbreton, 1998; Karabenick & Newman, 2009; Ryan & Shin, 2011). Help requests (and responses) can be characterized as "instrumental" when they are aimed at learning, or "executive," when they are aimed merely at completing tasks (Karabenick & Knapp, 1991). Effective help seeking has also been shown to be associated with better learning in educational technologies (Aleven, McLaren, Roll, & Koedinger, 2006; Renkl, 2002; Wood & Wood, 1999). In the present article the focus is on help seeking within an intelligent tutoring system, that is, a system that provides step-by-step guidance as students practice a complex problem-solving skill (VanLehn, 2006).

Most tutoring systems offer help, often in the form of on-demand contextual hints that explain how to solve the current problem step and why (for a review, see Aleven, Stahl, Schworm, Fischer, & Wallace, 2003). Students often can ask for multiple levels of hints for each step (Aleven et al., 2006; Karabenick & Newman, 2009; Mathews & Mitrovic, 2008). When asking for a hint, the student first receives a very general hint, intended to point out a good choice for the next problem step. The student can then ask for more elaborated hints. These hints offer instrumental help by giving increasingly detailed explanations of how to solve the step. The last level of hint often gives the answer for the step, essentially converting a too-challenging problem step into an annotated example. This last-level hint, often called "bottom-out hint", allows

* Corresponding author. Tel.: +1 778 938 4141; fax: +1 604 822 2102.
E-mail address: ido@phas.ubc.ca (I. Roll).

0959-4752/$ - see front matter © 2010 Elsevier Ltd. All rights reserved.
doi:10.1016/j.learninstruc.2010.07.004

Table 1
Hint levels in the Geometry Cognitive Tutor.

Hint level	Hint text	Type of hint
1.	"Enter the value of the radius of circle A"	Orientation
2.	"How can you calculate the value of the radius of circle A given the value of the diameter of circle A?"	Instrumental help
3.	"The radius of a circle is half of the diameter"	Instrumental help
4.	"The radius of circle A = 46.5"	Bottom-out hint

students to complete problem steps that are beyond their ability and to reflect on the solution process (Aleven et al., 2006).[1] Table 1 includes an example of the different levels of hints for a typical problem step in an intelligent tutoring system, with Level 4 being the bottom-out hint.

While effective help-seeking behavior has been shown to be important for successful learning with tutoring systems, students often make suboptimal use of the help resources these systems offer (Aleven et al., 2003). Research suggests two main forms of help misuse, namely, help avoidance (underuse) and help abuse (overuse). For example, students may avoid seeking help altogether, even when clearly in need of more guidance (Aleven et al., 2003). Conversely, students may ask for elaborated help when no or little help is needed, or use help in executive ways aimed at finding answers without thinking through the material (Aleven et al., 2006; Karabenick & Newman, 2009). One common form of executive help seeking is rapidly clicking through hints to reach the bottom-out hint to get the answer (Baker, Corbett, Roll, & Koedinger, 2008). Notably, students who are most in need of help make poorer decisions regarding their help-seeking behavior (Karabenick & Newman, 2009; Renkl, 2002; Wood & Wood, 1999). Given the evidence that maladaptive help-seeking with tutoring systems is widespread, attempting to improve students' help-seeking behavior in a tutoring system may achieve lasting positive effects on students' ability to learn from such systems (Mercier & Frederiksen, 2007; Roll, Aleven, McLaren, & Koedinger, 2007b).

1.2. Improving students' help-seeking behavior in tutoring systems

One approach to improving the quality of students' help seeking while solving problems with a tutoring system is by delegating some of the responsibility to the system, rather than to the students. For example, a number of systems use a contingent help mechanism in which the level of the hint is automatically adapted to the student's knowledge level, as estimated based on past performance (Luckin & du Boulay, 1999; Wood & Wood, 1999). Similarly, few tutors from the Cognitive Tutor family (although not the one described below) automatically present hints once a certain threshold number of errors is reached, or enforce a delay of 2 s between repeated hint requests.

While these approaches are likely to result in better help-seeking behavior while they are in effect, they are not likely to help students become better independent help-seekers.

A more complete solution would be to help students acquire better help-seeking skills. Robust learning of help-seeking skills would allow students to transfer these skills to novel learning situations, when no explicit support for help seeking is available. The research described in the present article examined whether immediate feedback of a metacognitive nature on students' help-seeking behavior could achieve the goal of robust improvement in students' help-seeking skills. This was done by adding automated feedback on students' help-seeking errors to a commercial software, the Geometry Cognitive Tutor (Koedinger & Corbett, 2006). At the cognitive (or domain) level, immediate feedback has been shown to improve learning (Corbett & Anderson, 2001). At the metacognitive level, however, the effect of feedback on students' learning is understudied. *Metacognitive feedback* is defined as feedback that is triggered by students' learning behavior (e.g., avoiding necessary help), and not by the accuracy of their responses at the domain level. Also, metacognitive feedback delivers metacognitive content, that is, it conveys information about desired learning behavior (e.g., advising the student to ask for a hint), rather than domain knowledge (Roll et al., 2007b). A small number of systems use metacognitive feedback to direct students in their learning process (Shute & Glaser, 1990; Wagster, Tan, Wu, Biswas, & Schwartz, 2007). However, to the best of our knowledge, the effect of metacognitive feedback on students' help-seeking behavior is yet to be fully evaluated in controlled studies. Furthermore, we are not aware of any evaluation of the long-term effect of help-seeking training within the context of a tutoring system.

1.3. Identifying help-seeking errors

To give metacognitive feedback, the tutoring system must be able to detect metacognitive errors in real time, without interrupting the learning process. This was done by evaluating students' actions using a metacognitive computational model of help seeking (Aleven et al., 2006). The help-seeking model evaluates help-seeking behaviors in a tutored step-based problem-solving environment such as the Geometry Cognitive Tutor. Unlike other models of help seeking that have been put forward in the literature (Mercier & Frederiksen, 2007; Nelson-Le Gall, 1981), this model is detailed enough to classify individual actions as help-seeking desired or undesired actions. Aleven et al. (2006) present a detailed comparison of the model to other frameworks.

Following the Adaptive Control of Thought – Rational (ACT-R) theory of cognition and learning (Anderson, 1993), the help-seeking model describes help-seeking actions as outcomes of desirable and undesirable rules, taking the form of if-then

[1] The term "bottom-out hint" refers to the most elaborated hints that convey the answer. Students view higher-level hints prior to receiving a bottom-out hint, and can re-read the higher-level hints after reading the bottom-out hint.

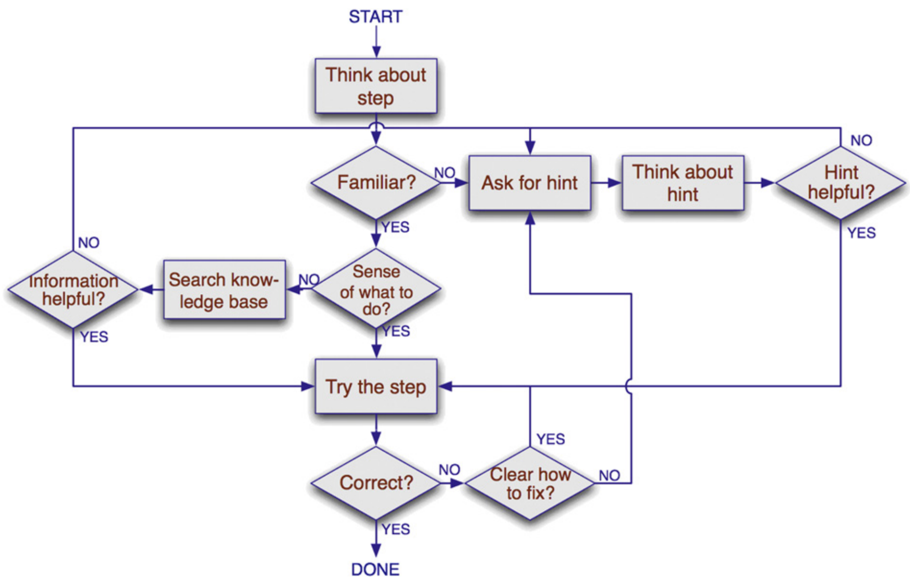

Fig. 1. The help-seeking model. This flowchart describes the main decision points of the complete help-seeking model. The full model is implemented using 80 if-then rules.

statements (see Fig. 1; Aleven et al., 2006). For example, according to the model, a student who does not have sufficient knowledge to solve a certain step should ask for enough help to understand the step (instrumental help-seeking). Conversely, asking for a bottom-out hint is not desirable when the student's knowledge level and help-seeking pattern suggest that the hint is to be used in an executive manner (for example, when "drilling down" to bottom-out hint without reading intermediate hint levels). Similarly, a student who has sufficient knowledge to solve a specific problem step should answer it without requesting help, and should do so deliberately (i.e., taking sufficient time to think). The help-seeking model predicts desirable metacognitive actions in any given problem-solving situation within a cognitive tutor environment. Likewise, in any given problem-solving situation within this environment, the help-seeking model is capable of predicting specific help-seeking errors. The help-seeking model is implemented using 80 if-then statements (or rules). These rules are inclusive, that is, each action that a student does in the cognitive tutor is labeled as either desired or undesired help-seeking behavior. This is done by comparing each action performed by the student against those predicted by the help-seeking model, so as to automatically assess a student's help-seeking behavior moment-by-moment, over extended periods of time. The undesirable forms of help-seeking behavior captured by the help-seeking model were shown to correlate with poor learning gains at the domain level across student populations and domains (Aleven et al., 2006).

1.4. Research questions and hypotheses

The aims of the present studies were to investigate the validity of the help-seeking model as implemented within the Geometry Cognitive Tutor and to evaluate the effect of metacognitive feedback that is based on the help-seeking model on students' help-seeking skills. The paper focuses on the direct effect of help-seeking support on help-seeking skills, and only to a lesser degree on the indirect effect on domain-learning gains. While the effect of given help on domain knowledge acquisition has been previously studied, the effect of automated metacognitive feedback on students' help-seeking behavior is yet to be evaluated.

The first research question focused on whether the help-seeking model, which identifies desired and undesired help-seeking behaviors, has construct validity. While the help-seeking model was shown to correlate with pre to posttest improvements, learning gains are only an indirect measure of metacognitive behavior, and as such, cannot be the sole source of assessment of metacognitive behavior (MacLeod et al., 1996). To establish construct validity, the current studies examined the relationship between poor help-seeking behavior in the cognitive tutor (according to the help-seeking model) and poor help-seeking behavior outside the tutoring environment (such as a paper posttest). The prediction was that the quality of students' help-seeking behavior, as evaluated by the help-seeking model, would correlate with other measures of help-seeking based on the paper posttest (Hypothesis 1).

The second research question examined whether automated metacognitive feedback, in the context of tutored problem-solving, leads to a more appropriate help-seeking behavior while this feedback is in effect. This was done by comparing students' behavior using the standard version of the Geometry Cognitive Tutor to one that is enhanced with the Help Tutor, a tutoring agent that provides immediate metacognitive feedback on students' help-seeking behavior. The prediction was that metacognitive feedback on help-seeking errors would enhance students' more appropriate help-seeking behavior (such as avoiding redundant bottom-out hints) as compared to help-seeking behavior of students who do not receive such feedback (Hypothesis 2).

The third research question examined whether the effect of the metacognitive feedback (in terms of improved help-seeking behavior) transfers to learning new domain content within the same learning environment, even after the support of the Help Tutor is no longer in effect. The hypothesis was that the acquired help-seeking skills would transfer to new domain-level content even in the absence of the help-seeking feedback. At the same time only limited transfer would be found to new tasks outside the Geometry Cognitive Tutor, such as a help-seeking dilemmas paper test (Hypothesis 3).

In what follows, first Study 1 is described, in which the construct validity of the help-seeking model and the effect of metacognitive feedback were evaluated. Then, Study 2 is described in which the transferability of the acquired help-seeking skills was examined.

2. Study 1: Validating the help-seeking model and evaluating the effect of the help-seeking feedback

Study 1 tested Hypotheses 1 and 2.

2.1. Method

2.1.1. Design

Half of the participating students worked with the Help Tutor, an enhanced version of the Geometry Cognitive Tutor that is capable of giving immediate metacognitive feedback on students' help-seeking errors (Help condition). The other half of the students worked with the unmodified Geometry Cognitive Tutor (Control condition). Students were assigned to condition under the constraint that the conditions were balanced with respect to students' previous achievement in the Geometry Cognitive Tutor class.

2.1.2. Participants

58 students who were enrolled in a Geometry Cognitive Tutor class participated in the study. Two classes, with 28 students, taught by one teacher, were from an inner-city school in Pittsburgh (97% minorities, 22% mathematical proficiency, 48% male to 52% female). A different teacher in a suburban school (2% minorities, 84% mathematical proficiency, 51% male to 49% female) taught the other two classes, with 30 students. All students were in Grades 10–11 (15–17 years old). Data regarding proficiency, ethnicity, and gender refers to the entire cohort and is representative of the classes we worked with.

2.1.3. Materials

During the study, all participating students worked on a unit in the Geometry Cognitive Tutor dealing with the geometric properties of circles. As intelligent tutoring systems typically do (VanLehn, 2006), the Geometry Cognitive Tutor provides step-by-step guidance with a complex problem-solving skill. It uses a cognitive model to track students in their various approaches through these problems, estimate their knowledge level, and provide domain-level feedback and on-demand hints regarding the steps of these problems. Fig. 2 shows the interface of the tutor used in this study. In the Scenario window the tutor presents a problem statement and the students type their answers (on the left). In the upper-right-hand corner students can see their estimated proficiency level on each skill. The Geometry Cognitive Tutor has two built-in

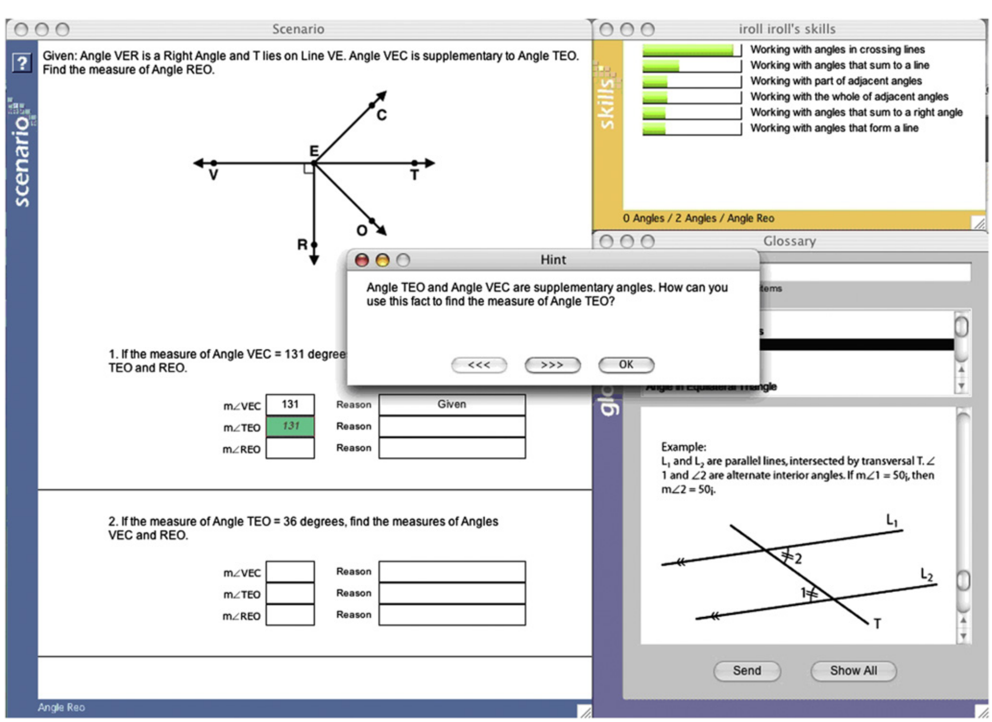

Fig. 2. The Geometry Cognitive Tutor interface.

help-seeking mechanisms: on-demand contextual hints and a glossary. Students can ask for contextual hints at increasing levels of detail, as described above, by clicking on the hint button (marked "?"). The second help resource is a searchable Glossary that contains all relevant theorems, illustrated with examples (bottom right corner).

In addition to the domain-level feedback offered by the Geometry Cognitive Tutor itself, half of the students in the study received metacognitive feedback from the Help Tutor. The Help Tutor is an add-on tutor agent to the Geometry Cognitive Tutor. It teaches help-seeking skills by giving metacognitive feedback on students' help-seeking errors in the context of learning a domain-specific problem-solving skill. The Help Tutor uses the help-seeking model described above to trigger its feedback. When the student commits a help-seeking error (for example, by "drilling down" to the bottom-out hint without reading intermediate hints), the Help Tutor gives immediate feedback (e.g., "No need to hurry so much. Take your time and read the hint carefully. Consider trying to solve this step without another hint. You should be able to do so"; see Fig. 3.). The Help Tutor messages include only domain-independent metacognitive content for several reasons: to encourage students to focus more on the metacognitive feedback (and not be distracted by domain content), to help students generalize the help-seeking skills, and to make the Help Tutor reusable with different Cognitive Tutors. The Help Tutor messages use the existing hint-window mechanism and are distinguished from regular hints in their font (color and type).

2.1.4. Instruments

To address the first research question, namely evaluating the construct validity of the help-seeking model, students' help-seeking behavior was assessed using a paper-and-pencil posttest. The posttest included 9 geometry problems with 18 problem steps (see example in Fig. 4). Most of these problems included a procedural component, in which students were asked to calculate the measure of an angle, a circle arc, or a circle chord, and a declarative component, in which students were asked to name the theorem they used. To evaluate help-seeking behavior within the posttest, independent of domain knowledge, all items in the test appeared in two versions, counter balanced between forms: with or without hints. Items without hints (i.e., conventional test items) were used to assess the difficulty level of the specific test item for the specific group of students. The other type of items included hints, as seen in Fig. 4. The difference between students' performance on identical items with and without hints (for the same group of students) is a direct measure of the quality of students' hint comprehension and application abilities.

A different type of item was used to assess students' ability to identify desired help-seeking behavior. In these items, students were asked to choose the most appropriate action in hypothetical help-seeking dilemmas (cf., Al-Hilawani, 2003). These situations described common situations that students encounter when working with the Geometry Cognitive Tutor and adhere to the help-seeking strategies put forward by the help-seeking model. For example, the hypothetical situation described in Fig. 5 suggests that the student knows the target

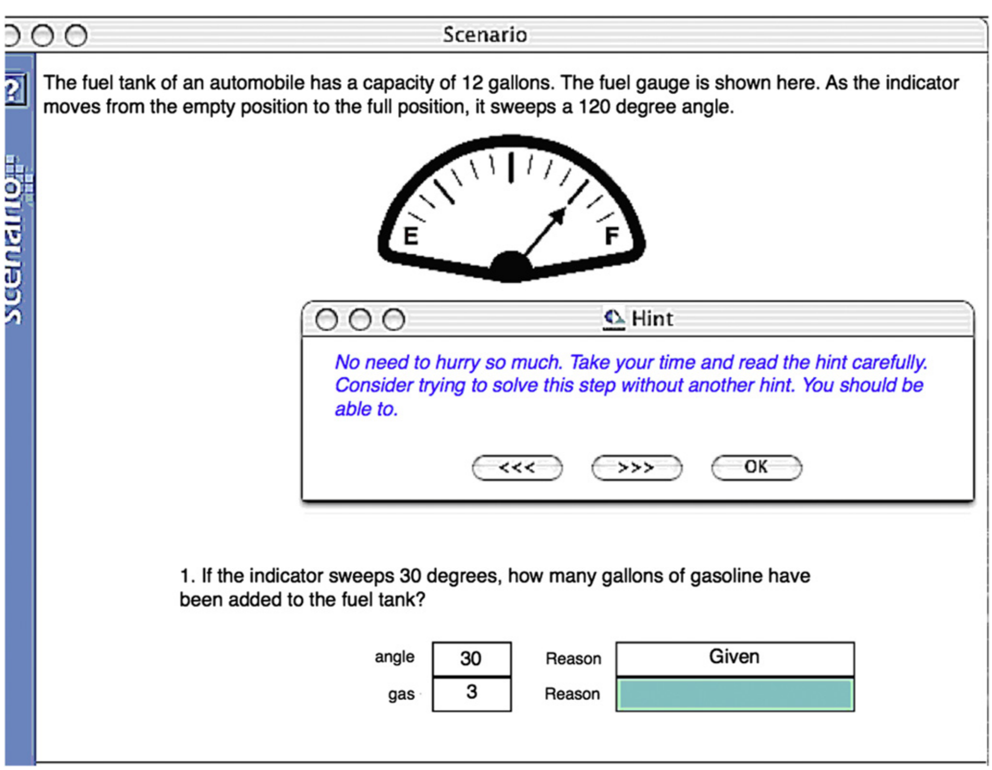

Fig. 3. The Help Tutor feedback messages.

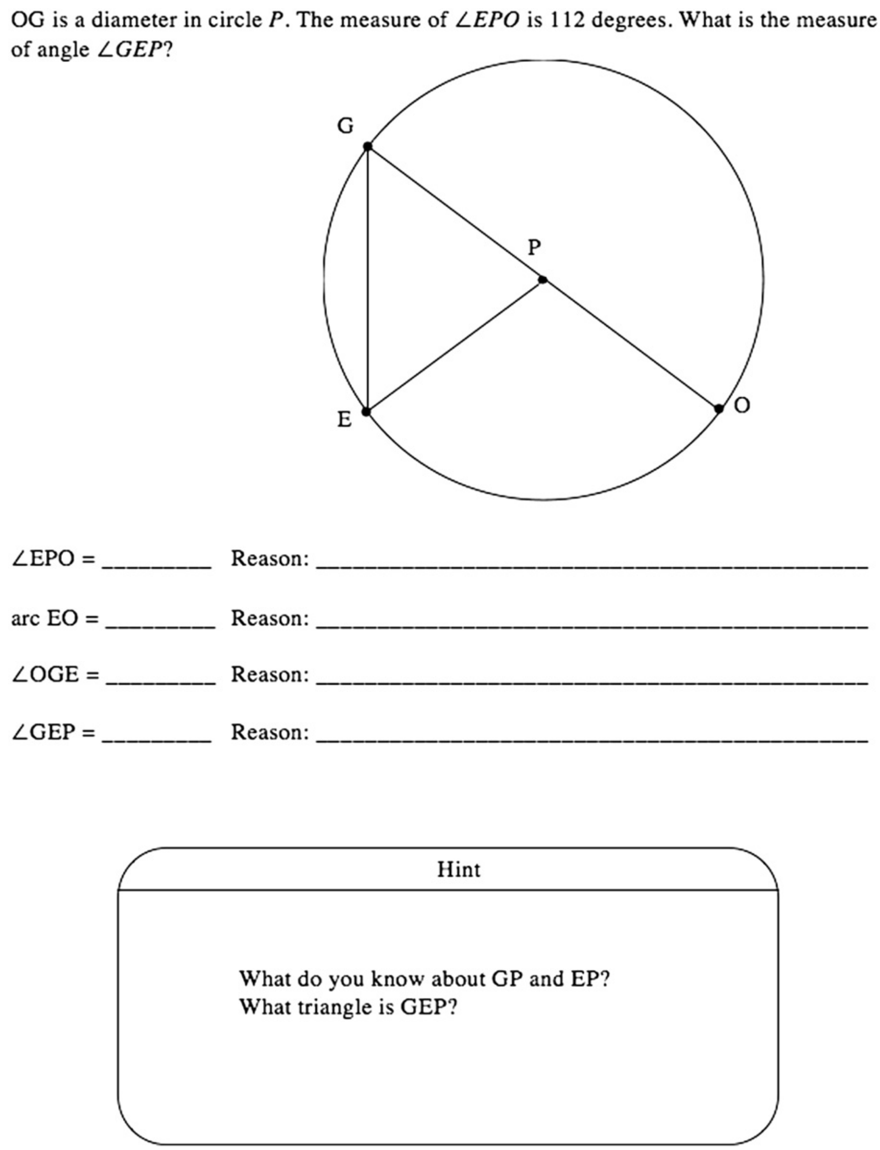

Fig. 4. Pre and posttest items.

material, and yet, has not tried to correct her own error. The help-seeking model suggests that at this point the student should attempt to review her calculations prior to asking for a hint. Note that these items do not always assume that more help is better. For example, the situation described in Fig. 5 suggests that less help is a more appropriate action. To reduce guessing, students were given 1 point for each correct answer and −1 for each wrong answer.

To address the second research question, the effect of the metacognitive feedback, students' help-seeking behavior while working with the Geometry Cognitive Tutor was evaluated using the help-seeking model. In addition to triggering feedback, the output of the help-seeking model can be used to calculate error rate for different types of help-seeking errors.

A secondary goal of the study was to evaluate the effect of metacognitive feedback on domain-level learning gains (in this case, the properties of circles). Since the domain results are reported elsewhere (Roll et al., 2006), as well as the fact that the present article focuses on metacognitive learning results (and not an indirect improvement in domain learning), the domain results are not included here. Yet, to convey a complete picture of the effect of the given metacognitive feedback, domain-learning outcomes are discussed in Section 2.3.2, the Effect of the Help Tutor.

2.1.5. Procedure

Students worked with the Geometry Cognitive Tutor for six 45-min sessions over three weeks. No instruction on help seeking or circles was given in advance. The tutor recorded detailed log files of students' interactions with the tutor. The recorded information included attempted solutions, use of help resources, and time stamps. Pre and posttests were collected on paper on the first and last day of the study. The students in the inner-city school did not complete the posttest and, thus, only data logged by the Cognitive Tutor is available for

1. You tried to answer a question that you know, but for some reason the tutor says that your answer is wrong. What should you do?

[ ] First I would review my calculations. Perhaps I can find the mistake myself?

[ ] The Tutor must have made a mistake. I will retype the same answer again.

[ ] I would ask for a hint, to understand my mistake.

Fig. 5. Hypothetical help-seeking dilemmas assessment.

students in these classes. These logs include online help-seeking and domain-level problem-solving data.

2.1.6. Data analysis

Data from the log files was used to calculate a help-seeking error rate for each student (i.e., what proportion of a student's actions were categorized as help-seeking errors by the model). The error rate for specific actions (e.g., faulty hint requests) is reported out of the overall instances of that action (e.g., out of all hint requests) and not out of all actions. The mean help-seeking error rate for each condition was computed by averaging the students' individual error rates.

Performance on the different measures was analyzed using a two-tailed equal-variance *t*-test. Comparisons between different types of measures were done using partial correlations. An alpha level of .05 was used in all statistical tests.

2.2. Results

Overall, students solved on average 128 problem steps per student ($SD = 61$). Students performed on average 685 actions per student ($SD = 962$), including correct and incorrect entries, hint requests, and glossary access.

2.2.1. Validation of the help-seeking model

The first aim of Study 1 was to evaluate whether actions assessed as help-seeking errors by the help-seeking model are associated with other forms of poor help-seeking behavior. The difference between students' performance on posttest items with no hints and items with hints offer a direct measure of how well students use hints in the paper tests. The online help-seeking error rate (that is, the proportion of actions labeled help-seeking errors) correlated with performance on items with hints on the paper posttest, controlling for performance on items with no hints on the same test (aggregated across conditions; partial $r = -.50$, $p < .01$).

Students' online help-seeking error rate also correlated significantly with their performance on the hypothetical dilemmas assessment. A median split analysis showed that students who made fewer help-seeking errors in the tutor scored significantly higher on the hypothetical dilemmas assessment (77% vs. 59% respectively), $t(28) = 2.2$, $p = .04$, Cohen's $d = 0.83$.

2.2.2. Effect of the metacognitive feedback on students' behaviour in the tutor

The second aim of Study 1 was to evaluate how the metacognitive feedback students received affected their online behavior within the Geometry Cognitive Tutor. Overall, the mean help-seeking error rate of students in the Help condition was significantly lower than that of the Control condition (14% vs. 16% accordingly), $t(64) = 2.0$, $p = .05$, Cohen's $d = 0.50$ (see Table 2). The Help Tutor had a different effect on different actions. Specifically, when asking for hints, students working with the Help Tutor made significantly fewer help-seeking errors. In fact, 26% of students' hint requests in the Help condition were categorized as faulty, compared to 36% in the Control condition, $t(57) = 5.7$, $p < .001$, Cohen's $d = 1.51$. The Help Tutor also had an effect on the level of hints students asked to see. Students in the Control group asked for the bottom-out hint on 70% of their hint sequences, while students in the Help condition asked for the bottom-out hint on only on 48% of their hint sequences, $t(64) = 4.3$, $p < .001$, Cohen's $d = 1.07$.

Help-seeking errors can also be made when trying a step, for example, when students attempt a step too quickly or make an attempt when the model determines they should seek help. There was no significant difference between the conditions in the rate of metacognitive errors related to solution attempts, $t(58) = .53$, *ns*. There was also no

Table 2
Multiple measures from Study 1 as a function of group. Top: Students' help-seeking error rates (%) and effect sizes (Cohen's *d*) while working with the Geometry Cognitive Tutor. Bottom: Students' posttest scores on items with and without embedded hints, and scores on the hypothetical help-seeking dilemmas paper assessment.

	Help group	Control group	Cohen's *d*
Help-seeking error rate			
Overall help-seeking error rate	14%*	16%	0.50
Hint	26%**	36%	1.50
Try-step	8%	9%	
Glossary	22%	23%	
% of hint sequences to reach bottom-out hint	48%**	70%	1.10
Help-seeking assessments during posttest			
Posttest items with hints	.41	.40	
Posttest items with no hints	.35	.31	
Hypothetical help-seeking dilemmas	.74	.64	

*$p < .05$; **$p < .01$.

significant difference between the conditions on the rate of faulty glossary searches, $t(49) = .27$, *ns*.

2.2.3. Effect of the help tutor on other measures of help seeing

There was no significant difference between conditions with regard to students' performance on the hypothetical help-seeking dilemmas, $F(1, 28) = 1.3$, *ns*. Nor was there a difference between conditions with regard to performance on items with embedded hints, $F(1, 28) = 0$, *ns*.

2.3. Discussion

2.3.1. Validation of the help-seeking model

Study 1 found that online help-seeking behavior according to the help-seeking model correlated with students' performance on posttest items with hints, controlling for performance on items with no hints. In other words, students who applied better help-seeking skills while working with the tutor (according to the help-seeking model) were better able to use the hints embedded in the paper test when controlling for domain knowledge (i.e., controlling for performance on items with no hints). Thus, the metacognitive behavior that was captured by the help-seeking model was consistent across environments. Furthermore, it was found that behavior according to the help-seeking model correlated with students' ability to identify desired help-seeking behavior in the hypothetical dilemmas assessment. The significant correlation between the help-seeking model and two independent measures of help seeking on paper confirms the construct validity of the help-seeking model, in support of Hypothesis 1.

2.3.2. The effect of the help tutor

The results of Study 1 also suggest that students in the Help condition improved several aspects of their help-seeking behavior, in support of Hypothesis 2. Mainly, the Help Tutor led students to limit the use of executive help (that is, to ask for fewer bottom-out hints). At the same time, the frequency of students' help avoidance was not reduced, compared to the control condition. Perhaps the Help Tutor was a victim of the same behavior it attempts to cure, that is, students who tend to avoid using the tutor's help at the domain level also ignored the Help Tutor's recommendations.

Notably, the Help Tutor did not force students into better help-seeking behavior. For example, a student who is determined to view the bottom-out hint is able to do so even if the system recommends against it. Therefore, the effect of the Help Tutor on the more common types of errors suggests that metacognitive feedback can be used to improve students' behavior in tutoring systems. Yet, the present study did not provide evidence that the Help Tutor helped students learn transferable help-seeking skills. Study 2 focused on replicating the effect of the Help Tutor and assessing the transferability of the improved help-seeking patterns.

As mentioned earlier, Study 1 also measured learning at the domain level. Results of domain-level assessments show that while students in both conditions improved significantly from pre to posttest, there was no effect of the Help Tutor on students' learning gains (for complete results see Roll et al, 2006). The lack of effect on domain knowledge raises an important question, namely why the improved help-seeking behavior did not lead to improved learning at the domain level. One possible explanation suggests that the Help Tutor messages were interpreted in the local context in which they were given. Students may have treated them as an additional task, and while they followed their recommendations, they did not engage in deliberate help-seeking. This behavior is sensible, given that domain-level hints are often specific to the problem step at hand. This explanation emphasizes a limitation of Study 1, that is, that knowledge of help seeking was not pitched as a valid learning goal, alongside knowledge of Geometry. In fact, students in the study received no explicit instruction on help seeking, and no connection was made between help-seeking behavior and success in learning the domain. To make this connection more explicit, Study 2 included additional help-seeking support in the form of direct instruction on help seeking. In addition, Study 1 did not stress the general nature of help-seeking skills. To help students realize the wide applicability of these skills, Study 2 included help-seeking support across multiple instructional units.

An alternative explanation for the lack of effect of the Help Tutor on domain-learning gains suggests that the metacognitive feedback imposed excessive cognitive load that interfered with learning the domain knowledge. To reduce cognitive load that is associated with help seeking, Study 2 included an environment in which students could engage in self-assessment and help-seeking behavior outside the Geometry Cognitive Tutor. During these self-assessment episodes students identified the domain-level skills on which they need assistance, thus bridging between the domain-independent help-seeking principles and the relevant domain-specific knowledge.

An additional limitation of Study 1 has to do with the student population that was studied. In the absence of posttest results from the lower-achieving school, the study used only data from the relatively high-achieving school to validate the help-seeking model (Research Question 1). However, given the previously reported high correlation between the help-seeking model and domain-level learning gains, it is likely that the validity of the Help Tutor is not limited to high performing students. Note that data from both schools was used to evaluate the effect of the Help Tutor on students' help-seeking behavior (Research Question 2).

3. Study 2: Evaluating the transferability of the improved help-seeking skills

Study 2 evaluated Hypothesis 3, namely, that students' improved help-seeking behavior will transfer to new content within the Geometry Cognitive Tutor. Hypothesis 3 also predicted that no transfer would be found to paper-and-pencil measures of help-seeking knowledge. As described above, few changes were made to the help-seeking support offered to students. First, a brief in class instruction on help seeking was given to students in the Help condition. The instruction included a 4-min video presentation with annotated examples of productive and faulty help-seeking behavior in the Geometry Cognitive Tutor. The goals of the instruction were threefold: to give students a better declarative understanding of desired help-seeking behavior, to improve students' dispositions towards seeking help, and to frame help-seeking knowledge as an important learning goal.

In addition to the Help Tutor and the help-seeking instruction, students in the Help condition worked with the Self-Assessment Tutor. The Self-Assessment Tutor was developed with multiple goals in mind: to make the connection between adaptive help-seeking and domain-learning salient, to help students get in the habit of self-assessing their ability, to help students evaluate their knowledge level on the target set of skills, and to offer students opportunities to practice their help-seeking skills outside the context of the Geometry Cognitive Tutor. The Self-Assessment Tutor capitalized on the idea that correctly assessing one's own ability was found to correlate with strategic use of help (Tobias & Everson, 2002). The Self-Assessment Tutor makes the self-assessment process explicit and scaffolds it, similarly to the learning process with the Metacognitive Instruction using a Reflection Approach system (MIRA; Gama, 2004). In the Self-Assessment Tutor the student first assesses her ability to solve a prototypical problem from the subsequent Geometry Cognitive Tutor section. She then attempts to solve the problem (on-demand domain-level hints are made available at this stage). Finally the student reflects upon her prior assessment and ability (see Fig. 6). Students used the Self-Assessment Tutor for about 5 min prior to each Geometry Cognitive Tutor section (each of the two units in this study had four sections).

The third and last change from Study 1 to Study 2 was that the Help Tutor in Study 2 was used in two instructional units. Between the two units, as well as after the second unit, students from both conditions used the unmodified version of the Geometry Cognitive Tutor (without the Help Tutor support). The Help Tutor support was extended to two instructional units (and months) to see whether a "double dosage" of help-seeking support would improve transfer, especially that students may identify similar patterns of desired help-seeking behavior across different domain-level units.

3.1. Method

3.1.1. Design

Study 2 spanned a period of four months and four instructional units. During Month 1 students learned about Angles. Month 2 was devoted to review of the previous eleven units in preparation for statewide exams. The topic of Month 3 was Quadrilaterals. During Month 4 students worked on different units of the tutor (mostly triangles and 3-D geometry). As in Study 1, Study 2 compared two conditions: Students in the Help condition received support for their help-seeking behavior, and students in the Control condition worked with the unmodified Geometry Cognitive Tutor as they would have without the study. Unlike Study 1, students in the Help condition had the support "turned off" for part of the study. Help condition students received support only during Months 1 and 3. During Months 2 and 4 they too worked with the unmodified Geometry Cognitive Tutor, similar to students in the Control condition. Comparing the help-seeking behavior of students from both conditions during Months 2 and 4 can be used to evaluate the transferability of the acquired help-seeking skills across domain-level topics within the same environment. Furthermore, comparing students' help-seeking behavior following Month 2 to that following Month 4 can be used to evaluate dosage effect, that is, whether receiving help-seeking support for two instructional units (during Months 1 and 3) yields different transfer results compared with receiving support during a single instructional unit (during Month 1 alone).

3.1.2. Participants

Study 2 took place in a rural vocational high school in Western Pennsylvania (3% minorities, 25% mathematical proficiency). Participants were 67 students, 10th and 11th graders (15–17 years old), who were enrolled in a Geometry Cognitive Tutor course and were accustomed to the tutor's interface. While no exact gender information is available, enrollment in these classes was 69% male to 31% female.

Since the study included teacher-led classroom instruction, assignment to conditions was done at the class level. Two classes (with 29 students) were assigned to the Help condition and two classes (38 students) were assigned to the Control condition. Two teachers taught the classes so that each teacher taught one class in each condition. Classes were assigned in consultation with the teachers, attempting to control for number of students and time of day. One class that was reported by the teachers to be of lower ability was assigned to the Help condition.

3.1.3. Materials

The help-seeking support offered to students in the Help condition during Months 1 and 3 included (a) a classroom instruction on help seeking, (b) the Self-Assessment Tutor, and (c) the Help Tutor.

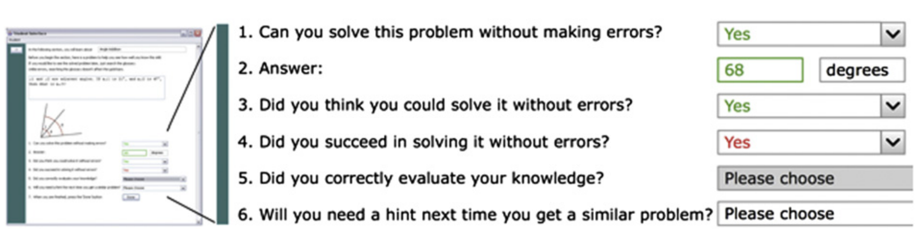

Fig. 6. The Self-Assessment Tutor.

3.1.4. Instruments

Students' behavior while receiving support (during Months 1 and 3) was assessed using the help-seeking model. Students' help-seeking behavior once support from the Help Tutor was no longer available (during Months 2 and 4) was assessed using log files from students' interaction with the unmodified Geometry Cognitive Tutor. This analysis focused on two key aspects of help-seeking behavior in the Geometry Cognitive Tutor. The first is the average hint level. A high average hint level suggests that students often ask for bottom-out hints, implying that they rely on more executive, rather than instrumental, help. Hence, one would expect the Help Tutor to reduce the average hint level students ask to see. The second aspect in the analysis is time to take action. Rapid actions usually suggest guessing or clicking through hints without reading their content (Aleven et al., 2006). Hence, one would expect the Help Tutor to encourage students to act more deliberately.

In addition, following the intervention, students completed a hypothetical help-seeking dilemmas assessment, identical to the one used in Study 1. Detailed descriptions of the domain-level assessments and results are provided elsewhere (Roll, Aleven, McLaren, & Koedinger, 2007a).

3.1.5. Procedure

As described above, Study 2 spanned a period of four months. Students in the Help condition received support during Months 1 and 3, but not 2 and 4. The support included 4 min of help-seeking instruction at the beginning of each unit. However, due to the low quality of the video equipment in the school, students could not notice the details of the instruction. Following the instruction students in the Help condition completed interleaved Geometry Cognitive Tutor + Help Tutor and Self-Assessment Tutor sessions, with the self-assessment sessions taking about 10% of the students' time. Each unit (during Months 1 and 3) included four sections, hence each help-seeking instruction was followed by four cycles of interleaving Geometry Cognitive Tutor and Self Assessment Tutor sections (see Fig. 7). There was no difference between conditions during Months 2 and 4. Students in both conditions worked with their respective versions of the Geometry Cognitive Tutor for two days a week. The remaining three days were used by the teacher to cover the same material in class, as is commonly done in Geometry Cognitive Tutor courses.

3.1.6. Statistical analysis

Analysis of the data was done using similar statistical procedures to the ones described in Study 1. To allow for easy comparisons between different instructional units, Cohen's *d* effect sizes are presented, in addition to raw values. An alpha level of .05 was used in all statistical tests.

3.2. Results

3.2.1. Help-seeking behavior within the tutor

While using the Geometry Cognitive Tutor, students in the Help condition committed significantly fewer help-seeking errors on hints, compared to their counterparts in the Control condition, that is, 13% vs. 20% in Month 1, $t(52) = 4.3$, $p < .001$, Cohen's $d = 1.2$, and 12% vs. 19% in Month 3, $t(56) = 5.5$, $p < .001$, Cohen's $d = 1.47$ (see Table 3). The Help Tutor had no effect on students' help-seeking error rate during solution attempts (e.g., help avoidance), and it increased students' rate of inappropriate glossary searches in Month 1, although this difference was only marginally significant (1% vs. 2%), $t(57) = 1.9$, $p = .06$, Cohen's $d = 0.50$. Due to a technical error, no glossary searches were logged in Month 3.

When asking for hints, Help students did not drill down as much as Control students. Students in the Help condition asked to see the bottom-out hint on only 45% of their hint sequences, compared to 60% of the Control students in Month 1, $t(52) = 2.3$, $p < .03$, Cohen's $d = 0.64$, and on 41% compared to 53% of the Control students during Month 3, $t(56) = 2.2$, $p = .03$, Cohen's $d = 0.59$. Furthermore, the Help Tutor significantly reduced the mean hint level in both months, that is, 3.1 vs. 3.7 in Month 1, $t(41) = 3.5$, $p < .001$, Cohen's $d = 1.10$, and 3.3 vs. 3.9 in Month 3, $t(45) = 4.4$, $p < .001$, Cohen's $d = 1.31$.

With respect to the time taken on the different types of actions, students in the Help condition did not take more time before asking for their first hint, but they took longer while reading hints and before asking for additional ones compared to the Control condition, that is, 6 vs. 5 s in Month 1, $t(51) = 3$, $p < .001$, Cohen's $d = 0.84$, and 8 vs. 6 s in Month 3, $t(56) = 2.2$, $p < .04$, Cohen's $d = 0.59$. Students in the Help condition took significantly more time also before making attempts compared to the students in the Control condition, that is, 16 vs. 14 s in Month 1, $t(65) = 2.5$, $p < .02$, Cohen's $d = 0.62$, and 22 vs. 19 s in Month 3, $t(62) = 2.7$, $p < .01$, Cohen's $d = 0.69$.

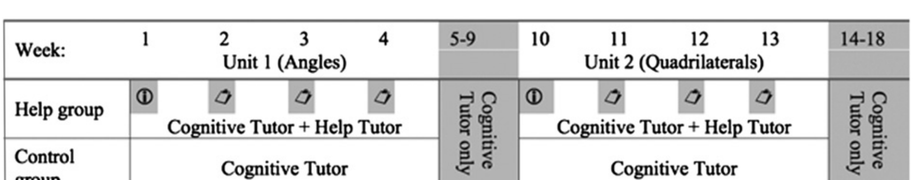

Fig. 7. Procedure of study 2. ⓘ – Declarative help-seeking instruction; [self-assessment icon] – Self-assessment episode. Table is not to scale, that is, self-assessment and instructional activities took about 10% of the class time.

Table 3
Effect sizes (Cohen's d) of differences between the Help and the Control groups in Study 2 as a function of month of tutoring.

Domain-level topic	Month 1: (Help group received help-seeking support) Angles	Month 2: (Both groups worked with the unmodified Cognitive Tutor) Different units	Month 3: (Help group received help-seeking support) Quadrilaterals	Month 4: (Both groups worked with the unmodified Cognitive Tutor) Different units
Measures using the help-seeking model				
Overall help-seeking error rate	0.1 (7% vs. 8%)	N/A	0.4 (13% vs. 14%)	N/A
Hint	1.2** (13% vs. 20%)	N/A	1.4** (12% vs. 19%)	N/A
Try-step	0.3 (7% vs. 8%)	N/A	0.0 (13% vs. 13%)	N/A
Glossary	−0.5 (2% vs. 1%)	N/A	N/A	N/A
% of drilling down to bottom-out hint	−0.6* (45% vs. 60%)	N/A	−0.6* (47% vs. 59%)	N/A
Depth of requested hints				
Average hint level	−1.1** (3.1 vs. 3.7)	0.1 (2.2 vs. 2.1)	−1.3** (3.3 vs. 3.9)	−0.5* (2.2 vs. 2.6)
Time before action				
Ask for first hint	0.2 (18 vs. 16 s)	0.3 (24 vs. 22 s)	0.2 (24 vs. 22 s)	0.2 (24 vs. 22 s)
Ask for additional hint	0.7* (6 vs. 5 s)	0.3 (11 vs. 10 s)	0.6* (8 vs. 6 s)	0.9** (12 vs. 8 s)
Try step	0.6* (16 vs. 14 s)	0.4 (25 vs. 22 s)	0.7** (22 vs. 19 s)	0.3 (24 vs. 22 s)

N/A = not available. $*p < .05$; $**p < .01$.
To support comparisons across units, values reported are effect sizes of the differences between conditions (Cohen's d). Numbers in parenthesis are raw values, Help vs. Control condition.

3.2.2. Transfer of help-seeking skills

The columns labeled Month 2 and Month 4 in Table 3 show data regarding the effect of the Help Tutor when students moved on to work on different units and the Help Tutor's feedback was no longer available (Research Question 3). Seven of the eight measures in these two columns were in the desired direction with the Help Tutor students carrying forward patterns of improved help-seeking behavior into the months with no help-seeking support.

As described in subchapter 3.2.1, the Help Tutor reduced students' average hint level while it was in effect, that is, during Months 1 and 3. As shown in the Month 2 column of Table 3, this effect did not carry over to Month 2, yet was marginally significant in Month 4, that is, mean hint level was 2.2 vs. 2.6, $t(51) = 1.9$, $p = .06$, Cohen's $d = -0.50$.

As described above, Help Tutor students took more time to read hints before asking for additional ones while the Help Tutor support was in effect. No difference in hint reading time was observed during Month 2, following the first use of the help tutor. However, a difference in hint reading time was found after the second unit with the Help Tutor, with students in the Help condition taking more time than students in the Control condition, that is, in Month 4 mean time was 12 vs. 8 s, $t(44) = 3.0$, $p < .01$, Cohen's $d = 0.90$.

Another effect that was observed while working with the Help Tutor was that students in the Help condition took longer than students in the Control condition before attempting a solution. This effect did not persist in the absence of the Help Tutor during Months 2 and 4. The time students took before asking for their first hint was not affected by the Help Tutor, and naturally, no effect was observed in its absence.

Between Months 3 and 4 students completed the hypothetical help-seeking dilemmas assessment. Students in the Help condition performed significantly better on this assessment; specifically, .36 in the Help condition and −.04 in the Control condition, $t(41) = 2.4$, $p = .02$, Cohen's $d = 0.75$ (to reduce guessing, wrong answers received a score of −1).

3.3. Discussion

Study 2 found that the help-seeking support led to an improvement in students' help-seeking behavior. While the support was in effect, students used hints more deliberately (i.e., took more time to request the next level, requested fewer hint levels, and fewer bottom-out hints). The support did not lead to a reduction in help avoidance, and, in contrast to Study 1, led to a marginally significant increase in the rate of errors related to the glossary. However, the low frequency of these errors (1–2% of all glossary searches) makes them less important.

The main finding of Study 2 is that students transferred their improved help-seeking skills to new units of the Geometry Cognitive Tutor where the help-seeking support was no longer in effect. The significant improvement in students' behavior in the unsupported environment was found only during Month 4, and not during Month 2. This finding suggests that students could extract the domain-independent nature of help-seeking skills only after they received support for their help-seeking behavior across two different instructional units (angles and quadrilaterals) and for a longer duration. Also, unlike Study 1, Study 2 found that students in the Help condition transferred their improved help-seeking skills to the hypothetical help-seeking dilemmas assessment. Multiple explanations can account for the transfer of help-seeking skills that was observed in Study 2 (but not in Study 1). First, Study 2 included an instruction on desired help-seeking behavior, delivered via a short video. While the quality of the video was too low for students to notice details, a clear message regarding the importance of help-seeking skills (and their relevance to the Geometry class) was conveyed. The Self-Assessment Tutor may have had a similar effect. Though the self-assessment episodes did not include explicit advice on how to seek help, they directed students' attention to help seeking, and helped them relate the general help-seeking principles to their specific knowledge gaps. Therefore, the transfer of the improved help-seeking skills may be attributed to the interaction between the attention-focusing manipulations (help-seeking instruction and self-assessment episodes) and the detailed feedback of the Help Tutor.

In addition to the described measures, Study 2 found no effect of the help-seeking support on domain learning during Months 1 and 3 (see Roll et al., 2007a for detailed results). Learning gains were not evaluated during Months 2 and 4. The fact that improved help-seeking behavior did not lead to improved learning requires an explanation. One explanation may be that the improvement in help-seeking behavior was not sufficient to measurably impact domain-level learning. However, one would expect to see at least a trend in the test scores. Another possible explanation for the lack of effect on domain learning may be excessive cognitive load when receiving support on both the domain and the metacognitive level. In this case, the Help Tutor may have led to improved learning during the periods in which it was no longer in effect, when students were not subject to the cognitive load imposed by its messages. Domain-learning gains were not measured during Months 2 and 4, and thus this hypothesis cannot be evaluated with the available data. An alternative explanation is detailed in the General discussion.

4. General discussion

The Help Tutor uses the help-seeking model to identify errors students make while working with the Geometry Cognitive Tutor. The help-seeking model was previously shown to capture, in real-time, help-seeking errors that are associated with poor learning (Aleven et al., 2006). Study 1 found that assessing students' help-seeking behavior using the help-seeking model also correlated with other direct measures of help-seeking behavior, namely, the use of hints embedded in paper tests, and identifying correct help-seeking strategies in hypothetical learning dilemmas. These findings support Hypothesis 1 and demonstrate the usefulness of unobtrusive moment-by-moment assessment of help-seeking using a detailed computational metacognitive model. Such assessment enables adaptation of support not only to students' domain knowledge, but also to their metacognitive behavior.

As for the effect of the metacognitive feedback, several aspects of students' hint usage improved significantly while working with the Help Tutor. Most notably, students asked for fewer bottom-out hints, often associated with executive help seeking (Aleven et al., 2006), and took more time to read the tutor's domain-level hints. This improvement on the most common forms of help-seeking errors supports Hypothesis 2. At the same time, no improvement in students' attempts was found (for example, rate of unskilled students guessing rapidly). It may be that students who tend to avoid seeking help at the domain level also do not pay attention to help offered at the metacognitive level.

The effect of the Help Tutor was not limited to the supported environment. Study 2 found that most aspects of the improved help-seeking behavior transferred to new tasks even in the absence of help-seeking support. These improvements consistently reached significance only after an extended period of time of working with the Help Tutor, across two instructional units. It was hypothesized that this effect is the outcome of the interaction between multiple factors. First, since the Help Tutor feedback appeared in two different areas of geometry learning, students could more easily extract the domain-independent help-seeking skills and thus transfer them better to the other subject-matter areas. Second, the manipulation directed students' attention to help-seeking in general (during the help-seeking instruction) and specifically to one's own help-seeking needs (during the self-assessment episodes), thus framing knowledge of help-seeking as a valuable learning goal alongside knowledge of Geometry. The addition of the help-seeking instruction and self-assessment episodes in Study 2 also led to a measurable transfer to a more conceptual task, that is, students' ability to identify desired help-seeking strategies. The evidence for transferable help-seeking skills supports Hypothesis 3. Yet, Hypothesis 3 did not anticipate the more conceptual transfer to the help-seeking dilemmas assessment.

Alongside the improvement to students' help-seeking skills, both studies did not show improvement to students' domain learning while receiving help-seeking support. One explanation is that the Help Tutor, when used in the context of domain learning, imposes excessive cognitive load. A different explanation is that while students in the Help condition were engaged in better learning behaviors, these do not necessarily suffice to yield better learning. The current focus on help seeking may be missing critical elements. Specifically, it may be that learning from hints is harder than perceived. For example, Renkl (2002) showed that automated hints (which he termed "instructional explanations") might hinder students' tendency to self-explain and thus learn. Reevaluating the existing help-seeking literature suggests that learning from given hints may indeed be a rather complex process. Very few experiments have manipulated help seeking in Intelligent Tutoring Systems to date. Those who did often found that interventions that increase reflection yield better learning: Dutke and Reimer (2000) found that principle-based hints are better than operative ones; Ringenberg and VanLehn (2006) found that analogous solved examples may be better than conventional hints; Schworm and Renkl (2002) found that deep reflection questions caused more learning compared with conventional hints; and Baker et al. (2006) showed that auxiliary exercises for students who misuse hints help them learn better. These results suggest that in order to learn from given help, applying the help should be followed by reflection on the process. Additional support to this hypothesis can be found in Shih, Koedinger, and Scheines (2008) who showed that bottom-out hints were found most useful to students who spontaneously self-explained following the hint. In the absence of reflection, hints may not be very useful for learning. Indeed, Baker et al. (2006) found that reducing help abuse (and other gaming behaviors) might not contribute to learning gains by itself, a similar result to the findings presented in this study. While the work described in this article focused on *help seeking*, it may lack a complementary support for *reflection* on the application of the help.

While Studies 1 and 2 demonstrated the benefits of automated metacognitive feedback, they have several limitations. First, Study 2 confounded several manipulations (the Help Tutor, the help-seeking instruction, and the self-assessment episodes). As a result, the relative role of each manipulation cannot be determined. In addition, the studies evaluated help-seeking support within a specific type of tutoring system, that is, coached problem-solving environments. Study 1 found that students' help seeking within the environment correlates with help-seeking behavior in other contexts. Yet, it is not clear whether assessing help-seeking behavior and giving appropriate feedback in more open-ended tasks is doable and useful. While Study 2 found transfer to new domain knowledge, the Help Tutor is yet to pass the ultimate test of a far transfer of the help-seeking skills to different learning environments.

To conclude, the two studies demonstrate that detailed metacognitive models of help seeking can be used by a tutoring system to assess students' moment-to-moment learning behavior unobtrusively. Furthermore, metacognitive feedback, driven by a model of help seeking, can help students improve their behavior in the tutoring system. Last, comprehensive help-seeking support was shown not only to improve students' help-seeking behavior while supported, but also to help students acquire better help-seeking skills in a meaningful way that transfers across topics within the same tutoring environment. This support combined immediate feedback on help-seeking errors, general help-seeking instruction, and self-assessment episodes that bridged between the general principles and the detailed feedback. These results suggest that educational technologies can be used to achieve robust improvement in students' general learning skills. At the same time, they demonstrate that measurable improvement to students' learning behavior does not necessarily translate to immediate benefits at the domain level.

Acknowledgements

We thank Ryan Baker, Jo Bodnar, Ido Jamar, Terri Murphy, Sabine Lynn, Kris Hobaugh, Kathy Dickensheets, Grant McKinney, EJ Ryu, and Christy McGuire for their help. This work was supported by the Pittsburgh Science of Learning Center, which is supported by the National Science Foundation (#SBE-0354420), by a Graduate Training Grant awarded by the Department of Education (#R305B040063), and by a grant from the National Science Foundation (#IIS-0308200).

References

Al-Hilawani, Y. (2003). Measuring students' metacognition in real-life situations. *American Annals of the Deaf, 148*(3), 233–242.

Aleven, V., McLaren, B., Roll, I., & Koedinger, K. (2006). Toward meta-cognitive tutoring: a model of help seeking with a cognitive tutor. *International Journal of Artificial Intelligence and Education, 16*(2), 101–128.

Aleven, V., Stahl, E., Schworm, S., Fischer, F., & Wallace, R. M. (2003). Help seeking and help design in interactive learning environments. *Review of Educational Research, 73*(2), 277–320.

Anderson, J. R. (1993). *Rules of the mind*. Mahwah, NJ: Erlbaum.

Arbreton, A. (1998). Student goal orientation and help-seeking strategy use. In S. A. Karabenick (Ed.), *Strategic help seeking: Implications for learning and teaching* (pp. 95–116). Mahwah, NJ: Erlbaum.

Baker, R. S. J.d., Corbett, A. T., Koedinger, K. R., Evenson, E., Roll, I., Wagner, A. Z., Naim, M., Raspat, J., Baker, D. J., & Beck, J. (2006). Adapting to when students game an intelligent tutoring system. In M. Ikeda, K. D. Ashley, & T. W. Chan (Eds.), *Proceedings of the 8th international conference on intelligent tutoring systems* (pp. 392–401). Berlin: Springer.

Baker, R. S. J.d., Corbett, A. T., Roll, I., & Koedinger, K. R. (2008). Developing a generalizable detector of when students game the system. *User Modeling and User-Adapted Interaction, 18*(3), 287–314.

Dutke, S., & Reimer, T. (2000). Evaluation of two types of online help for application software. *Journal of Computer Assisted Learning, 16*(4), 307–315.

Corbett, A. T., & Anderson, J. R. (2001). Locus of feedback control in computer-based tutoring: impact on learning rate, achievement and attitudes. In J. Jacko, A. Sears, M. Beaudouin-Lafon, & R. Jacob (Eds.), *Proceedings of CHI'2001 conference on human factors in computing systems* (pp. 245–252). New York: ACM.

Gama, C. (2004). Metacognition in interactive learning environments: the reflection assistant model. In J. C. Lester, R. M. Vicario, & F. Paraguaçu (Eds.), *Proceedings 7th international conference on intelligent tutoring systems* (pp. 668–677). Berlin: Springer.

Karabenick, S. A., & Knapp, J. R. (1991). Relationship of academic help seeking to the use of learning strategies and other instrumental achievement behavior in college students. *Journal of Educational Psychology, 83*(2), 221–230.

Karabenick, S. A., & Newman, R. S. (2009). Seeking help: generalizable self-regulatory process and social-cultural barometer. In M. Wosnitza, S. A. Karabenick, A. Efklides, & P. Nenniger (Eds.), *Contemporary motivation research: From global to local perspectives* (pp. 25–48). Goettingen, Germany: Hogrefe & Huber.

Koedinger, K. R., & Corbett, A. T. (2006). Cognitive tutors: technology bringing learning science to the classroom. In K. Sawyer (Ed.), *The Cambridge handbook of the learning sciences* (pp. 61–78). New York: Cambridge University Press.

Luckin, R., & du Boulay, B. (1999). Ecolab: the development and evaluation of a Vygotskian design framework. *International Journal of Artificial Intelligence in Education, 10*(2), 198–220.

MacLeod, W. B., Butler, D. L., & Syer, K. D. (1996, April). *Beyond achievement data: Assessing changes in metacognition and strategic learning*. Paper presented at the Annual Meeting of the American Educational Research Association, New York. http://ecps.educ.ubc.ca/faculty/Butler/Confer/AERA%25201996%2520metacognition.pdf, Retrieved on (13 August, 2010).

Mathews, M., & Mitrovic, T. (2008). How does students' help-seeking behaviour affect learning? In E. Aimeur & B. P. Woolf (Eds.), *Proceedings of the 9th international conference on intelligent tutoring systems* (pp. 363–372). Berlin: Springer.

Mercier, J., & Frederiksen, C. H. (2007). Individual differences in graduate students' help-seeking process in using a computer coach in problem-based learning. *Learning and Instruction, 17*(2), 180–203.

Nelson-Le Gall, S. (1981). Help-seeking: an understudied problem-solving skill in children. *Developmental Review, 1*(3), 224–246.

Newman, R. S. (1994). Academic help seeking: a strategy of self-regulated learning. In D. H. Schunk & B. J. Zimmerman (Eds.), *Self-regulation of learning and performance: Issues and educational applications* (pp. 293–301). Hillsdale, NJ: Erlbaum.

Pintrich, P. R. (2000). The role of goal orientation in self-regulated learning. In M. Boekaerts, P. R. Pintrich, & M. Zeidner (Eds.), *Handbook of self-regulation* (pp. 451–502). San Diego, CA: Academic.

Renkl, A. (2002). Worked-out examples: instructional explanations support learning by self-explanations. *Learning and Instruction, 12*(5), 529–556.

Ringenberg, M. A., & VanLehn, K. (2006). Scaffolding problem solving with annotated, worked-out examples to promote deep learning. In M. Ikeda, K. D. Ashley, & T. W. Chan (Eds.), *Proceedings of the 8th international conference on intelligent tutoring systems* (pp. 624–634). Berlin: Springer.

Roll, I., Aleven, V., McLaren, B. M., & Koedinger, K. R. (2007a). Can help seeking be tutored? Searching for the secret sauce of metacognitive tutoring. In R. Luckin, K. R. Koedinger, & J. Greer (Eds.), *Proceedings of the 13th international conference on artificial intelligence in education* (pp. 203–210). Amsterdam: IOS.

Roll, I., Aleven, V., McLaren, B. M., & Koedinger, K. R. (2007b). Designing for metacognition – applying cognitive tutor principles to the tutoring of help seeking. *Metacognition and Learning, 2*(2), 125–140.

Roll, I., Aleven, V., McLaren, B. M., Ryu, E., Baker, R. S., & Koedinger, K. R. (2006). The help tutor: does metacognitive feedback improve students' help-seeking actions, skills and learning? In M. Ikeda, K. D. Ashley, & T. W. Chan (Eds.), *Proceedings of the 8th international conference on intelligent tutoring systems* (pp. 360–369). Berlin: Springer.

Ryan, A. M., & Shin, H. (2011). Help-seeking tendencies during early adolescence: An examination of motivational correlates and consequences for achievement. *Learning and Instruction, 21*(2), 247–256.

Shih, B., Koedinger, K. R., & Scheines, R. (2008). A response time model for bottom-out hints as worked examples. In R. S. J.d Baker, T. Barnes, & J. E. Beck (Eds.), *Proceedings of the 1st international conference on educational data mining* (pp. 117–126). Montreal: Canada.

Shute, V. J., & Glaser, R. (1990). A large-scale evaluation of an intelligent discovery world: Smithtown. *Interactive Learning Environments, 1*(1), 51–77.

Schworm, S., & Renkl, A. (2002). Learning by solved example problems: instructional explanations reduce self-explanation activity. In W. D. Gray, & C. D. Schunn (Eds.), *Proceedings of the 24th annual conference of the cognitive science society* (pp. 816–821). Mahwah, NJ: Erlbaum.

Tobias, S., & Everson, H. (2002). *Knowing what you know and what you don't: Further research on metacognitive knowledge monitoring*. New York: College Board.

VanLehn, K. (2006). The behavior of tutoring systems. *International Journal of Artificial Intelligence in Education, 16*(3), 227–265.

Wagster, J., Tan, J., Wu, Y., Biswas, G., & Schwartz, D. L. (2007). Do learning by teaching environments with metacognitive support help students develop better learning behaviors? In D. S. McNamara & J. G. Trafton (Eds.), *Proceedings of the 29th annual meeting of the cognitive science society* (pp. 695–700). Austin, TX: Cognitive Science Society.

Wood, H. A., & Wood, D. J. (1999). Help seeking, learning, and contingent tutoring. *Computers and Education, 33*(2), 153–169.