

Towards Computer-Based Tutoring of Help-Seeking Skills

Vincent Aleven

Bruce M. McLaren

Kenneth R. Koedinger

In today's economic and technological environment, individuals continually face the challenge of acquiring new knowledge and skills. To be successful, people must be "intelligent novices" (Mathan & Koedinger, 2003), able to get up to speed quickly in a new domain. Meta-cognitive skills are often regarded as key to being a good learner (Bransford, Brown, & Cocking, 2000; Palincsar & Brown, 1984; White & Frederiksen, 1998). While these skills are crucial in today's society, they are addressed insufficiently in the current educational system.

One such meta-cognitive skill is help seeking. The ability to seek help at appropriate times from appropriate sources and to learn from the received help is important simply because it is not possible for the school system to prepare students for all future skill needs. Also, it is often very difficult to learn a new set of skills by oneself, without any help, say, by reading a textbook and doing the exercises, or by searching the Internet and integrating information found in diverse sources. Typically, it is more effective and efficient to selectively enlist the help of a more experienced individual or even to post a query to a specialized forum or mailing list on the Internet (Keefer & Karabenick, 1998). Developmental psychologists view help seeking as a key strategy for developing independent ability and skills (Nelson-LeGall, 1981). A number of studies have shown that individuals who seek help more effectively have better learning outcomes.

However, there is considerable evidence that many individuals do not seek help effectively (Nelson-LeGall, 1987) or avoid seeking help altogether (Ryan, Gheen, & Midgley, 1998; Ryan, Pintrich, & Midgley, 2001). A number of studies have shown that students with greater prior knowledge of a given domain exhibit more effective help-seeking behavior (e.g., Miyake & Norman, 1979; Nelson-Le Gall, 1987; Nelson-Le Gall, Kratzer, Jones, & DeCooke, 1990; Puustinen, 1998). The troubling consequence of this finding is that those students who are most in need of help are the least likely to get it at appropriate times (Karabenick & Knapp, 1988a). There are social barriers to seeking help, such as fear of being seen as incompetent or not getting full credit for task completion (Nelson-LeGall, 1981) and the expectation that one should solve problems independently (VanderMeij, 1988). Other barriers include inaccurate self-assessment, difficulty in judging when it might pay off to persist in trying to solve a task by oneself rather than seek help, and an inability to formulate good questions. Some forms of help require that one reads and understands “technical” text and evaluates how to apply the general knowledge found in textual sources to the problem at hand, which can be challenging. Finally, many students tend to be performance-oriented (e.g., focused on finishing assigned problems) rather than learning-oriented, thus leading to less effective help seeking (Arbreton, 1998; Karabenick, 2003; Ryan & Pintrich, 1997) and learning (Dweck, 1989).

While help seeking has been studied extensively in social contexts, it has been studied only to a very limited degree in the context of computer-based interactive learning environments (ILEs), even though such systems are becoming increasingly commonplace at many levels of education (Aleven, Stahl, Schworm, Fischer, & Wallace,

2003; Karabenick & Knapp, 1988b). By ILEs we mean a broad range of instructional software including for example computer-assisted instruction (CAI, see Eberts, 1997; Gibbons & Fairweather, 1998; Larkin & Chabay, 1992), intelligent tutoring systems, (Corbett, Koedinger, & Anderson, 1997), authentic problem-solving environments (CTGV, 1997; Slotta & Linn, 2000), systems geared towards the adaptive presentation of instructional multi-media materials (Brusilovsky, 2001; Dillon & Gabbard, 1998), and systems focused on guided discovery learning (de Jong & van Joolingen, 1998). The current chapter focuses on help seeking within one particular type of ILE, namely, Cognitive Tutors (Anderson, Corbett, Koedinger, & Pelletier, 1995). These types of systems are designed to support “guided learning by doing.” They offer a rich problem-solving environment, monitor students as they work through problems in this environment, and provide various forms of guidance including hints and feedback. Cognitive Tutors have been proven to be effective in raising students’ test scores in actual classrooms. As of the spring of 2004 they were in use in approximately 1,700 schools nationwide. Thus they are a prime example of a successful technological educational innovation.

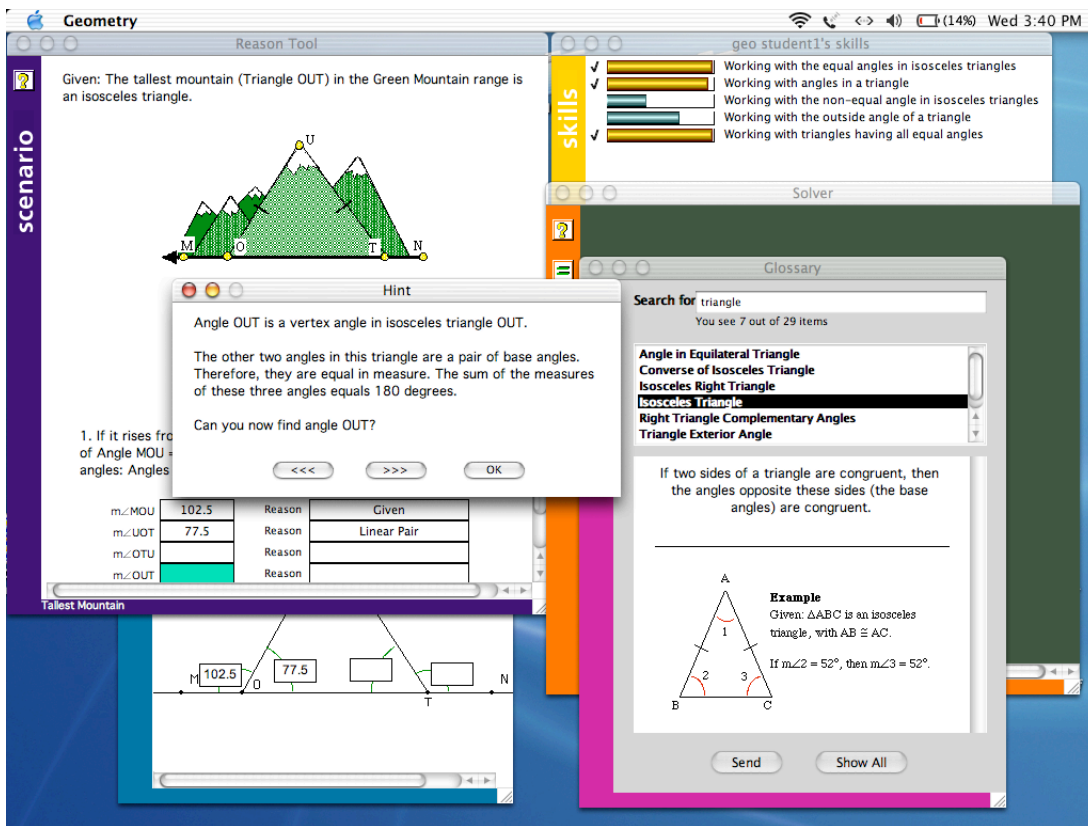


Figure 1: The Geometry Cognitive Tutor

While different ILEs are based on very different pedagogical approaches, a feature common to many systems is on-demand help. For example, the Geometry Cognitive Tutor (Aleven & Koedinger, 2002; Koedinger, Corbett, Ritter, & Shapiro, 2000), shown in Figure 1, offers two types of on-demand help, in addition to other forms of tutorial guidance such as feedback on students’ stepwise problem solutions. We shall call these two types of help *context-sensitive* and *de-contextual* help, reflecting whether or not the help content is tailored, by the system, to the specific learning context. First, at the student’s request, the tutor provides context-sensitive hints with information tailored toward the student’s specific goal within the problem at hand (see the small window

entitled “Hint” in the middle of Figure 1). Typically, multiple levels of hints are available for any given problem-solving step. Second, the tutor offers de-contextual help in the form of a Glossary window, which the student can use to browse a set of relevant problem-solving principles and to selectively display more information about each (see the window labeled “Glossary” at the bottom-right of Figure 1). Other systems provide similar de-contextual help facilities, as well as a variety of others, for example, hyperlinked textbooks, custom-designed hyperlinked background material (Slotta & Linn, 2000), or links to relevant lecture materials (Mandl, Gräsel, & Fischer, 2000)¹.

¹ The distinction between context-sensitive and de-contextual help made above does not align neatly with the difference between *executive* and *instrumental* help seeking that has long been made in the literature (see, e.g., Nelson-LeGall, 1985). Executive help seeking has been defined as focused mainly on supporting performance or completing a task, whereas instrumental help seeking episodes are concerned with the acquisition of new skills or knowledge. Given that an ILE’s goal is to support learning, ideally, the help facilities of any given ILE would channel students into instrumental help seeking and would not allow executive help seeking, or perhaps allow it only as a last resort, to allow students who get stuck on one step at least to make progress on other steps. The help facilities in the Geometry Cognitive Tutor were designed with these goals in mind. The context-sensitive hints sequences are meant to encourage instrumental help seeking by first presenting hints that explain why the answer is the way it is, before actually giving the answer. However, they do not completely disallow executive help seeking: a student could simply ignore the hint levels that provide explanations and pay attention only to the hints that provide answers. The tutor’s de-contextual help, on the other hand, (i.e., the on-line Glossary of geometry knowledge) never provides answers directly and thus does not seem open to such executive help-seeking strategies. Yet there is a cost: when information is not tailored to the problem-solving step at hand, there is more work to do for the student, which may lower the likelihood that the help-seeking episode will be successful or the likelihood that the student will select this type of help. Ultimately, the type of help that tends to be used

on-demand help is a key feature of many ILEs, it is important to understand to what extent students use it effectively (e.g., to seek help in an instrumental rather than executive manner) and to what extent the proficient use of help facilities helps students learn better. Certainly the literature on help seeking in social contexts, as indicated above, has illustrated the impact of proficient help seeking (or lack thereof) on student learning. On the other hand, researchers in the field of ILEs have only just begun to study help seeking as an important influence on student learning (for a recent overview of the literature, see Aleven et al., 2003). Some studies have provided evidence that the proficient use of the help facilities offered by an ILE can lead to higher learning outcomes (Renkl, 2002; Wood, 2001; Wood & Wood, 1999). For example, Renkl (2002) found that adding on-demand help to a system that presents students examples to study leads to better learning outcomes. Wood and Wood (1999) studied help seeking with a small-scale intelligent tutoring system and found that students with lower prior knowledge tended to have higher learning gains if they used help more frequently.

However, those same studies, as well as others, have presented evidence that students often use the help facilities of ILEs in ways that are not conducive to learning (Mandl, Gräsel, & Fischer, 2000). For example, Renkl (2002) identified a group of students with low learning gains who did not use help very frequently. Wood and Wood (1999) found that students with lower prior knowledge were the least likely to use help in an adaptive manner. Unfortunately, such students arguably needed help the most. Aleven and Koedinger (2000), in a classroom study involving the Geometry Cognitive Tutor of

more effectively (i.e., in a more instrumental fashion) is an empirical question. Later in the chapter we present data related to that question.

Figure 1, found evidence of widespread help abuse (e.g., over-use of help) and to a lesser degree of help avoidance. One goal of this chapter is to present data from this study, documenting students' help use with the Geometry Cognitive Tutor. We believe this evidence to be especially compelling, as it comes from actual classroom use of the tutors.

To interpret these results and adequately address the causes of ineffective help seeking in ILEs, it is important to understand both the system's and the student's contribution to the problem. With respect to the system's contribution, it is important to understand the influence on students' learning of various factors in the design of help systems. Only a small number of studies have been done in this area. These studies found that the effectiveness of on-demand help often depends on the content of the help messages given by the system. For example, in a study in which the subjects tried to learn, by exploration, how to use a simple computer drawing program, Dutke and Reimer (2000) found that explanations of what to do were less effective in supporting learning than explanations of the functioning of the drawing program that the students were trying to master (Dutke and Reimer used the term operative help and function-oriented help, respectively)².

² One could view the central question of this study as: "Which type of help is more instrumental, in the Nelson-LeGall (1985) sense, operative help or function-oriented help?" Dutke and Reimer's results show that in a context of learning software by task-oriented discovery, function-oriented help is more instrumental. We note further that the distinction between operative help and function-oriented help is orthogonal to that made earlier between context-sensitive and de-contextual help. The function-oriented help and operative help used in the Dutke and Reimer study is not context-sensitive in the way that the hints of the Geometry Cognitive Tutor are – that is, neither type of help refers to specific information in the

Similarly, studies of the content of feedback messages given by ILEs (i.e., messages that were given at the system's initiative rather than the student's) have shown an effect of such factors as whether the help message draws attention to the students' goals (McKendree, 1990), the degree of interactivity (i.e., the extent to which students are asked to answer questions when receiving system feedback), and whether the help is couched in abstract or concrete terms (Arroyo, Beck, Beal, Wing, & Woolf, 2001; Arroyo, Beck, Woolf, Beal, & Schultz, 2000).

A second goal of this chapter is to present research aimed at understanding how two additional factors related to the help content may have helped shape students' help-seeking behavior with the Geometry Cognitive Tutor, namely, (1) whether the system's help messages focus solely on domain-specific skills and knowledge or whether they also mix in advice at the meta-cognitive level, and (2) the number of levels of help given to the students.

In addition to studying how system design may contribute to good or poor help use, it is important to understand the student's contribution. As mentioned, many factors may contribute to poor help-seeking behavior. Our research focuses on the hypothesis that students would learn more effectively, across a range of domains, if they had better help-seeking skills. From this viewpoint, help seeking is a meta-cognitive skill that must be acquired and that may be open to instruction. This view is consistent with studies that show the success of instructional programs that focus on teaching other types of meta-cognitive skills, such as self-explanation (Bielaczyc, Pirolli, & Brown, 1995),

given problem or explains how a problem-solving principle applies. It is not clear however that greater context-sensitivity would be appropriate in the given discovery learning context.

comprehension monitoring (Palincsar & Brown, 1984), monitoring problem-solving progress (Schoenfeld, 1987), and self-assessment (White & Frederiksen, 1998). In developing instruction on help-seeking skills, we aim to take advantage of the unique qualities of Cognitive Tutors: Given their proven effectiveness in teaching domain-specific skills and knowledge, it is plausible that they could also help students learn to seek help more effectively. We hypothesize that such support for help seeking would help students to learn better at the domain level but also help them to become better help seekers and thus better *future learners*. The question whether an ILE can be effective in fostering meta-cognitive skills such as help seeking has not, to our knowledge, been addressed in the literature on ILEs. One possible exception, in which some interesting but preliminary results were reported, is the work of Luckin and Hammerton (2002) on “meta-cognitive scaffolding.”

The third goal of this chapter is to present ongoing work aimed at extending the Geometry Cognitive Tutor so that it provides tutoring on students’ help-seeking behavior, in addition to assistance with geometry problem solving. The extension to the tutor, called the Help-Seeking Tutor Agent, will be able to evaluate at any point during a tutoring session whether the student could benefit by asking for help, in light of how difficult the given step is estimated to be at that point in time for that particular student. It will comment when the student makes unnecessary or too-rapid help requests and also when the student refrains from asking for help in situations where it would likely be beneficial. This approach will be implemented within the “model tracing” paradigm that underlies the Cognitive Tutor technology (Anderson et al., 1995). In this paradigm, a tutoring system uses a cognitive model, essentially a simulation model of student

thinking that can be executed on the computer, to monitor students' individual approaches to problems and provide guidance as needed. For example, the Geometry Cognitive Tutor employs a cognitive model of problem-solving skill in the domain of geometry to monitor students' solution steps and to provide hints and feedback. Similarly, the Help-Seeking Tutor Agent uses a cognitive model of help seeking to evaluate students' help-seeking behavior. Thus, a key challenge in creating this tutor is the implementation of a detailed model of adaptive help-seeking behavior. We have created an initial model that shares some traits with models of help seeking that have been presented in the literature on help seeking (Nelson-LeGall, 1981; Newman, 1998) but is considerably more detailed. Further, it is a computational model, meaning that it can be executed on a computer to make (real-time) predictions about help seeking and to provide tutoring.

In summary, we will present data about students' use of the help facilities of the Geometry Cognitive Tutor and document the difficulties that students have in using help effectively. We then present progress made toward addressing those difficulties with respect to both solution approaches outlined above, including preliminary data investigating the influence of two factors on students' help use, namely, the mixing of cognitive and meta-cognitive advice, and the number of hint levels. We also present our initial cognitive model of help seeking that will form the basis for extending the Tutor so that it helps students become better help seekers. Finally, we describe an experiment to evaluate whether this tutor actually improves students' help-seeking ability and whether they learn better as a result.

The Geometry Cognitive Tutor

Cognitive Tutors are a form of intelligent instructional software, designed to help students as they learn a complex cognitive skill (Anderson, et al, 1995). This type of tutoring software provides support for learning by doing in the form of context-sensitive hints, feedback, and individualized problem selection. Carefully integrated with classroom instruction, Cognitive Tutors for high-school algebra and geometry have been shown to be about one standard deviation better than traditional classroom instruction (Koedinger, et al., 1997; 2000) and have also been shown to be instrumental in enhancing student motivation (Schofield, 1995). Currently, Cognitive Tutors for Algebra I, Algebra II, and Geometry are disseminated nationwide by a company spun off from our research group (see <http://www/carnegielearning.com>). As of the 2003-2004 school year, the Algebra I Cognitive Tutor is used in 1,500 schools around the country, the Geometry Cognitive Tutor in 350 schools.

The Geometry Cognitive Tutor™, shown in Figure 1, is an integrated part of a full-year high-school geometry course (Alevan & Koedinger, 2002). Both the tutor and the curriculum were developed by our research group. Following guidelines of the National Council of Teachers of Mathematics (NCTM, 1989), the tutor presents geometry problems with a real-world flavor. In these problems, students are given a description of a problem situation plus a diagram and are asked to calculate unknown quantities such as angle measures, segment measures, the area of 2D shapes, etc. The tutor offers a rich problem-solving environment in which students can work with both a tabular and diagrammatic representation of the problem and get assistance from a symbolic equation solver tutor if their solution requires it.

As students enter the values for the unknown quantities into the table (shown on the left in Figure 1), the tutor provides feedback indicating whether the entry is correct or not. In the Angles unit, the unit of the tutor curriculum dealing with the geometric properties of angles (most of the data presented in this chapter pertain to this unit), students are asked also to provide brief explanations for each numeric answer by indicating which geometry theorem or definition justifies it (e.g., “Triangle Sum”). They can either select explanations from the tutor’s online Glossary, described further below, or they can type the explanation into the tutor’s table. Students are allowed to move on to the next problem only when they provide correct answers and explanations for all steps.

As mentioned, the tutor provides two forms of on-demand help. The tutor’s online Glossary lists important theorems and definitions of geometry and illustrates each with a simple example diagram (Figure 1, center). This type of help is de-contextual: the information that is presented is not tailored to the specific problem that the student is working on. The students can use the Glossary freely as they work with the tutor. The Glossary also functions as a menu from which students can select the reasons that justify their steps, saving them the effort of typing the name of the theorem. Given that the de-contextual help does not directly provide answers, one might regard it as instrumental help as defined by Nelson-LeGall (1985). Ultimately, though, whether this type of help (or any type of help) is instrumental in helping students learn is an empirical question, the answer of which depends not only on the help itself but also on whether and how students use it.

Table 1: Example Hint Sequence in the Geometry Cognitive Tutor

Hint Text	General Hint Plan
1. As you can see in the diagram, Angles LGH and TGH are adjacent angles. Together they form line HI. How can you use this fact to find the measure of Angle TGH?	Identify key problem feature (“adjacent angles”)
2. Look in the Glossary for reasons dealing with adjacent angles.	Recommend use of Glossary to look up relevant info
3. Some rules dealing with adjacent angles are highlighted in the Glossary. Which of these reasons is appropriate? You can click on each reason in the Glossary to find out more.	Help reduce Glossary search
4. If two angles form a linear pair, the sum of their measures is 180 degrees. Angle TGH and Angle LGH form a linear pair.	Identify relevant problem-solving principle
5. The measure of Angle TGH plus the measure of Angle LGH is equal to 180 degrees.	} Discuss how principle can be applied
6. The measure of Angle TGH is equal to 180 degrees minus the measure of Angle LGH.	
7. $m\angle TGH = 180^\circ - m\angle LGH$	

Further, the tutor provides context-sensitive hints at the student’s request with information on how to complete the next step in the given problem (Figure 1, window in the middle of the figure). For each step, between 5 to 8 levels of hints are available, depending on the specific geometry skill involved. Each hint level provides increasingly more specific advice (see Table 1). The hints were designed not just to communicate to the student which problem-solving principle is applicable and how it applies, but also to communicate a strategy for how they might find an applicable principle, if they were not able to recall it from memory. The early hints in each sequence suggest that the student undertake a Glossary search using a cue identified in the problem (e.g., that the problem

involves an isosceles triangle or supplementary angles). An intermediate hint highlights a few potentially relevant geometry rules in the Glossary, in case the student was not successful in determining which rule is applicable. The later hints in each sequence discuss how the rule can be applied to find the targeted angle measure. Each hint sequence “bottoms out” by stating a simple algebraic expression that describes how to find the unknown quantity. In other words, the bottom-out hint typically comes very close to giving students the answer for the current step in the problem. Thus, the hint sequences progress from indirect help (hints, really), to direct help. One might say that the earlier help levels provide instrumental help (Nelson-LeGall, 1985), whereas the last level (the bottom-out hints) provide executive (Nelson-LeGall, 1985), or expedient (Butler, 1998), help. As before, we caution that judgments as to whether a particular type of help is instrumental or are better made after the fact, based on data relating help use and learning outcomes.

Like all Cognitive Tutors (Anderson, et al., 1995; Koedinger, et al., 1997), behind the scene the Geometry Cognitive Tutor uses a cognitive model to monitor students as they solve problems. The cognitive model represents the skills targeted in the instruction as a set of production rules, following the ACT-R theory of cognition and learning (Anderson & Lebière, 1998). An example of such a skill would be applying the triangle sum theorem to calculate an unknown angle in a triangle, given the measures of the other two angles in the triangle. The cognitive model is an executable simulation of student problem solving. That is, the system can run the model to solve problems step by step in the same manner that students must learn. This executable cognitive model is key to the tutor’s ability to monitor and evaluate students’ activities, for which it uses an algorithm

called “model tracing.” When the student attempts to answer a tutor question (e.g., enters a numeric value into the tutor’s table), the system runs the model to find out what steps it would take in the same situation and gives feedback to the student by comparing the student’s solution step to those generated by the model. Similarly, when the student requests a hint, the system runs the model to generate possible steps, selects the best one, and then produces hint text using templates attached to the production rule in the model that are involved in generating that step.

Cognitive Tutors keep track of the student’s knowledge growth over time by means of a Bayesian algorithm called knowledge tracing (Corbett & Anderson, 1995). The algorithm is invoked at each problem-solving step to update the tutor’s estimates of the probability that the student knows the skills involved in that step. Whether the estimate is incremented or decremented depends on whether the student was able to complete the step without errors and hints. The tutor uses the estimated probability of skill mastery to select problems and make pacing decisions on an individual basis. It displays these estimates in a “skillmeter” window, shown at the top right in Figure 1, to inform the students about their progress. Skills for which the tutor’s estimate exceeds a certain threshold (set to .95) are considered to be “mastered” and are “ticked off” in the skillmeter window. Once the student has reached mastery for all skills, the tutor advances her to the next section of the curriculum. Students tend to be quite aware of this advancement criterion and often keep a close eye on their skillmeter window

Students’ help-seeking patterns

In this section, we present data about students’ help-seeking behavior while using the Geometry Cognitive Tutor. We assess this behavior by comparing it to reasonable

predictions about that behavior. We also report how various measures of help use correlate with students' learning outcomes.

Data collection

The data presented in this section were collected during a classroom study whose goal it was to assess the added value of having students explain their problem-solving steps (Aleven & Koedinger, 2002). The participants were assigned to two conditions, each working with slightly different versions of the tutor. Students in the Explanation condition explained their problem-solving steps (i.e., justified their steps by citing the relevant geometry rule, as described above). Students in the Problem-Solving condition used a tutor version that did not require that they explain their steps (i.e., the table in the tutor interface did not have boxes for entering “Reasons”). Since in this chapter we are not primarily interested in the effect of having students explain, we report the results of both conditions together, except where noted.

The study took place in the course of regular instruction with the Geometry Cognitive Tutor in a suburban school in the Pittsburgh area. The study involved the students in two periods of a Geometry Cognitive Tutor course taught by the same teacher and his assistant. The students were mostly 10th-graders, that is, 15 and 16-year olds. A total of 41 students completed the experiment. During the study, about 50% of the time was devoted to classroom and small-group activities. The other 50% of the time the students worked on the tutor's Angles unit, one of the six units that make up the full-year curriculum of the tutor. The students completed a pre-test and post-test before and after their work on the tutor. These tests included questions similar to those encountered on the tutor and included transfer problems as well. The pre-test and post-test data indicate that

there were significant learning gains, attributable to the combination of work on the tutor and classroom instruction (Aleven & Koedinger, 2002).

As the students worked with the tutor, detailed logs of the student-tutor interactions were collected. The log data were analyzed to explore patterns of students' use of the tutor's help facilities, that is, the tutor's context-sensitive hints and the tutor's (de-contextual) on-line Glossary.

Use of context-sensitive help

We looked at a number of variables reflecting the students' use of the tutor's on-demand hints. As mentioned, these hints are context-sensitive: they provide information tailored to the particular problem-solving step that the student is working on. We expected that students, to the extent that they were good help seekers, would request help from the tutor when faced with a step that looked unfamiliar or when they had made an error, as indicated by the tutor's feedback, and were not sure how to fix it. We predicted further that good help seekers would use the hints deliberately, that is, read them carefully in order to decide whether they now had enough information to try the step with reasonable confidence or whether they needed to look at a more detailed hint. Finally, we expected that the students, as they became more proficient with a given skill, would request less detailed help from the tutor (i.e., would request to see fewer hint levels).

As it turned out, the students used the tutor's context-sensitive help on 29% of the answer steps and 22% of the explanation steps. By "step" we mean a subgoal in the problem, or equivalently an entry in the tutor's answer sheet. Students requested to see a hint before their first attempt at answering on 12% of the answer steps and 9% of the explanation steps. We do not have a firm basis for making a numerical prediction as to

what the optimal rate of help use *should* be. However, it is clear that help use and errors should to some degree be balanced, that is, should be roughly equal. If students made many errors without asking for help, this would not be good. But neither would it be good if students asked for hints very often while making very few errors. In our data, students made one or more errors on 36% of the numeric answer steps and 37% of the reason steps. Thus, while the frequency of help use was somewhat lower than the error frequency, help use and errors seemed reasonably balanced. After all, it is quite likely that some errors are slips that can be fixed quite easily without requiring assistance from the system. On the other hand, the expectation that students, when faced with an unfamiliar step, would not guess an answer but would request help first was not fully borne out. Regardless of the precise criterion one uses for deeming a step to be “unfamiliar,” it seems unlikely that only one out of every 10 steps would be unfamiliar when students were working in a novel domain. As further evidence of help avoidance, the students had a tendency to resist asking for a hint, even after making multiple errors on a given step. Details can be found in Aleven and Koedinger (2000).

We also looked for evidence that the tutor’s help messages helped to improve students’ performance with the tutor, that is, that they made it easier for students to complete tutor problems (see Table 2). Overall, students made fewer errors on answer attempts that followed a hint request than on steps where their first action was an answer attempt (82% correct v. 62% correct). Further, after a student made one or two errors on a step, asking for help as the next action reduced both the number of subsequent errors as well as the time needed to complete the step (Aleven & Koedinger, 2000). Thus, the help messages aided performance. Similar evidence that the use of context-sensitive help in an

Table 2: The effect of hints on subsequent performance (i.e., performance on answer attempts immediately following a hint)

%Correct after intermediate hint	%Correct after bottom-out hint	%Correct after any hint	%Correct when first action is answer attempt
64%	87%	82%	62%

intelligent tutoring system aids performance was found in other studies as well (Anderson, Conrad, & Corbett, 1989; McKendree, 1990). It does not necessarily follow however that students' help use led to more efficient or effective learning. In order to support that conclusion, one would need to compare learning results obtained with tutor versions with and without hints.

Finally, we looked at a number of variables that relate to how deliberately students used the help messages (see Table 3). The students spent one second or less with as much as 68% of the hints at intermediate levels, meaning that their next action after requesting the hint (often this was the next hint request) occurred in less than one second. By "intermediate levels" we mean all levels but the last. Also, students spent an average of 2.3 seconds per intermediate hint, which seems inadequate to read and interpret it and assess whether one knows enough to enter the answer. Further, on 81% of the steps on which students requested hints, they requested to see all hint levels including the bottom-out hint. As mentioned, the tutor's bottom-out hints stopped just short of handing students the answer. These data point to frequent use of a strategy in which students click their way through all hint levels as quickly as they can until they reach the bottom-out hint. In other words, the majority of the time, the students choose to use the tutor's context-sensitive hints in (what appears to be) highly executive fashion. This "gaming"

behavior has been studied in other work with Cognitive Tutors (Baker, Corbett, & Koedinger, 2004).

Summing up the findings with respect to students' use of the tutor's context-sensitive hints, while the students appeared to use the tutor's context-sensitive help with appropriate frequency and while this type of help assisted students in completing the tutor, on closer scrutiny, students' help use left much to be desired. In particular, the expectation that the students would use the context-sensitive hints in a deliberate manner was not met. Quite the contrary, the students very frequently used hints to get answers quickly, without careful reflection on the answer. In addition, there was evidence that the students avoided hint use at moments in which it would have been appropriate, namely, when faced with unfamiliar steps or after making multiple errors on a step. The observed preoccupation with bottom-out hints is undesirable. In order to learn from being given the right answer (in this case, the bottom-out hint), one needs to construct an explanation of why the right answer is right (Anderson, et al., 1989). The finding that students focused on the bottom-out hints was not consistent with our expectation that students would need increasingly less detailed help as they became more proficient. The finding is also inconsistent with results from an earlier study with a Cognitive Tutor for Lisp Programming, which showed that the average number of hints requested was 1.5 out of 3 hints (Anderson, personal communication).

Use of de-contextual help

Table 3: Data on the deliberateness of students' hint use

Condition	%Bottom-Out	Time per Intermediate Hint	%Intermediate Hints used Undeliberately
Explanation	87%	1.9s	69%
Problem-Solving	79%	2.5s	68%
Overall	81%	2.3s	68%

We also looked at the frequency with which the students used the tutor's de-contextual help, the on-line Glossary of geometry knowledge. As mentioned, the students could browse this resource freely during their work with the tutor. We expected that for any skill targeted in the instruction, good help seekers would gradually rely less on the tutor's on-demand hints and instead would use the Glossary more often. By using the Glossary, they would avoid the negative update of the tutor's estimate of their skill mastery that comes with hint use. Thus, Glossary use would contribute to their goal of getting their skills "checked off" by the tutor.

As it turned out, however, students used the tutor's on-line Glossary on only 2.7% of the numeric answer steps (see Table 4). They used the Glossary more frequently when giving reasons for their answers, namely, on 43% of the reason steps. The large discrepancy is likely to be due to the fact that on reason steps, the Glossary served as a menu: the students could enter reasons by selecting the name of a geometry rule from the Glossary. Much of the Glossary use on reason steps however appeared to be rapid selection, rather than deliberate reading and interpretation of information about geometry. On only 15% of the reason steps did students spend at least one second with at least one Glossary item, a very minimal criterion for deliberate use. It is clearly not possible for a

Table 4: Frequency of use of de-contextual help (percentage of steps on which the on-line Glossary was used). Deliberate use means that the student inspected at least one Glossary item for at least one second.

	All use	Deliberate use only
Numeric Answer Steps	2.7%	2.0%
Reason Steps	43%	15%

student in one second to learn much from a description of a geometry theorem she has not seen before. However, it does seem possible in that amount of time to visually recognize the example diagram.

Clearly, our expectation that the Glossary would be a convenient and useful resource was not borne out. The frequency of Glossary use on numeric answers steps was very low, even if one considers that Glossary use would be appropriate only on steps for which students have some, but incomplete knowledge.

Correlation between help use and learning

In order to study the relation between help seeking and learning outcomes, we focused on the measures derived from the student-tutor interactions (which we shall call “process measures”) listed in Table 5. Some means shown in this table are different from those reported above because Table 5 shows the average of the per-student average whereas the numbers shown earlier are averages over all steps of all students.

Table 5: Process measures related to the student-tutor interaction. All measures pertain to numeric answer steps only. Explanation steps are not included.

Process measure	Description	Mean and st. dev.
Success frequency	Percentage of steps that the student got right without making any errors or requesting any hints.	55 ± 18
Help frequency	Percentage of steps on which student asked for help	27 ± 19
Error frequency	Percentage of steps on which student made one or more errors	36 ± 12
Successful use of help to avoid error	Number of steps where student used help and did not make an error divided by the number of incorrect steps (i.e., steps where the student used help or made an error)	18 ± 14
Use of help after an error	Percentage of incorrect steps on which the students used help after an error was made (i.e., when students make an error, how often do they ask for help)	40 ± 20

It is difficult to predict whether the frequency of help use will correlate positively with learning. On the one hand, one would predict that greater help use would lead to higher learning gains, since students who use help more frequently have more opportunity to benefit from the explanations given in the help messages. On the other hand, frequent help use is likely to be a sign of students' experiencing difficulty during their work on the tutor. It is likely that those who are in trouble more often during training tend to learn less, which would lead to a negative correlation between help seeking and learning. We have no good basis for predicting which of these opposing influences will be stronger. We note that Wood and Wood (1999), in a study involving a small-scale intelligent tutoring system for factoring quadratic expressions, found a positive correlation between help use and learning, for students with lower prior knowledge.

Table 6: Partial correlation coefficients for correlations between process measures and post-test, when prior domain knowledge and pre-test are partialled out

	<i>r</i>	<i>p</i>
Success frequency	.70	.0001
Help frequency	-.61	.0001
Error frequency	-.66	.0001
Successful use of help to avoid error	-.26	ns
Use of help after error	-.54	.005

We considered correlations between process measures and post-test score, with prior geometric knowledge and the pre-test scores partialled out (see Table 6). Prior geometry knowledge was measured by means of a test administered after the student completed the previous tutor unit, which dealt with the area of various 2-dimensional geometric figures. Most of the process measures correlated significantly with the post-test scores. Students who were most successful during training, making fewer errors and completing more steps without errors or hints, tended to do best on the post-test. Students who used help more often tended to do worse, although not if the help was used successfully to solve a step without errors.

Thus, unlike Wood and Wood (1999), we found a negative correlation between help seeking and learning. The negative correlation implies that any learning advantages due to the more frequent use of help messages were not sufficient to enable the more frequent help users, presumably the learners who experienced more difficulties during the learning process, to overcome these difficulties and learn as well as, or better than, students who used help less frequently.

Implications

In short, while there was evidence that students used the tutor's context-sensitive help facility with appropriate frequency and that the help messages helped students in completing problems, there is considerable room for improvement in students' help use. First, there was evidence of widespread "gaming" of the system, with students using the tutor's hints to get to answers as quickly as possible (see also Baker et al., 2004). Second, there was evidence, although not quite as strong, of hint avoidance, that is, of students not requesting a hint when it would likely have been beneficial. Also, the students infrequently used the tutor's on-line Glossary. Under these circumstances, it is hardly surprising that there was a negative correlation between help seeking and learning. Help use aimed at getting answers without reflection on the reasons behind those answers has little potential to improve learning. Therefore, quite possibly, the negative correlation merely reflects the fact that those who are in trouble more often during a learning process tend to have lower learning outcomes.

It is an interesting question why the results reported here differ from those found by Wood and Wood (1999), who as mentioned found a positive correlation between help seeking and learning and did not report the kind of hint abuse that was rampant in our study. There are considerable differences in scale between the studies. While QUADRATIC, the system used by Wood and Wood, comprises about 1-1.5 hours of instruction, the Geometry Cognitive Tutor addresses a full-year high-school geometry course. The data reported here relate to the third of six major units that made up the curriculum at the time of the study, which takes about 8.5 hours of work on the tutor. Prior to the study, the students had already spent at least that amount of time working on

the previous two tutor units. It is possible that certain tendencies in students' help-seeking behavior do not develop until students work on the tutor over an extended period of time. It would obviously take some time and use for students to fully absorb the fact that the bottom-out hint virtually provides the answer to a step. Another significant difference may be the fact that the current study is a classroom study, whereas the Wood and Wood study took place in a lab. Finally, there were a number of differences between the help systems used in the two studies. For example, Wood and Wood had five levels of hints whereas the Geometry Cognitive Tutor has between 5 and 8 levels of hints for any given step. Also, in QUADRATIC, the system used in the Wood and Wood study, the initial hint level (i.e., the hint level given when the student first asks for help on a step) is contingent upon the student's previous performance, whereas in the Geometry Cognitive Tutor, help always starts at the first level. So, for instance, the approach in the Geometry Cognitive Tutor might encourage rapid hint selection since hints are not customized to performance; i.e., many students might find the early hints to be unhelpful. The current data do not allow us to distinguish between these different explanations.

We see two broad possible explanations for the observed ineffective help-seeking behavior. First, the problem may have been mainly with the system. Perhaps the kinds of hints provided by the tutor were too difficult for the population of students who participated in the study, that is, the hints may have been outside their zone of proximal development (Vygotsky, 1978). For example, the hints and Glossary may require greater mathematical reading ability than the given student population has. In addition, it may have been too difficult to find relevant information in the Glossary (a search facility was added only later). One must know some terminology in order to search effectively.

Students may not always have picked up this terminology during classroom instruction. Further, it may have been difficult to interpret information found in the Glossary and judge whether that information is relevant to the problem at hand. Better hints and a Glossary that is easier to search may have helped.

A second set of explanations focuses on the students. The observed poor hint use may reflect a conscious minimum-effort strategy on the part of students or it may be the result of students' being performance-oriented as opposed to learning-oriented (see e.g., Arbretton, 1998; Butler, 1998; Karabenick, 2003; Newman & Karabenick, 2004). Alternatively, the poor help use may be the result of poor meta-cognitive skill. Possibly, these students did not have the habit to try to understand their own solutions to problems or to understand the help that was given. Maybe these students were not in the habit of looking up things they do not know. With some practice, these students may have been able to use the Geometry Cognitive Tutor more effectively to learn skills involved in geometry problem solving. In the remainder of the chapter, we describe on-going research aimed at better understanding and addressing these potential causes.

Investigation of two factors potentially influencing help use

In the current section, we consider two factors related to the content of the tutor's help messages that may have contributed to students' poor help use, namely, (a) the mixing of advice at the cognitive and meta-cognitive level and (b) the length of the hint sequences. The nature of this work is exploratory; we use data that were not collected specifically to isolate the influence of these two factors. The comparisons therefore are not entirely free of confounds, as we note below. Nonetheless, the comparisons are very useful, since they help zero in on hypotheses that are worthwhile to test in a more

carefully controlled study. Without such exploratory work, one is faced with a very wide set of potential hypotheses related to the system design without much guidance as to which ones are most worthy of further study (as argued further in Alevan et al., 2003).

Effect of mixing cognitive and meta-cognitive advice

The hint sequences of the Geometry Cognitive Tutor, as illustrated in Table 1, mix advice at the cognitive and meta-cognitive level. In other words, these hints not only explain how to apply a problem-solving principle to a given problem-solving step, they also advise students on how to find an applicable problem-solving principle, by searching the Glossary. There was a good reason for advising students to consult the Glossary, namely, that the students learn to use a resource that is like many reference resources found in the real world. In retrospect, however, the resulting hint sequences may have become too involved and may have imposed too much cognitive load, perhaps contributing to the observed poor hint use.

We explore this hypothesis by comparing the hint data between the two conditions in our experiment, the Problem-Solving condition and the Explanation condition. The hints in the Problem-Solving condition mixed advice at the cognitive level with advice at the meta-cognitive level to a far lesser extent than the hints in the Explanation condition. Consider the hint sequence shown in Table 1, which was given only to students in the Explanation condition. The second and third hint level provide instructions for using the Glossary. These two hint levels provide advice at the meta-cognitive level. They explain how the student should go about seeking help, rather than helping directly with the problem-solving step at hand. In the corresponding hint sequence given to the students in the Problem-Solving condition, these two hint levels

were omitted. Instead, the following text was added to the first hint of the sequence: “You can look it up in the Glossary.” Thus, the comparison is less than ideal: while the hint sequences in the Problem-Solving condition provide considerably less meta-cognitive advice than those in the Explanation condition, they do provide some. Further, these hint sequences are shorter than those in the Explanation condition, confounding the comparison. Nonetheless, the comparison is useful for exploratory purposes, as argued above

Comparing the hint data between the two conditions (see Table 6), we see that the students in the Problem-Solving condition used bottom-out hints less frequently than students in the Explanation condition (79% of steps with hints v. 87% of steps with hints). The difference is not statistically significant. Further, the students in the Problem-Solving condition spent somewhat more time with the intermediate hints than the students in the Explanation condition (2.5s v. 1.9s) and had a higher rate of correct answers on their attempts that immediately followed an intermediate hint (69% v. 49%). On the other hand, the rate of deliberate use of intermediate hints (as before, defined as the percentage of hint messages that the student examined for more than 1 second) was the same in each condition (68% in the Problem-Solving condition, 69% in the Explanation condition).

Thus, these data provide some support for the notion that the mixing of advice at the cognitive and meta-cognitive level may have led to less deliberate help use. But this mixing was not the only cause, as the “non-mixed” hint sequences were not used in a very deliberate manner, either.

Table 7: Example hint sequence in the Area unit of the Geometry Cognitive Tutor

Hint Text	General Hint Plan
<p>1. The area (A) of a trapezoid can be found using the formula: $A = 1/2 * h * (b + s)$, where h is the height, b is the longer base and s is the shorter base.</p> <p>Which segment in the diagram is the height? Which is the longer base? Which is the shorter base?</p>	<p>Provide problem-solving principle, prompt for mapping of principle to current problem</p>
<p>2. The trapezoid area formula is: $A = 1/2 * h * (b + s)$, where h is the height, b is the longer base and s is the shorter base.</p> <p>In this particular diagram, the height (h) is segment NE, the longer base (b) is segment EV, and the shorter base (s) is segment NA.</p>	<p>Provide problem-solving principle, map to geometric elements of current problem</p>
<p>3. Plug the given values for the height (NE= 304), the longer base (EV= 500), and the shorter base (NA= 207) into the trapezoid area formula: $A = 1/2 * h * (b + s)$</p>	<p>Provide problem-solving principle and show values to be substituted</p>

Impact of shorter hint sequences with less specific bottom-out hints

Another factor that may have invited poor help use is the length of the hint sequences. As mentioned, the hints sequences in the Geometry Cognitive Tutor contained 5-8 hint levels. Possibly, the students may have felt that the tutor should get to the point more quickly, which may have been the root cause of their tendency to skip all hint levels but the last. To test this hypothesis, we have begun to look at data about the student-tutor interactions of different tutor units. For example, we looked at data about student-tutor interactions for the Area unit of the Geometry Cognitive Tutor, which deal with the area of 2D geometric figures. This data set was collected in a different school in a different year, compared to the data from the Angles unit presented earlier.

The hint sequences in the Area unit, illustrated in Table 7, have a different underlying hint plan, compared to those in the Angles unit. First, the hint sequences in the

Area unit do not include any meta-cognitive advice. That is, a problem-solving principle is given without any advice on how in general a student should go about finding relevant knowledge when faced with an area problem. Further, the hint sequences in the Area unit are much shorter than those in the Angles unit: most sequences have 3 levels. The first level provides the problem-solving principle (often, a formula for finding the area of a particular geometric figure), the second level explains how to map the principle to the step at hand, and the third provides the relevant values to substitute into the formula.

The data indicate that the students working on the Area unit used the hints in a much more deliberate manner. Students working on the Area unit:

- requested to see bottom-out hints on 53% of the steps with hints (Angles: bottom-out hints on 81% of answer steps with hints).
- spent an average of 26s per intermediate hint and 22s per bottom-out hint (Angles: 2.3 and 6.2 seconds)
- used only 2.5% of intermediate hints and 10% of bottom-out hints undeliberately, meaning that they looked at them for less than 1 second (Angles: 68% and 54%).

The data show also that in the Area unit, the percentage correct on answer attempts following a hint is lower than the overall percent correct (59% v. 64%). In other words, there is not the same direct evidence as there was in the Angles unit that the hints help performance. Thus, the data overwhelmingly show that the students in the Area unit used the hints in a far more deliberate manner than the students in the Angles unit. There is no evidence of widespread gaming of the system in the Area unit, as opposed to the Angles unit. There is considerably less evidence that the hints helped performance.

Implications

While neither comparison presented in this section was entirely free of confounds (the second comparison involved a different group of students in a different school), taken together the two comparisons strongly suggest that it is worthwhile to investigate whether more effective help-seeking behavior would result in the Angles unit if the hint sequences were shortened and did not mix cognitive and meta-cognitive advice. We plan to test this hypothesis in a controlled experiment with two versions of the Angles unit, the current version and a version with shortened hint sequences (typically 3 levels per skill) that provide only cognitive advice (i.e., no advice on how to search the Glossary).

A model of good help-seeking behavior

The data about student-tutor interactions with the Geometry Cognitive Tutor indicate that there is considerable room for improvement in students' help-seeking strategies. In a second line of research, we plan to extend the tutor with a Help-Seeking Tutor Agent aimed at improving students' help use and help-seeking skills, like the strategy tutor agent proposed in (Ritter, 1997). The Help-Seeking Tutor Agent will give feedback to students on the way they use the tutor's help facilities, so as to get them to use hints and Glossary at appropriate times and with appropriate frequency.

In order to implement the Help-Seeking Tutor Agent, we are developing a (prescriptive) model of help seeking with the tutor. This is no easy matter. The relation between help seeking and learning is not fully understood. Wood and Wood (1999) report examples of individual students whose learning results (working with a computer tutor) seem to be at odds with their help-seeking behavior. For example, they described a

student who seemed initially to be overusing the help facilities, yet ultimately had positive learning outcomes.

As a framework for our inquiry, and as a basis for building a computational model, we have developed a conceptual model of help seeking with a computer tutor, shown in Figure 2. The model is based on our experience and intuition with intelligent tutors – in particular, the Geometry Cognitive Tutor – and is informed by models such as Nelson-LeGall and Newman (Nelson-LeGall, 1981; Newman, 1994); see also Gross and McMullen, 1983 for a more detailed model. Such a model needs to be confirmed empirically and we are actively engaged in doing this, as explained later in the chapter.

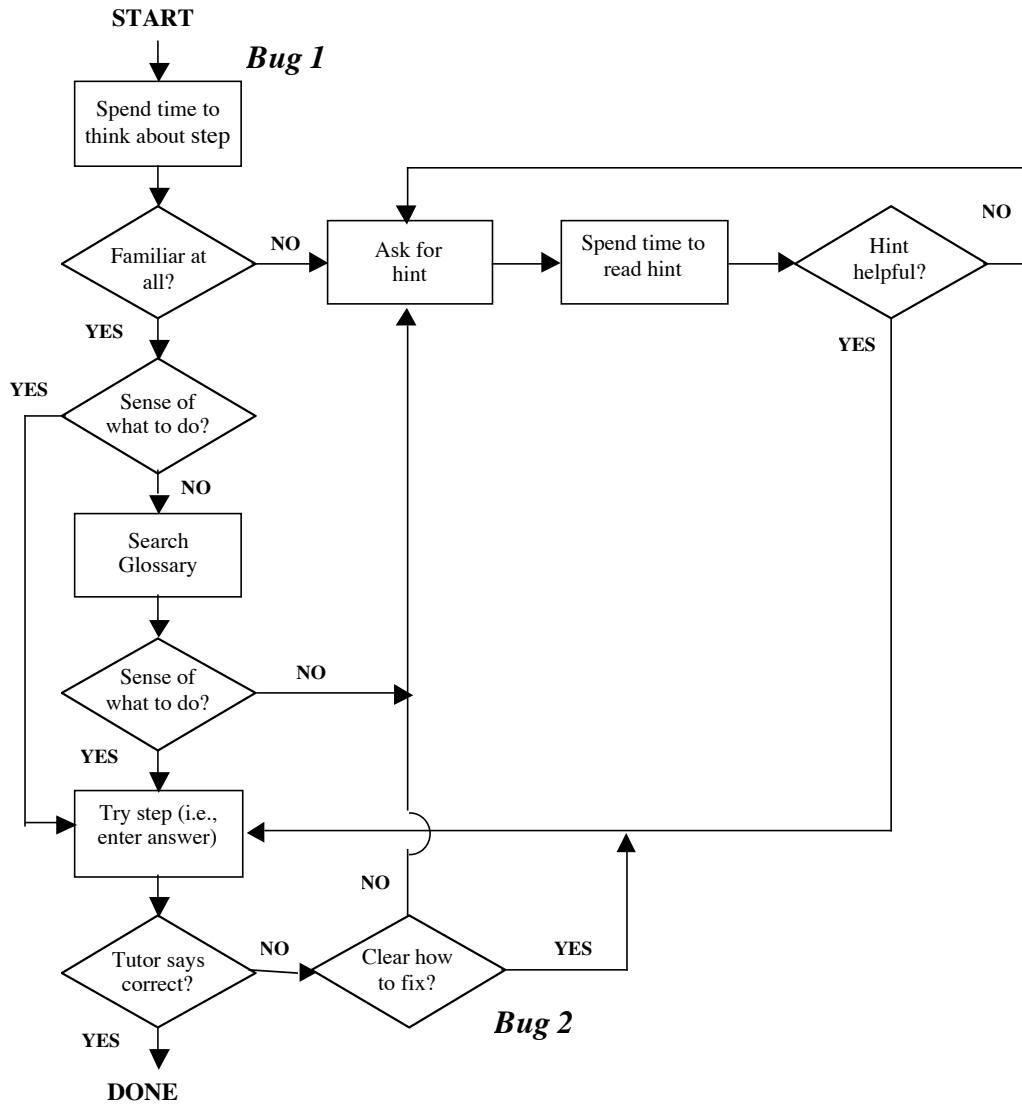


Figure 2: A model of good help-seeking behavior for a student using an intelligent tutor (The “Bug 1” and “Bug 2” labels mark examples of where violations of the model can occur, as will be discussed later in the chapter.)

According to the conceptual model, the ideal student behaves as follows: faced with a step in a tutor problem, the student has three choices: try the step, go the Glossary, or request a hint. If, after spending time considering a problem-solving step, the step looks familiar to the student (i.e., “Familiar at all?” in the flowchart) and she has a good idea of how to solve the step (“Sense of what to do?”), she should proceed with an

attempt to solve (“Try Step”). If on the other hand, the step looks familiar but the student does not have a clear sense of what to do, she should use the Glossary to explore definitions and formulas that may be helpful (“Search Glossary”). After use of the Glossary, the learner should again reflect on whether she has a good sense of how to tackle the step (the second “Sense of what to do?” from the top of the flowchart). If so, she should attempt a solution (“Try Step”); if not, she should ask for help (“Ask for hint”). If from the start the step is not recognizable (i.e., “Familiar at all?” in the flowchart), the student should request help (“Ask for hint”). After reading the hint carefully and deliberately (“Spend time to read hint”), the learner should then decide whether the hint provides enough information to attempt the step (“Hint Helpful?”) or whether another hint is needed at this point (the loop back to “Ask for hint” in the flowchart).

After the student tries a step, the tutor will provide feedback to indicate whether the attempt was correct or not. If the tutor feedback indicates the step is correct, the learner is done with this cycle of help seeking (“DONE”). On the other hand, if the tutor indicates that the step is incorrect, the student should ponder whether it is clear what to do next (“Clear how to fix?”). At this stage the student is expected to either ask for help, if she is unclear what to do (“Ask for Hint”), or attempt to solve the step again (the loop back to “Try Step”).

Thus, the model recommends that students in the initial phase of the acquisition of a skill use hints, then as they become more proficient with the skill, switch to using the Glossary, and finally, try to use the skill without the use of hints or the Glossary. It may not be immediately clear why the Glossary is included in the ideal help-seeking behavior,

since the context-sensitive hints provide more information than can be found in the Glossary and require less effort on the part of the student. The Glossary is included in this strategy because learning to use a resource like the Glossary is useful in its own right. Skills learned with the Glossary may transfer to a wide range of readily available sources of help in the real world that students must learn to use effectively, such as a web browser, a bookshelf with textbooks, manuals, an encyclopedia, the on-line help facilities of many software application packages, etc. These sources are within reach for many people for a large portion of the time. Sources of context-sensitive help on the other hand, while more directly helpful in supporting performance and (perhaps) learning, tend to be more rare and less readily available—the knowledgeable colleague, the teacher, the local guru one floor up, the help line, etc. Some of these sources, for example, the knowledgeable colleague, can be tapped only with limited frequency without running the risk of overburdening the source and possibly losing it as a result. Within the tutor, this rationale is reflected by the fact that the students can use the Glossary freely, but there is a cost associated with the use of context-sensitive hints, at least when hints are requested before an error has been made: As mentioned, the tutor’s knowledge-tracing algorithm increments the estimate of skill mastery only when a step is completed without errors or hints.

Using the help-seeking model to tutor students

In order to build a computational model of the help-seeking flowchart, we had to refine and make concrete some of the abstract components of the flowchart. For example, in order to model the decision points labeled “Familiar at all?” and “Sense of what to do?”, which represent acts of self-assessment by the student, we needed to

implement a test that determines how well a particular student knows a particular skill at a particular point in time. We considered item response theory (Hambleton & Swaminathan, 1985), but determined that it is not a suitable, since it does not track the effect of learning over time. Instead, we decided to rely on information provided by the Cognitive Tutor's Bayesian knowledge-tracing algorithm. As mentioned, for each skill targeted in the instruction, such as applying the triangle sum theorem to find the measure of one of the angles in a triangle, this algorithm computes a probability that the student has mastered the skill, based on the student's performance so far on problem steps that involve the skill (Corbett & Anderson, 1995). Thus, in our model of help seeking, the decisions "Familiar at all?" and "Sense of what to do?" are implemented by comparing the relevant estimate of skill mastery to pre-defined "skill" thresholds. For example, if a student's current estimated level for the skill involved in this step exceeds a probability threshold, currently set to 0.4, our model assumes "Familiar at all?" = YES. For "Sense of what to do?", a more advanced knowledge level than "Familiar at all?", the threshold is set to 0.6.

For the "Clear how to fix?" step of the flowchart, which determines the student's course of action after making an error, our detailed implementation prescribes that a student with a higher estimated skill level for the particular step (i.e., a skill level above the "Sense of what to do?" threshold of 0.6), re-try a step after missing it once, but that mid or lower skilled students (equal to or above and below the "Familiar at all?" threshold of 0.4, respectively) should ask for a hint. In the future we plan to elaborate "Clear how to fix?" by using heuristics that capture ways that students respond to

particular types of common errors. For example, the model may be able to recognize certain types of errors as slips that are easy to fix.

Our implementation of the “Hint Helpful?” step predicts that students with different skill levels for a particular step need different amounts of help. Thus, once a hint is first chosen, a student with high skill for the given step is predicted to need 1/3 of the hints, a mid-skill student 2/3 of the hints, and a low-skill student all of the hints. These are, of course, rough initial estimates of the amount of help required by students with varying skills. We will refine this as we empirically evaluate our computational model. Ultimately, this particular element of the flowchart is about reading comprehension: The extent to which a hint is helpful depends a great deal on how well a student understood the language of the hint. Thus, in the future we anticipate using results from the reading comprehension literature and also empirically evaluate tutor data to estimate the difficulty of understanding the tutor’s hints.

The help-seeking flowchart is also imprecise in prescribing what the student should do in highly complex and multi-step scenarios. For instance, the model is unclear about what a student who repeatedly misses a step or overuses the glossary should do. We’ve addressed this by introducing thresholds that check for a maximum number of solution attempts or glossary searches that are expected. Production rules that appeal to these thresholds have been developed. Thus, for example, a student who exceeds the maximum number of attempts (currently set to 3), while also having seen all the available hints, is expected to ask the teacher for help. When Cognitive Tutors are used in schools, teachers are available to provide help beyond what the Cognitive Tutors can give (see e.g., Schofield, 1995). The option to ask the teacher for help is not made explicit in our

help-seeking model (i.e., there are no production rules that model this step) because such activities could not be monitored by the Help-Seeking Tutor Agent.

The various threshold defaults are intuitively plausible but need to be empirically validated. Also, it is likely we will need more fine-grained thresholds, for instance, different “thinking time” thresholds for high, medium, and low skill students. One of the goals of our empirical study, described below, is to refine and extend the use of thresholds in our model.

The computational model of help seeking

The model described in the previous section forms the basis for a Help Seeking Tutor Agent that will serve as an adjunct to an existing Cognitive Tutor. After each student action with the tutor (e.g., requesting a hint, trying a step, inspecting an item in the Glossary), the Help Seeking Tutor Agent evaluates the action “at the meta-cognitive level,” that is, with respect to its model of help seeking. As long as the student’s help-seeking behavior conforms to the model (i.e., follows the prescribed steps of the flowchart), no feedback is provided. Deviations from the flowchart may lead to feedback from the Help Seeking Tutor Agent.

The current initial version of the model consists of 57 production rules. Four key pieces of information are evaluated each time the help-seeking model is run: (1) whether the student took sufficient time to consider their action, (2) whether the action taken by the student is appropriate given the student’s mastery (or lack thereof) of the skill involved in the step – for example, a low-skill student should ask for a hint, a high-skill student should try the step, etc., (3) what the student has already done with respect to this step (e.g., tried and failed multiple times, viewed all the hints) and (4) if the step was

attempted, whether the student got it right. Thirty-two of the rules are “bug rules,” which reflect deviations of the ideal help-seeking behavior (or “meta-cognitive bugs”). When the student’s behavior matches one of the bug rules, the Help Seeking Tutor Agent provides feedback to the student pointing out the observed help-seeking error.

Example of a student who does not need help

Two examples illustrate how our computational model of help seeking operates. In the first example we show how the model handles a situation in which a student appropriately does not seek help from the tutor. Suppose a student is presented with a step in a problem involving a skill that she masters quite well. The student carefully ponders the step and then attempts a solution. The tutor feedback indicates that the student’s solution step is correct. The student’s action corresponds to the following path through the flow chart of help-seeking behavior shown in Figure 2: (1) “Spend time to think about step” = YES, (2) the YES path from both “Familiar at all?” and “Sense of what to do?,” (3) doing a “Try step, and (4) “Tutor Says Correct?” = YES. Thus, the student’s help-seeking behavior conforms to the model and the Help Seeking Tutor Agent remains silent.

Figure 3 shows the tree of rules explored by the production rule interpreter as it searched for rules matching the student’s behavior in this example case. The nodes of the tree represent individual rules considered by the interpreter as it attempts to match the student’s action to the ideal action prescribed by the model. The “CHAIN” nodes mark choice points where multiple rules are considered. The matching rule chain is shown in blue/gray; in this case, the student’s behavior correctly followed the model. Rules that

were considered but that did not match the observed student action are shown in white boxes.

Thus, the matching rule chain contains four rules: an initial rule that starts the chain (“start-new-metacog-cycle”), a rule that identifies the student as having spent an adequate

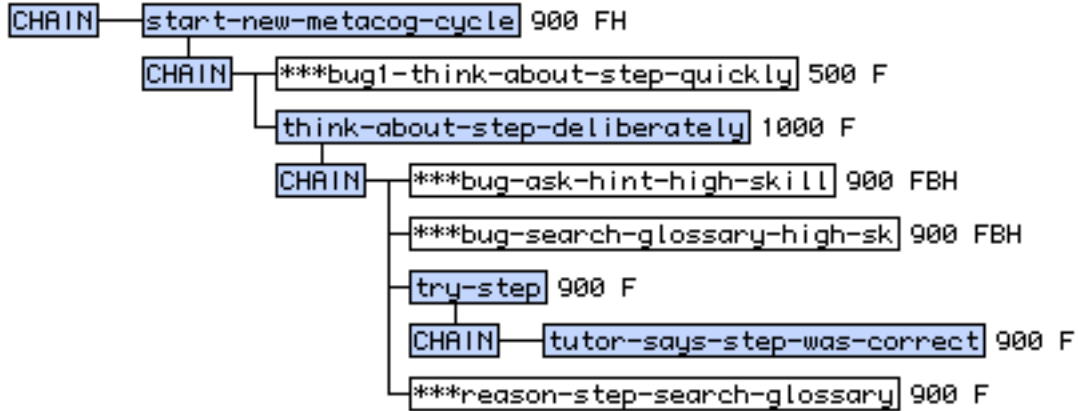


Figure 3: A chain of rules that conforms to the help-seeking model of Figure 2.

amount of time (“think-about-step-deliberately”), a rule that indicates that the student tries the step, as anticipated by her high mastery of the skill at that point in time (“try-step”), and, finally, a fourth rule that indicates that the student got the step right (“tutor-says-step-was-correct”). In this instance, the student has acted as expected and the help-seeking model allows the student to continue with no advice or intervention. This can be seen by the fact that none of the matching rules (i.e., those in blue/gray) are preceded by “bug” –rules that are marked as indicating deviations from the model. This approach of matching a student’s behavior to the ideal model is known as “model tracing” (Anderson et al., 1995). Two of the matching rules are shown in Figure 4.

Rule: think-about-step-deliberately

If the student is engaged in a meta-

Rule: try-step

If the student is engaged in a meta-cognitive

cognitive problem	problem
And the current subgoal is to think about the step	And the current subgoal is to decide what action to take
And the student spent at least min-thinking-time to think about the step	And the students' estimated mastery level for the skill involved in the current step is greater than min-solvable-level
Then	Then
Remove the subgoal (next subgoal is to decide what action to take)	Try step
	Set a subgoal to evaluate the result

Figure 4: Example of rules matching expected behavior by students.

Example of a student who abuses help

Now, consider another student, unlike the high-skill student in the previous example, who is faced with an unfamiliar problem-solving step. Suppose further that our student has already tried and missed the step once before. In a second attempt at this step our student, without spending adequate time thinking, ventures an answer and gets it wrong again. In doing so, the student deviates from the help-seeking model in two ways: she does not spend enough time thinking about the step (a meta-cognitive error marked as “Bug 1” in Figure 2) and in spite of the fact that it is not clear how to fix the problem from the previous step, she does not ask for a hint (marked as “Bug 2”). The combination of these errors amounts to what we call a Try-Step Abuse bug. The bug rules allow the tutor to provide feedback to the student on these meta-cognitive bugs.

Figure 5 shows the tree of rules explored in this meta-cognitive cycle. Note that, unlike the previous example, an initial rule fires signaling that a solution has already been attempted (“consider-next-action”). This rule indicates that in order to interpret the

student action we must not start at the top of the flow chart (indicated by “START” in Figure 2), but rather, we need to start from the decision point immediately after the prior failed solution attempt (“Clear how to fix?”). The next rule indicates the student acted too quickly (“bug1-think-about-step-quickly”), another rule indicates that the student was not expected to try the step, given her low mastery of the skill at that point in time (“bug1-try-step-low-skill”), and, finally, a rule fires indicating that the student missed the step (“bug-tutor-says-step-wrong”). The Help-Seeking Tutor Agent provides the following feedback message, corresponding to the multiple bugs identified in the chain: “Take your time. Think carefully. A hint might help you tackle this difficult step.” The bug rules corresponding to the student acting too quickly and trying the step when they should not have are shown in Figure 6.

Note that getting the step wrong in this situation (i.e., the fact that the answer to the geometry problem is incorrect) is not in itself a deviation from the help-seeking model. At a meta-cognitive level, evaluation of the student’s behavior is focused on whether the student conforms to expected (and ideal) meta-cognitive behavior. A bug at the cognitive level does not necessarily constitute a bug at the meta-cognitive level. However, whether the student gets the step right does matter in the following way: When the student gets a step right, even if there are bugs at the meta-cognitive level, the Help-Seeking Tutor Agent will remain silent.

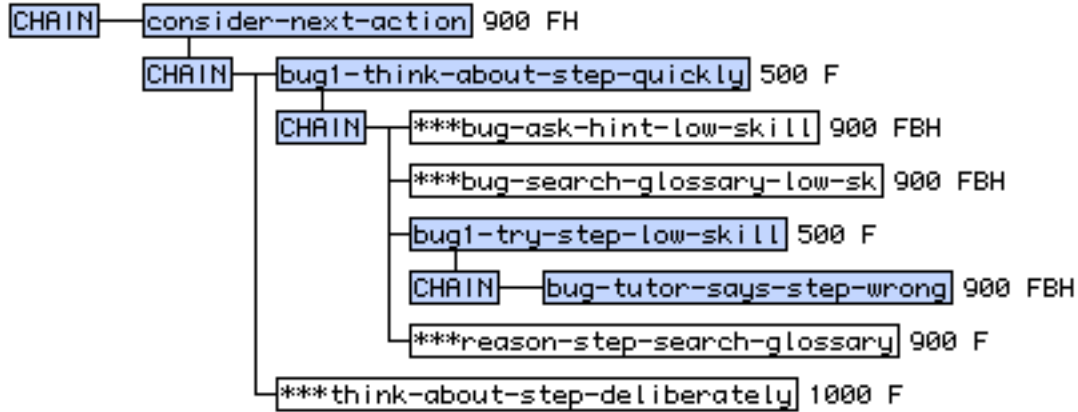


Figure 5: A chain of rules in the Meta-Cognitive Model that deviate from the ideal help-seeking model of Figure 2.

The help-seeking model uses information passed from the cognitive model to perform its reasoning. For instance, the particular skill involved in a step, the estimated mastery level of a particular student for that skill, the number of hints available for that step, and whether or not the student got the step right, are passed from the cognitive to the meta-cognitive model. Meta-cognitive model tracing takes place after cognitive model tracing. In other words, when a student inputs a value to the tutor, that value is first evaluated at the cognitive level before it is evaluated at the meta-cognitive level.

Rule: Bug1-think-about-step-quickly

If the student is engaged in a meta-cognitive problem

And the current subgoal is to think about the step

And the student spent less than min-thinking-time to think about the step

Then

Rule: Bug1-try-step-low-skill

If the student is engaged in a meta-cognitive problem

And the current subgoal is to decide what action to take

And the students' estimated mastery level for the skill involved in the current step is less than min-familiarity-level

Remove the subgoal (next subgoal is to decide what action to take)

And the student has not seen all the hints yet for the current step

Then

Try step

Set a subgoal to evaluate the result

Bug message: “Take your time. Think carefully.”

Bug message: “A hint might help you tackle this difficult step.”

Figure 6: Example of bug rules matching unproductive help-seeking behavior by students.

An important consideration in the development of the Help-Seeking Tutor Agent was to make it modular and useable in conjunction with a variety of Cognitive Tutors. The Help-Seeking Tutor Agent will be a plug-in agent applicable to a wide range of Cognitive Tutors with limited customization. We have attempted to create meta-cognitive rules that are applicable to *any* Cognitive Tutor, not to a specific tutor. Certainly, there will be some need for customization, as optional supporting tools (of which the Glossary is but one example) will be available in some tutors and not others, but many elements of the model will be shared across tutors.

Comparison with Nelson-LeGall’s and Newman’s model of help seeking

Our approach to help seeking has some overlap and similarity to earlier research focused on help seeking in social contexts, such as classrooms. In this section we

compare our model with the general steps in the help-seeking model originally presented by Nelson-LeGall (1981) and later elaborated by Newman (1994):

1. Become aware of a need for help
2. Decide to seek help
3. Identify potential helper(s)
4. Use strategies to elicit help
5. Evaluate help-seeking episode

It is informative to consider the similarities and differences between the Nelson-LeGall model and the help-seeking model presented in this chapter. In the Nelson-LeGall / Newman model, a learner first must become aware that the task is difficult or that she is stuck and in need of help. The ability to assess task difficulty, monitor task progress, and evaluate one's own comprehension and knowledge are important meta-cognitive functions (Nelson LeGall, 1981; Newman, 1998). In the next step of the Nelson-LeGall / Newman model, learners must consider all available information and decide whether to seek help. In our help-seeking model, student awareness of the difficulty of the task and need for help is represented by the ideal help-seeking student asking the questions "Familiar at all?" and "Sense of what to do?" In other words, our model, like the Nelson-LeGall / Newman model, assumes that students with good meta-cognitive skills will monitor their knowledge and be aware of situations in which help is needed. As described above, our model provides a very precise criterion to determine whether to seek help. The Nelson-LeGall / Newman model encompasses a broad range of factors (besides self-assessment of skill level) that influence the decision to seek help (e.g., fear of embarrassment) but models them in less detail.

In the next step, the learner must find a suitable helper. In most classrooms, the teacher or a fellow student could serve this role. In a social context, selecting the helper may depend on the age of the learner and include the perceived competence and sensitivity of the helper (Nelson-LeGall, 1981; Knapp & Karabenick, 1988). In our help-seeking model, finding a suitable helper entails the student's determining which of the available software help facilities is most appropriate in the current context. In the case of the Geometry Cognitive Tutor, this amounts to a choice between the context-sensitive hints and the Glossary. As mentioned, the choice depends on the student's mastery of the skill involved in the step. A good help seeker will opt for context-sensitive hints early on, when she deems her mastery of the given skill to be low and for the Glossary when the step is familiar. In the worst case (e.g., repeatedly missing a step, although all hints have been seen), our model assumes the student will seek help from the teacher. This is not depicted explicitly in the model of Figure 2, but it is part of the computational version of the model.

Next in the Nelson-LeGall/Newman model, the learner must decide *how* to request help, influenced by their knowledge and skills of discourse (Newman, 1998). Essentially, the request must match the task demands. There is no real counterpart to this step in our model of help seeking. Since the student is working with well-defined software functions with specific means to request help, there is no need in our model to consider how to request help.

Finally, in the last step of the Nelson-LeGall / Newman model, the learner reflects upon the help-seeking event to decide if it was helpful and to determine whether further help is required. In our model this behavior is captured by the student asking "Hint

helpful?” after receiving a hint and by asking “Sense of what to do?” after using the Glossary.

In summary, there is overlap between the Nelson-LeGall / Newman model and our model of help seeking. Perhaps the most significant divergence between the models is in the level of specificity. Our model is grounded within an intelligent tutor with specific help functions available to the student. Thus our model is commensurately more detailed but also more bound to a specific tutoring context. The Nelson-LeGall / Newman model does not aspire to the same level of detail or specificity nor does it aspire to computational rendering. Ultimately, we intend to empirically test our model with real students in connection with a variety of actual Cognitive Tutors.

Planned evaluation studies of the Help-Seeking Tutor Agent

Our goal is to both (a) improve the current model of help seeking and (b) use it for tutoring students to be better help seekers. We plan to do empirical work to achieve both objectives. First, we plan to refine the initial model presented above, as we examine data about student tutor/interactions. For example, various elements of the model may need to be adjusted or refined (for example, the skill threshold discussed above) if it turns out that the model does not distinguish well between the productive and unproductive help-seeking behavior or that students deviate so often from the model that it would not be feasible for the Help Seeking Tutor Agent to comment on all deviations.

To assist in the evaluation of the frequency with which students deviate from the model, we are developing a *bug taxonomy* that characterizes the range of help-seeking bugs the model can produce. The taxonomy will enable us to aggregate help-seeking problems in broader categories than the fine-grained level of individual bug rules. As an

example, the bug “Help Abuse” category, intended to cover general situations in which the student misuses hints provided by the Cognitive Tutor, has a sub-category called “Clicking Through Hints” which occurs when a student rapidly proceeds from one hint to the next, without sufficient time to read the hints. Another sub-category “Ask Hint when Skill indicates Trying” characterizes situations in which a student with sufficient skill asks for a hint instead of giving the step a try. Bug instances are categorized in the taxonomy by their specific sequence of rule firings (assuming at least one “bug rule” fires, such as the rule firings illustrated in Figure 5). One useful way we intend to use the taxonomy is to process actual logs of student-tutor interactions (such as those from which the data presented earlier in this chapter were extracted) and determine the percentage of bugs that manifest in each category. This processing can be done automatically, by running the help-seeking model to evaluate the student actions stored in the logs. The resulting data will inform us as to which categories of bugs are most urgently in need of tutoring and which rarely occur. It will also help us to calibrate the various tests and thresholds in the model.

To evaluate whether the model accurately distinguishes between productive and unproductive help-seeking behavior, we will study whether meta-cognitive behavior that conforms to the model correlates with learning outcomes. We will study also how the various bug categories relate to learning. To a degree, however, the proof will be in the pudding. That is, we will know that the model is accurate when the Help-Seeking Tutor Agent actually helps students to learn better. We plan to evaluate whether this is the case in a controlled experiment in which we compare students’ learning outcomes obtained

with a Cognitive Tutor with the Help-Seeking Tutor Agent “plugged in” against those obtained with the same Cognitive Tutor running without the Help-Seeking Tutor Agent.

The planned experiment evaluating the added value of the Help-Seeking Tutor Agent will indicate whether there is a causal relation between help seeking and learning in an ILE. We know of only one previous study that has addressed the causal relationship, namely, Renkl’s (2002) study, in which he compared two versions of a system for example studying, one with and one without on-demand help. This study did find a causal relation. Our study will be different in that it looks at the causal relation within a problem-solving context and will be carried out in an actual classroom. It is different also in that it includes an intervention aimed at improving students’ help-seeking behavior. Further, we plan to address a key question, namely, whether students will be better *future* learners as a result of being tutored on help seeking. The Cognitive Tutors provide a convenient platform to investigate such questions, since students use them in yearlong courses. Thus, we plan to study students’ use of help throughout a course to determine if it improves by exposure to the Help-Seeking Tutor.

Conclusion

Many ILEs use on-demand help. As these types of systems are rapidly becoming widespread, it is important to understand the factors influencing the effectiveness of help use and their relation with learning. Examining data from one of the units of the Geometry Cognitive Tutor, we found that widespread gaming of the system (i.e., the use of hints to find answers without taking time to reflect on why the answer is what it is) leads to a negative correlation between help seeking and learning, clearly a sign that the tutor’s on-demand hints are not as effective as intended.

To improve the situation, we plan improve the hints themselves, for example by shortening the hint sequences and not mixing cognitive and meta-cognitive advice. More importantly, we are developing the Help-Seeking Tutor Agent, which will provide guidance with respect to students' help-seeking behavior. An initial computational model of help seeking we have developed provides a precise notion of good help-seeking behavior. Using this model to monitor students' help-seeking behavior, the Help Seeking Tutor will catch gaming and other unproductive help-seeking behavior. This approach is especially promising in light of the proven effectiveness of Cognitive Tutors, in teaching domain-specific skills and knowledge: When used as an adjunct to classroom instruction, the Cognitive Tutors lead to better learning than typical classroom instruction. It is hypothesized that the Cognitive Tutor approach will have similar advantages when used for instruction at the meta-cognitive level. Our planned evaluation will focus on the hypothesis that the Help Seeking Tutor will help students to learn better at the domain level, even after the support for help seeking has been removed.

Acknowledgements

The research presented in this chapter is sponsored by NSF Award No. IIS-0308200. Previously, it was sponsored by an NSF grant to the Center for Interdisciplinary Research on Constructive Learning (CIRCLE), a research center located at the University of Pittsburgh and Carnegie Mellon University. See <http://www.pitt.edu/~circle>.

References

- Aleven, V., & Koedinger, K. R. (2002). An Effective Meta-cognitive Strategy: Learning by Doing and Explaining with a Computer-Based Cognitive Tutor. *Cognitive Science*, 26(2), 147-179.
- Aleven, V. & Koedinger, K. R. (2000). Limitations of student control: Do students know when they need help? In G. Gauthier, C. Frasson, & K. VanLehn (Eds.), *Proceedings of the 5th International Conference on Intelligent Tutoring Systems, ITS 2000* (pp. 292-303). Berlin: Springer Verlag.
- Aleven, V., McLaren, B., Roll, I., & Koedinger, K. (2004). Toward tutoring help seeking: Applying cognitive modeling to meta-cognitive skills. In J. C. Lester, R. M. Vicario, & F. Paraguaçu (Eds.), *Proceedings of the 7th International Conference on Intelligent Tutoring Systems, ITS 2004* (pp. 227-239). Berlin: Springer Verlag.
- 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., Conrad, F. G., & Corbett, A. T. (1989). Skill Acquisition and the LISP Tutor. *Cognitive Science*, 13, 467-505.
- Anderson, J. R., Corbett, A. T., Koedinger, K. R., & Pelletier, R. (1995). Cognitive Tutors: Lessons Learned. *The Journal of the Learning Sciences*, 4, 167-207.
- Anderson, J. R., & Lebière, C. (1998). *The Atomic Components of Thought*. 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.
- Arroyo, I., Beck, J. E., Beal, C. R., Wing, R., & Woolf, B. P. (2001). Analyzing students' response to help provision in an elementary mathematics intelligent tutoring system. In R. Luckin (Ed.), *Papers of the AIED-2001 Workshop on Help Provision and Help Seeking in Interactive Learning Environments* (pp. 34-46).
- Arroyo, I., Beck, J. E., Woolf, B. P., Beal, C. R., & Schultz, K. (2000). Macro-adapting Animalwatch to gender and cognitive differences with respect to hint interactivity and symbolism. In G. Gauthier, C. Frasson, and K. VanLehn (Ed.), *Proceedings of the 5th International Conference on Intelligent Tutoring Systems, ITS 2000* (pp. 574-583). Berlin: Springer.
- Baker, R. S., Corbett, A. T., & Koedinger, K. R. (2004). Detecting Student Misuse of Intelligent Tutoring Systems. In J. C. Lester, R. M. Vicario, & F. Paraguaçu (Eds.), *Proceedings of the 7th International Conference on Intelligent Tutoring Systems, ITS 2004* (pp. 531-540). Berlin: Springer Verlag.
- Bielaczyc, K., Pirolli, P. L., & Brown, A. L. (1995). Training in Self-Explanation and Self-Regulation Strategies: Investigating the Effects of Knowledge Acquisition Activities on Problem Solving. *Cognition and Instruction*, 13, 221-252.
- Bransford, J. D., Brown, A. L. & Cocking, R. R. (2000). *How people learn: Brain, mind, experience, and school*. Washington: National Academic Press. Bransford, J. D., Brown, A. L. & Cocking, R. R. (2000). *How people learn: Brain, mind, experience, and school*. Washington: National Academic Press.

- Brusilovsky, P. (2001). Adaptive hypermedia. *User Modeling & User-Adapted Interaction, 11*, 87-110.
- Butler, R. (1998). Determinants of help seeking: Relations between perceived reasons for classroom help-avoidance and help-seeking behaviors in an experimental context. *Journal of Educational Psychology, 90*, 630-643.
- Corbett, A. T., & Anderson, J. R. (1995). Knowledge tracing: Modeling the acquisition of procedural knowledge. *User Modeling and User-Adapted Interaction, 4*, 253-278.
- Corbett, A. T., Koedinger, K. R., & Anderson, J. R. (1997). Intelligent tutoring systems. In M. Helander & T. K. Landauer & P. Prabhu (Eds.), *Handbook of human-computer interaction, second, completely revised edition* (pp. 849-874). New York: Elsevier Science Publishers.
- CTGV (Cognition and Technology Group at Vanderbilt). (1997). *The Jasper Project: Lessons in Curriculum, Instruction, Assessment, and Professional Development*. Mahwah, N.J.: Erlbaum.
- Dillon, A., & Gabbard, R. (1998). Hypermedia as an educational technology: A review of the quantitative research literature on learner comprehension, control, and style. *Review of Educational Research, 68*, 322-349.
- Dutke, S., & Reimer, T. (2000). Evaluation of two types of online help information for application software: Operative and function-oriented help. *Journal of Computer-Assisted Learning, 16*, 307-315.
- Dweck, C. S. (1989). Motivation. In: A. Lesgold & R. Glaser (Eds.), *Foundations for a psychology of education* (pp. 87-136). Hillsdale, NJ,: Erlbaum.

- Eberts, R. E. (1997). Computer-based instruction. In M. Helander & T. K. Landauer & P. Prabhu (Eds.), *Handbook of human-computer interaction, second, completely revised edition* (pp. 825-847). New York: Elsevier Science Publishers.
- Gibbons, A. S., & Fairweather, P. G. (1998). *Computer-based instruction: Design and development*. Englewood Cliffs, NJ: Educational Technology Publications.
- Gross, A. E., & McMullen, P. A. (1983). Models of the help-seeking process. In J. D. Fisher, N. Nadler, & B. M. DePaulo (Eds.), *New directions in helping* (Vol. 2, pp. 45-61). New York: Academic Press.
- Hambleton, R. K., & Swaminathan, H. (1985). *Item Response Theory: Principles and Applications*. Boston: Kluwer.
- de Jong, T. & van Joolingen, W. R. (1998). Scientific Discovery Learning With Computer Simulations of Conceptual Domains. *Review of Educational Research*, 68(2), 179-201.
- Karabenick, S. A. (2003). Seeking help in large college classes: A person centered approach. *Contemporary Educational Psychology*, 28, 37-58.
- Karabenick, S. A., & Knapp, J. R. (1988a). Help seeking and the need for academic assistance. *Journal of Educational Psychology*, 80, 406-408.
- Karabenick, S. A., & Knapp, J. R. (1988b). Effects of computer privacy on help seeking. *Journal of Applied Social Psychology*, 18, 461-472.
- Keefer & Karabenick (1998). Help Seeking in the Information Age. In S. A. Karabenick (Ed.), *Strategic help seeking: Implications for learning and teaching* (pp. 219-250). Mahwah, NJ: Erlbaum.

- Knapp, J. R., & Karabenick, S. A. (1988). Incidence of formal and informal academic help seeking in higher education. *Journal of College Student Development, 29*, 223-227.
- Koedinger, K. R., Anderson, J. R., Hadley, W. H., & Mark, M. A. (1997). Intelligent Tutoring Goes to School in the Big City. *International Journal of Artificial Intelligence in Education, 8*, 30-43.
- Koedinger, K. R., Corbett, A. T., Ritter, S., & Shapiro, L. (2000). Carnegie Learning's Cognitive Tutor™: Summary Research Results. White paper. Available from Carnegie Learning Inc., 372 N. Craig Street, Suite. 101, Pittsburgh, PA 15213, info@carnegielearning.com, <http://www.carnegielearning.com>.
- Larkin, J. H., & Chabay, R. W. (Eds.). (1992). *Computer-assisted instruction and intelligent tutoring systems: Shared goals and complementary approaches*. Hillsdale, NJ: Lawrence Erlbaum.
- Luckin, R., & Hammerton, L. (2002). Getting to Know Me: Helping Learners Understand their Own Learning Needs through Metacognitive Scaffolding. In S. A. Cerri, G. Gouardères, & F. Paraguaçu (Eds.), *Proceedings of Sixth International Conference on Intelligent Tutoring Systems, ITS 2002* (pp. 759-771). Berlin: Springer Verlag, 2002.
- Mandl, H., Gräsel, C. & Fischer, F. (2000). Problem-oriented learning: Facilitating the use of domain-specific and control strategies through modeling by an expert. In W. J. Perrig & A. Grob (Eds.), *Control of Human Behavior, Mental Processes and Consciousness* (pp. 165-182). Mahwah: Erlbaum.

- Mathan, S., & Koedinger, K. R. (2003). Recasting the feedback debate: Benefits of tutoring error detection and correction skills. In U. Hoppe, F. Verdejo, & J. Kay (Eds.), *Proceedings of the 11th International Conference on Artificial Intelligence in Education, AI-ED 2003*. Amsterdam: IOS Press.
- McKendree, J. (1990). Effective Feedback Content for Tutoring Complex Skills. *Human Computer Interaction*, 5, 381-413.
- Miyake, N., & Norman, D. A. (1979). To ask a question, one must know enough to know what is not known. *Journal of Verbal Learning and Verbal Behavior*, 18, 357-364.
- NCTM (1989). Curriculum and Evaluation Standards for School Mathematics. National Council of Teachers of Mathematics. Reston, VA: The Council. See also <http://standards-e.nctm.org/index.htm>.
- Nelson-Le Gall, S. (1981). Help-seeking: An understudied problem-solving skill in children. *Developmental Review*, 1, 224-246.
- Nelson-Le Gall, S. (1985). Help-Seeking Behavior in Learning. *Review of Research in Education*, 12, 55-90.
- Nelson-Le Gall, S. (1987). Necessary and unnecessary help seeking in children. *Journal of Genetic Psychology*, 148, 53-62.
- Nelson-Le Gall, S., Kratzer, L., Jones, E., & DeCooke, P. (1990). Children's self-assessment of performance and task-related help seeking. *Jrnl. of Experimental Child Psychology*, 49, 245-263.

- Newman, R. S. (1998). Adaptive help seeking: A role of social interaction in self-regulated learning. In S. A. Karabenick (Ed.), *Strategic help seeking. Implications for learning and teaching* (pp. 13-37). Mahwah: Erlbaum.
- Newman, R. S. (1994). Adaptive 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. 283-301). Hillsdale: Erlbaum.
- Newman, R. S., & Karabenick, S. A. (2004, April). *Achievement goal theory and students' help-seeking orientations: A test of competing models*. Annual meeting of the American Educational Research Association, San Diego, CA.
- Palincsar, A. S., & Brown, A. L. (1984). Reciprocal teaching of comprehension-fostering and comprehension-monitoring activities. *Cognition and Instruction*, 1, 117-175.
- Puustinen, M. (1998). Help-seeking behavior in a problem-solving situation: Development of self-regulation. *European Journal of Psychology of Education*, 13(2), 271-282.
- Renkl, A. (2002). Learning from worked-out examples: Instructional explanations supplement self-explanations. *Learning & Instruction*, 12, 529-556.
- Ritter, S. (1997). Communication, cooperation and competition among multiple tutor agents. In B. du Boulay & R. Mizoguchi (Eds.), *Artificial Intelligence in Education, Proceedings of AI-ED 97 World Conference* (pp. 31-38). Amsterdam: IOS Press.
- Ryan, A. M., Gheen, M. H., & Midgley, C. (1998). Why do some students avoid asking for help? An examination of the interplay among students' academic efficacy,

teachers' social-emotional role, and the classroom goal structure. *Journal of Educational Psychology*, 90(3), 528-535.

Ryan, A. M., & Pintrich, P. R. (1997). "Should I ask for help?" The role of motivation and attitudes in adolescents' help seeking in math class. *Journal of Educational Psychology*, 89(2), 329-341.

Ryan, A. M., Pintrich, P. R., & Midgley, C. (2001). Avoiding seeking help in the classroom: Who and why? *Educational Psychology Review*, 13(2), 93-114.

Schoenfeld, A. H. (1987). What's All the Fuss About Meta-cognition? In A. H. Schoenfeld (Ed.), *Cognitive Science and Mathematics Education* (pp. 189-215). Hillsdale, NJ: Erlbaum.

Schofield, J.W. (1995). *Computers and classroom culture*. Cambridge, UK: Cambridge University Press.

Slotta, J. D., & Linn, M. C. (2000). The Knowledge Integration Environment: Helping students use the Internet effectively. In M. J. Jacobson & R. B. Kozma (Eds.), *Innovations in science and mathematics education: Advanced designs for technologies of learning*. Mahwah, NJ, Erlbaum: 193-226.

Van der Meij, H. (1988). Constraints on question asking in classrooms. *Journal of Educational Psychology* 80(3), 401-405.

Vygotsky, L. S. (1978). *Mind in Society*. Cambridge: Harvard University Press.

White, B., & Frederiksen, J. (1998). Inquiry, modeling, and metacognition: Making science accessible to all students. *Cognition and Instruction*, 16(1), 3-117.

Wood, D. (2001). Scaffolding, contingent tutoring, and computer-supported learning.

International Journal of Artificial Intelligence in Education, 12.

Wood, H. A., & Wood, D. J. (1999). Help seeking, learning and contingent tutoring.

Computers and Education, 33, 153-169.