

2004

Toward Tutoring Help Seeking; Applying Cognitive Modeling to Meta-Cognitive Skills

Vincent Aleven
Carnegie Mellon University

Bruce M. McLaren
Carnegie Mellon University

Odo Roll
Carnegie Mellon University

Kenneth Koedinger
Carnegie Mellon University

Follow this and additional works at: <http://repository.cmu.edu/hcii>

This Conference Proceeding is brought to you for free and open access by the School of Computer Science at Research Showcase @ CMU. It has been accepted for inclusion in Human-Computer Interaction Institute by an authorized administrator of Research Showcase @ CMU. For more information, please contact research-showcase@andrew.cmu.edu.

Toward tutoring help seeking

Applying cognitive modeling to meta-cognitive skills

Vincent Aleven, Bruce McLaren, Ido Roll, Kenneth Koedinger

Human-Computer Interaction Institute, Carnegie Mellon University
{aleven, bmclaren}@cs.cmu.edu, {idoroll, koedinger}@cmu.edu

Abstract. The goal of our research is to investigate whether a Cognitive Tutor can be made more effective by extending it to help students acquire help-seeking skills. We present a preliminary model of help-seeking behavior that will provide the basis for a *Help-Seeking Tutor Agent*. The model, implemented by 57 production rules, captures both productive and unproductive help-seeking behavior. As a first test of the model's efficacy, we used it off-line to evaluate students' help-seeking behavior in an existing data set of student-tutor interactions. We found that 72% of all student actions represented unproductive help-seeking behavior. Consistent with some of our earlier work (Aleven & Koedinger, 2000) we found a proliferation of hint abuse (e.g., using hints to find answers rather than trying to understand). We also found that students frequently avoided using help when it was likely to be of benefit and often acted in a quick, possibly undeliberate manner. Students' help-seeking behavior accounted for as much variance in their learning gains as their performance at the cognitive level (i.e., the errors that they made with the tutor). These findings indicate that the help-seeking model needs to be adjusted, but they also underscore the importance of the educational need that the Help-Seeking Tutor Agent aims to address.

Introduction

Meta-cognition is a critical skill for students to develop and an important area of focus for learning researchers. This, in brief, was one of three broad recommendations in a recent influential volume entitled "How People Learn," in which leading researchers survey state-of-the-art research on learning and education (Bransford, Brown, & Cocking, 2000). A number of classroom studies have shown that instructional programs with a strong focus on meta-cognition can improve students' learning outcomes (Brown & Campione, 1996; Palincsar & Brown, 1984; White & Frderiksen, 1998). An important question therefore is whether instructional technology can be effective in supporting meta-cognitive skills. A small number of studies have shown that indeed it can. For example, it has been shown that self-explanation, an important meta-cognitive skill, can be supported with a positive effect on the learning of domain-specific skills and knowledge (Aleven & Koedinger, 2002; Conati & VanLehn, 2000; Renkl, 2002; Trafton & Trickett, 2001).

This paper focuses on a different meta-cognitive skill: help seeking. The ability to solicit help when needed, from a teacher, peer, textbook, manual, on-line help system, or the Internet may have a significant influence on learning outcomes. Help seeking

has been studied quite extensively in social contexts such as classrooms (Karabenick, 1998). In that context, there is evidence that better help seekers have better learning outcomes, and that those who need help the most are the least likely to ask for it (Ryan et. al, 1998).

Help seeking has been studied to a lesser degree in interactive learning environments. Given that many learning environments provide some form of on-demand help, it might seem that proficient help use would be an important factor influencing the learning results obtained with these systems. However, there is evidence that students tend not to effectively use the help facilities offered by learning environments (for an overview, see Alevan, Stahl, Schworm, Fischer & Wallace, 2003). On the other hand, there is also evidence that when used appropriately, on-demand help can have a positive impact on learning (Renkl, 2000; Schworm & Renkl, 2002; Wood, 2001; Wood & Wood, 1999) and that different types of help (Dutke & Reimer, 2000) or feedback (McKendree, 1990; Arroyo et al., 2001) affect learning differently.

Our project focuses on the question of whether instructional technology can help students become better help seekers and, if so, whether they learn better as a result. Luckin and Hammerton (2002) reported some interesting preliminary evidence with respect to “meta-cognitive scaffolding.” We are experimenting with the effects of computer-based help-seeking support in the context of Cognitive Tutors. This particular type of intelligent tutor is designed to support “learning by doing” and features a cognitive model of the targeted skills, expressed as production rules (Anderson, Corbett, Koedinger, & Pelletier, 1995). Cognitive Tutors for high-school mathematics have been highly successful in raising students’ test scores and are being used in 1700 schools nationwide (Koedinger, Anderson, Hadley, & Mark, 1997).

As a first step toward a Help-Seeking Tutor Agent, we are developing a model of the help-seeking behavior that students would ideally exhibit as they work with the tutor. The model is implemented as a set of production rules, just like the cognitive models of Cognitive Tutors. The Help-Seeking Tutor Agent will use the model, applying its model-tracing algorithm at the meta-cognitive level to provide feedback to students on the way they use the tutor’s help facilities. In this paper, we present an initial implementation of the model. We report results of an exploratory analysis, aimed primarily at empirically validating the model, in which we investigated, using an existing data set; to what extent students’ help-seeking behavior conforms to the model and whether model conformance is predictive of learning.

Initial test bed: The Geometry Cognitive Tutor

Although our help-seeking model is designed to work with any Cognitive Tutor, and possibly other intelligent tutors as well, we are initially testing it within the Geometry Cognitive Tutor, shown in Figure 1.

The Geometry Cognitive Tutor was developed in our lab as an integrated component of a full-year geometry high-school curriculum. It is currently in routine use in 350 schools around the country. The combination of tutor and curriculum has been shown to be more effective than classroom instruction (Koedinger, Corbett, Ritter, & Shapiro, 2000). Like other Cognitive Tutors, the Geometry Cognitive Tutor uses a

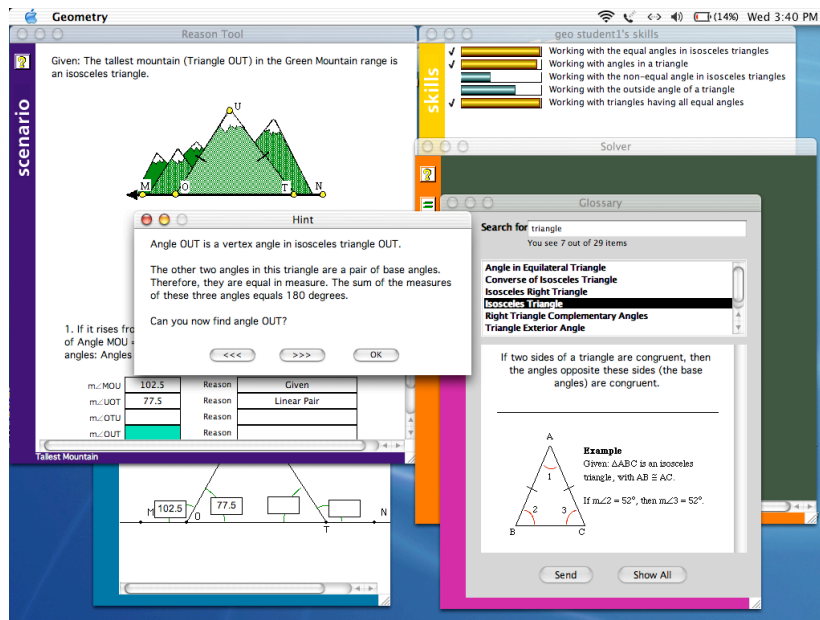


Figure 1: The Geometry Cognitive Tutor

cognitive model of the skills to be learned. It uses an algorithm called model tracing to evaluate the student's solution steps and provide feedback (Anderson et al., 1995).

The Geometry Cognitive Tutor offers two different types of help on demand. At the student's request, context-sensitive hints are provided at multiple levels of detail. This help is tailored toward the student's specific goal within the problem at hand, with each hint providing increasingly specific advice. The Geometry Cognitive Tutor also provides a less typical source of help in the form of a de-contextualized Glossary. Unlike hints, the Glossary does not tailor its help to the user's goals; rather, at the student's request, it displays information about a selected geometry rule (i.e., a theorem or definition). It is up to the student to search for potentially relevant rules in the Glossary and to evaluate which rule is applicable to the problem at hand.

Cognitive Tutors keep track of a student's knowledge growth over time by means of a Bayesian algorithm called knowledge tracing (Corbett & Anderson, 1995). At each problem-solving step, the tutor updates its estimates of the probability that the student knows the skills involved in that step, according to whether the student was able to complete the step without errors and hints. A Cognitive Tutor uses the estimates of skill mastery to select problems and make pacing decisions on an individual basis. These estimates also play a role in the model of help seeking, presented below.

A model of desired help-seeking behavior

Design

As part of our investigation into the help-seeking behavior of students, we have designed and developed a preliminary model of ideal help-seeking behavior, shown in

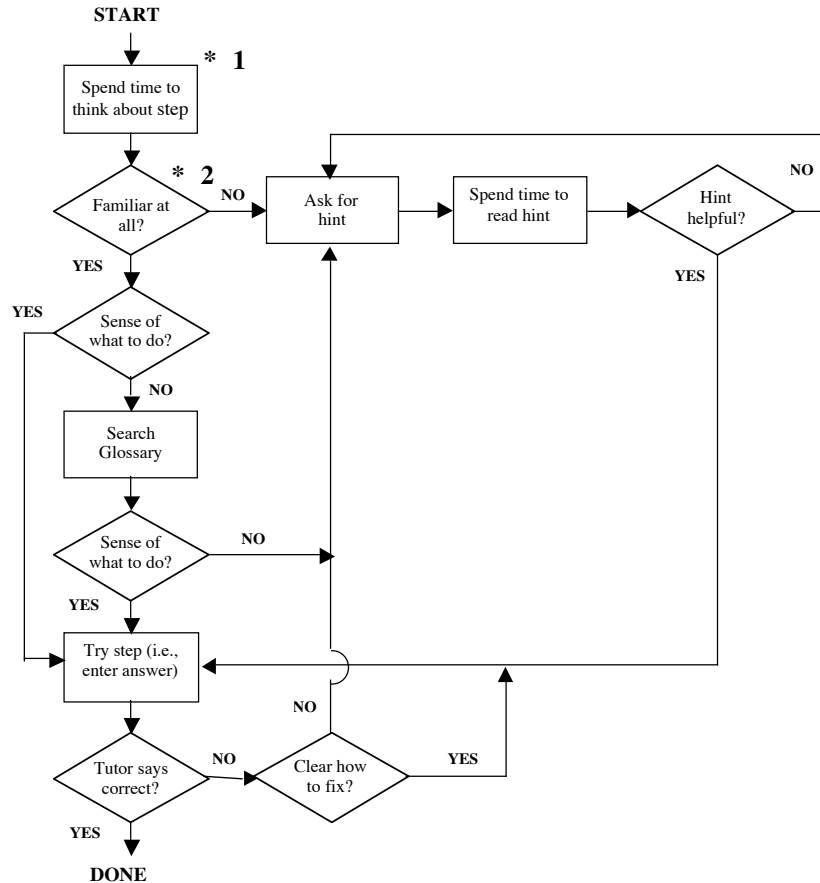


Figure 2: A model of help-seeking behavior (The asterisks “*” indicate examples of where violations of the model can occur. To be discussed later in the paper.)

Figure 2. This model shares some general traits with models of social help seeking put forward by Nelson-LeGall’s (1981) and Newman’s (1994). We believe our model is a contribution to the literature on help seeking because it is more fine-grained than existing models and will eventually clarify poorly understood relations between help seeking and learning.

According to the model, the ideal student behaves as follows: If, after spending some time thinking about a problem-solving step, a step does not look familiar, the student should ask the tutor for a hint. After reading the hint carefully, she should decide whether a more detailed hint is needed or whether it is clear how to solve the step. If the step looks familiar from the start, but the student does not have a clear sense of what to do, she should use the Glossary to find out more. If the student does have a sense of what to do, she should try to solve it. If the tutor feedback indicates that the step is incorrect, the student should ask for a hint unless it was clear how to

fix the problem. The student should think about any of her actions before deciding on her next move.

For implementation, we had to refine and make concrete some of the abstract elements of the flowchart. For example, the self-monitoring steps *Familiar at all?* and *Sense of what to do?* test how well a particular student knows a particular skill at a particular point in time. Item response theory (Hambleton & Swaminathan, 1985) is not a suitable way to address this issue, since it does not track the effect of learning over time. Instead, as a starting point to address these questions, we use the estimates of an individual student's skill mastery derived by the Cognitive Tutor's knowledge-tracing algorithm. The tests *Familiar at all?* and *Sense of what to do?* compare these estimates against pre-defined thresholds. So, for instance, if a student's current estimated level for the skill involved in the given step 0.4, our model assumes *Familiar at all?* = YES, since the threshold for this question is 0.3. For *Sense of what to do?*, the threshold is 0.6. These values are intuitively plausible but need to be validated empirically. One of the goals of our experiments with the model, described below, is to evaluate and refine the thresholds.

The tests *Clear how to fix?* and *Hint helpful?* also had to be rendered more concrete. For the *Clear how to fix?* test, the help-seeking model prescribes that a student with a higher estimated skill level (for the particular skill involved in the step, at the particular point in time that the step is tried), should re-try a step after missing it once, but that mid or lower skilled students should ask for a hint. In the future we plan to elaborate *Clear how to fix?* by using heuristics that catch some of the common types of easy-to-fix slips that students make. Our implementation of *Hint Helpful?* assumes that the amount of help a student needs on a particular step depends on their skill level for that step. Thus, a high-skill student, after requesting a first hint, is predicted to need 1/3 of the available hint levels, a mid-skill student 2/3 of the hints, and a low-skill student all of the hints. However, this is really a question of reading comprehension (or self-monitoring thereof). In the future we will use basic results from the reading comprehension literature and also explore the use of tutor data to estimate the difficulty of understanding the tutor's hints.

Implementation

We have implemented an initial version of the help-seeking model of Figure 2. The current model consists of 57 production rules. Thirty-two of the rules are "bug rules," which reflect deviations of the ideal help-seeking behavior and enable the help-seeking tutor to provide feedback to students on such deviations. The model is used to evaluate two key pieces of information each time it is invoked in the process of model-tracing at the meta-cognitive level: (1) whether the student took sufficient time to consider his or her action, (2) whether the student appropriately used, or did not use, the tutor's help facilities at the given juncture in the problem-solving process.

As an example, let us consider a student faced with an unfamiliar problem-solving step in a tutor problem. Without spending much time thinking about the step, she ventures an answer and gets it wrong. 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 "*" 1" in Figure 2) and in spite of the fact that the




Figure 3: A chain of rules in the Meta-Cognitive Model

Figure 3 shows the tree of rules explored by the model-tracing algorithm as it searched for rules matching the student’s help-seeking behavior (or in this situation, lack thereof). Various paths in the tree contain applicable rules that did not match the student’s behavior (marked with “***”), including most notably a rule that represents the “ideal” meta-cognitive behavior in the given situation (“think-about-step-deliberately”). The rule chain that matched the students’ behavior is highlighted. This chain includes an initial rule that starts the meta-cognitive cycle (“start-new-metacog-cycle”), a subsequent bug rule that identifies the student as having acted too quickly (“bug1-think-about-step-quickly”), a second bug rule that 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 that reflects the fact that the student answered incorrectly (“bug-tutor-says-step-wrong”). The feedback message in this case, compiled from the two bug rules identified in the chain, is: “Slow down, slow down. No need to rush. Perhaps you should ask for a hint, as this step might be a bit difficult for you.” The bug rules corresponding to the student acting too quickly and trying the step when they should not have are shown in Figure 4.

The fact that the student got the answer wrong is not in itself considered to be a meta-cognitive error, even though it is captured in the model by a bug rule (“bug-tutor-says-step-wrong”). This bug rule merely serves to confirm the presence of bugs captured by other bug rules, when the student’s answer (at the cognitive level) is wrong. Further, when the student answer is correct, (at the cognitive level) no feedback is given at the meta-cognitive level, even if the student’s behavior was not ideal from the point of view of the help-seeking model.

The help-seeking model uses information passed from the cognitive model to perform its reasoning. For instance, the skill involved in a particular 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 enters a value to the tutor, that value is first evaluated at the cognitive level before it is evaluated at the meta-cognitive level. An important consideration in the development of the Help-Seeking Tutor was to make it modular and useable in conjunction with a variety of Cognitive Tutors. Basically, the Help-Seeking Tutor Agent will be a plug-in agent applicable to a range of Cognitive Tutors with limited customization. We have attempted to create rules that are applicable to *any* Cognitive Tutor, not to a specific tutor. Certainly,

step is not familiar to her, she does not ask for a hint (marked as “* 2”). The students’ errors will match bug rules that capture unproductive help-seeking behavior, allowing the tutor to provide feedback.

<p>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 <i>min-thinking-time</i> to think about the step Then Remove the subgoal (next subgoal is to decide what action to take)</p>	<p>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 <i>min-familiarity-level</i> And the student has not seen all the hints yet for the current step Then Try step Set a subgoal to evaluate the result</p>
<p>Bug message: "Slow down, slow down. No need to rush."</p>	<p>Bug message: "Perhaps you should ask for a hint, as this step might be a bit difficult for you."</p>

Figure 4: Example bug rules matching unproductive help-seeking behavior.

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.

A taxonomy of help-seeking bugs

In order to compare students' help-seeking behavior against the model, we have created a taxonomy of errors (or *bugs*) in students' help-seeking behavior, shown in Figure 5. The taxonomy includes four main categories. First, the "Help Abuse" category covers situations in which the student misuses the help facilities provided by the Cognitive Tutor. This occurs when a student spends too little time with a hint ("Clicking Through Hints"), when a student requests hints (after some deliberation) when they are knowledgeable enough to either try the step ("Ask Hint when Skilled Enough to Try Step") or use the Glossary ("Ask Hint when Skilled Enough to Use Glossary"), or when a student overuses the Glossary ("Glossary Abuse"). Recall from the flow chart in Figure 2 that a student with high mastery for the skill in question should first try the step, a student with medium mastery should use the Glossary, and a student with low mastery should ask for a hint.

Second, the category "Try-Step Abuse" represents situations in which the student attempts to hastily solve a step and gets it wrong, either when sufficiently skilled to try the step ("Try Step Too Fast") or when less skilled ("Guess Quickly when Help Use was Appropriate").

Third, situations in which the student could benefit from asking for a hint or inspecting the Glossary, but chose to try the step instead, are categorized as "Help Avoidance". There are two bugs of this type – "Try Unfamiliar Step Without Hint Use" and "Try Vaguely Familiar Step Without Glossary Use."

Finally, the category of "Miscellaneous Bugs" covers situations not represented in the other high-level categories. The "Read Problem Too Fast" error describes hasty reading of the question, when first encountered followed by a rapid help request. "Ask

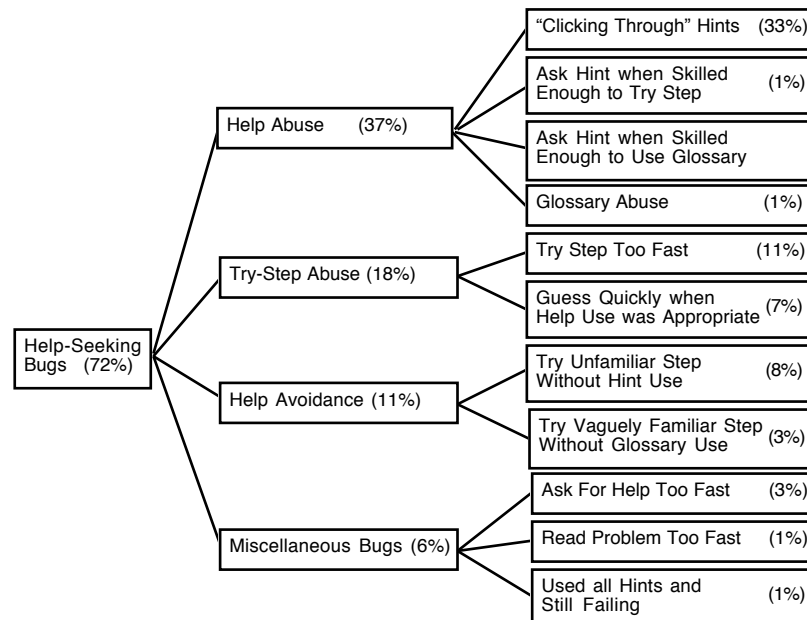


Figure 5: A taxonomy of help-seeking bugs. The percentages indicate how often each bug occurred in our experiment.

for Help Too Fast” describes a similar situation in which the student asks for help too quickly after making an error. The “Used All Hints and Still Failing” bug represents situations in which the student has seen all of the hints, yet cannot solve the step (i.e., the student has failed more than a threshold number of times). In our implemented model, the student is advised to talk to the teacher in this situation.

In general, if the student gets the step right at the cognitive level, we do not consider a meta-cognitive bug to have occurred, regardless of whether the step was hasty or the student’s skill level was inappropriate.

Comparing the model to students’ actual meta-cognitive behavior

We conducted an empirical analysis to get a sense of how close the model is to being usable in a tutoring context and also to get a sense of students’ help-seeking behavior. We replayed a set of logs of student-tutor interactions, comparing what actually happened in a given tutor unit (viz., the Angles unit of the Geometry Cognitive Tutor), without any tutoring on help seeking, with the predictions made by the help-seeking model. This methodology might be called “model tracing after the fact” – it is not the same as actual model tracing, since one does not see how the student might have changed their behavior in response to feedback on their help-seeking behavior. We determined the extent to which students’ help-seeking behavior conforms to the model. We also determined the frequency of the various categories of meta-cognitive bugs described above. Finally, we determined whether students’ help-seeking behavior (that is, the degree to which they follow the model) is predictive of their learning results.

Table 1: Correlation between the bug categories and Learning Gain

	Help Abuse	Try-Step Abuse	Help Avoidance	Misc. Bugs	Total Meta-cognitive Bugs
Learning Gain	-0.46	0.02	-0.10	-0.47	-0.61

The data used in the analysis were collected during an earlier study in which we compared the learning results of students using two tutor versions, one in which they explained their problem-solving steps by selecting the name of the theorem that justifies it and one in which the students solved problems without explaining (Alevan & Koedinger, 2002). For purposes of the current analysis, we group the data from both conditions together. Students spent approximately 7 hours working on this unit of the tutor. The protocols from interaction with the tutor include data from 49 students, 40 of whom completed both the Pre- and Post-Tests. These students performed a total of approximately 47,500 actions related to skills tracked by the tutor.

The logs of the student-tutor interactions were replayed with each student action (either an attempt at answering, a request for a hint, or the inspection of a Glossary item) checked against the predictions of the help-seeking model. Actions that matched the model's predictions were recorded as "correct" help-seeking behavior, actions that did not match the model's predictions as "buggy" help-seeking behavior. The latter actions were classified automatically with respect to the bug taxonomy of Figure 5, based on the bug rules that were matched. We computed the frequency of each bug category (shown in Figure 5) and each category's correlation with learning gains. The learning gains (LG) were computed from the pre- and post-test scores according to the formula $LG = (Post - Pre) / (1 - Pre)$, mean 0.41; standard deviation 0.28).

The overall ratio of help-seeking errors to all actions was 72%; that is, 72% of the students' actions did not conform to the help-seeking model. The most frequent errors at the meta-cognitive level were Help Abuse (37%), with the majority of these being "Clicking Through" hints (33%). The next most frequent category was Try Step Abuse (18%), which represents quick attempts at answering steps. Help Avoidance – not using help at moments when it was likely to be beneficial – was also quite frequent (11%), especially if "Guess quickly when help was needed" (7%), arguably a form of Help Avoidance as well as Try-Step Abuse, is included in both categories.

The frequency of help-seeking bugs was correlated strongly with the students' overall learning ($r = -0.61$ with $p < .0001$), as shown in Table 1. The model therefore is a good predictor of learning gains – the more help-seeking bugs students make, the less likely they are to learn. The correlation between students' frequency of success at the cognitive level (computed as the percentage of problem steps that the student completed without errors or hints from the tutor) and learning gain is about the same ($r = .58$, $p = .0001$) as the correlation between help-seeking bugs and learning. Success in help seeking and success at the cognitive level were highly correlated ($r = .78$, $p < .0001$). In a multiple regression, the combination of help-seeking errors and errors at the cognitive level accounted only for marginally more variance than either one alone.

We also looked at how the bug categories correlated with learning (also shown in Table 1). Both Help Abuse and Miscellaneous Bugs were negatively correlated with

learning with $p < 0.01$. These bug categories have in common that the students avoid trying to solve the step. On the other hand, Try Step Abuse and Help Avoidance were not correlated with learning.

Discussion

Our analysis sheds light on the validity of the help-seeking model and the adjustments we must make before we use it for “live” tutoring. The fact that some of the bug categories of the model correlate negatively with learning provides some measure of confidence that the model is on the right track. The correlation between Hint Abuse and Miscellaneous Bugs and students’ learning gain supports our assumption that the help-seeking model is valid in identifying these phenomena. On the other hand, the model must be more lenient with respect to help-seeking errors. The current rate of 72% implies that the Help-Seeking Tutor Agent would intervene (i.e., present a bug message) in 3 out of every 4 actions taken by a student. In practical use, this is likely to be quite annoying and distracting to the student. Another finding that may lead to a change in the model is the fact that Try-Step Abuse did not correlate with learning. Intuitively, it seems plausible that a high frequency of incorrect guesses would be negatively correlated with learning. Perhaps the threshold we used for “thinking time” is too high; perhaps it should depend on the student’s skill level. This will require further investigation. Given that the model is still preliminary and under development, the findings on students’ help seeking should also be regarded as subject to further investigation.

The finding that students often abuse hints confirms earlier work (Aleven & Koedinger, 2000; Aleven, McLaren, & Koedinger, to appear; Baker, Corbett, & Koedinger, in press). The current analysis extends that finding by showing that help abuse is frequent relative to other kinds of help-seeking bugs and that it correlates negatively with learning. However, the particular rate that was observed (37%) may be inflated somewhat because of the high frequency of “Clicking Through Hints” (33%). Since typically 6 to 8 hint levels were available, a single “clicking-through” episode – selecting hints until the “bottom out” or answer hint is seen – yields multiple actions in the data. One would expect to see a different picture if the clicking episodes were clustered into a single action.

Several new findings emerged from our empirical study. As mentioned, a high help-seeking error rate was identified (72%). To the extent that the model is correct, this suggests that students generally do not have good help-seeking skills. We also found a relatively high Help Avoidance rate, especially if we categorize “Guess Quickly when Help Use was Appropriate” as a form of Help Avoidance (18% combined). In addition, since the frequency of the Help Abuse category appears to be inflated by the high prevalence of Clicking Through Hints, categories such as Help Avoidance are correspondingly deflated. The significance of this finding is not yet clear, since Help Avoidance did not correlate with learning. It may well be that the model does not yet successfully identify instances in which the students should have asked for help but did not. On the other hand, the gross abuse of help in the given data set is likely to have lessened the impact of Help Avoidance. In other words, given that

the Help Avoidance in this data set was really Help Abuse avoidance, the lack of correlation with learning is not surprising and should not be interpreted as meaning that help avoidance is not a problem or has no impact on learning. Future experiments with the Help-Seeking Tutor Agent may cast some light on the importance of help avoidance, in particular if the tutor turns out to reduce the Help Avoidance rate.

It must be said that we are just beginning to analyze and interpret the data. For instance, we are interested in obtaining a more detailed insight into and understanding of Help Avoidance. Under what specific circumstances does this occur? We also intend to investigate in greater detail how students so often get a step right even when they answer too quickly, according to the model. Finally, how different would the results look if clicking through hints is considered a single mental action?

Conclusion

We have presented a preliminary model of help seeking which will form the basis of a Help-Seeking Tutor Agent, designed to be seamlessly added to existing Cognitive Tutors. To validate the model, we have run it against pre-existing tutor data. This analysis suggests that the model is on the right track, but is not quite ready for “live” tutoring, in particular because it would lead to feedback on as much as three-fourths of the students’ actions, which is not likely to be productive. Although the model is still preliminary, the analysis also sheds some light on students’ help-seeking behavior. It confirms earlier findings that students’ help-seeking behavior is far from ideal and that help-seeking errors correlate negatively with learning, underscoring the importance of addressing help-seeking behavior by means of instruction.

The next step in our research will be to continue to refine the model, testing it against the current and other data sets, and modifying it so that it will be more selective in presenting feedback to students. In the process, we hope to gain a better understanding, for example, of the circumstances under which quick answers are fine or under which help avoidance is most likely to be harmful. Once the model gives satisfactory results when run against existing data sets, we will use it for live tutoring, integrating the Help-Seeking Tutor Agent with an existing Cognitive Tutor. We will evaluate whether students’ help-seeking skill improves when they receive feedback from the Help-Seeking Tutor Agent and whether they obtain better learning outcomes. We will also evaluate whether better help-seeking behavior persists beyond the tutor units in which the students are exposed to the Help-Seeking Tutor Agent and whether students learn better in those units as a result. A key hypothesis is that the Help-Seeking Tutor Agent will help students to become better learners.

Acknowledgments

The research reported in this paper is supported by NSF Award No. IIS-0308200.

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 Student 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. M., & Koedinger, K. R. (to appear). Towards Computer-Based Tutoring of Help-Seeking Skills. In S. Karabenick & R. Newman (Eds.), *Help Seeking in Academic Settings: Goals, Groups, and Contexts*. Mahwah, NJ: Erlbaum.
- Aleven, V., Stahl, E., Schworm, S., Fischer, F., & Wallace, R.M. (2003). Help Seeking in Interactive Learning Environments. *Review of Educational Research*, 73(2), 277-320.
- 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.
- 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).
- Baker, R. S., Corbett, A. T., & Koedinger, K. R. (in press). Detecting Student Misuse of Intelligent Tutoring Systems. In *Proceedings of the 7th International Conference on Intelligent Tutoring Systems. ITS 2004*.
- Bransford, J. D., Brown, A. L., & Cocking, R. R. (Eds.) (2000). *How People Learn: Brain, Mind, Experience, and School*. Washington, CD: National Academy Press.
- Brown, A. L., & Campione, J. C. (1996). Guided Discovery in a Community of Learners. In K. McGilly (Ed.), *Classroom Lessons: Integrating Cognitive Theory and Classroom Practice* (pp. 229-270). Cambridge, MA: The MIT Press.
- Conati C. & VanLehn K. (2000). Toward computer-based support of meta-cognitive skills: A computational framework to coach self-explanation. *International Journal of Artificial Intelligence in Education*, 11, 398-415.
- Corbett, A. T. & Anderson, J. R. (1995). Knowledge tracing: Modeling the acquisition of procedural knowledge. *User Modeling and User-Adapted Interaction*, 4, 253-278.
- 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.
- Hambleton, R. K. & Swaminathan, H. (1985). *Item Response Theory: Principles and Applications*. Boston: Kluwer.
- Karabenick, S. A. (Ed.) (1998). *Strategic help seeking. Implications for learning and teaching*. Mahwah: Erlbaum.
- 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., 1200 Penn Avenue, Suite 150, Pittsburgh, PA 15222, E-mail: info@carnegielearning.com, Web: <http://www.carnegielearning.com>
- Luckin, R., & Hammerton, L. (2002). Getting to know me: Helping learners understand their own learning needs through meta-cognitive 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.
- McKendree, J. (1990). Effective feedback content for tutoring complex skills. *Human Computer Interaction*, 5, 381-413.
- Nelson-LeGall, S. (1981). Help-seeking: An understudied problem-solving skill in children. *Developmental Review*, 1, 224-246.
- 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, NJ: Erlbaum.
- Palincsar, A. S., & Brown, A. L. (1984). Reciprocal teaching of comprehension-fostering and comprehension monitoring activities. *Cognition and Instruction*, 1, 117-175.
- Renkl, A. (2002). Learning from worked-out examples: Instructional explanations supplement self-explanations. *Learning and Instruction*, 12, 529-556.
- 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.
- Schworm, S. & Renkl, A. (2002). Learning by solved example problems: Instructional explanations reduce self-explanation activity. In W. D. Gray & C. D. Schunn (Eds.), *Proceeding of the 24th Annual Conference of the Cognitive Science Society* (pp.816-821). Mahwah, NJ: Erlbaum.
- Trafton, J.G., & Trickett, S.B. (2001). Note-taking for self-explanation and problem solving. *Human-Computer Interaction*, 16, 1-38.
- 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., & Wood, D. (1999). Help seeking, learning and contingent tutoring. *Computers and Education*, 33, 153-169.