Modeling Helping Behavior in an Intelligent Tutor for Peer Tutoring

Erin WALKER[a], Nikol RUMMEL[b], and Kenneth R. KOEDINGER[a]
[a] *Human-Computer Interaction Institute, Carnegie Mellon University, USA*
[b] *Insitute of Psychology, University of Freiburg, Germany*

Abstract. Giving effective help is an important collaborative skill that leads to improved learning for both the help-giver and help-receiver. Adding intelligent tutoring to student interaction may be one effective way of assisting students in giving and receiving better help. However, such systems have proven difficult to implement, in part due to the challenges of modeling productive dialogue in a collaborative activity. We present a theoretical model of good helping behavior in a peer tutoring context, and validate the model using student tutoring data, linking optimal and buggy behaviors to learning outcomes. We discuss the implications of the model with respect to providing intelligent tutoring for peer tutoring.

Keywords. Cognitive tutors, peer tutoring, adaptive collaborative learning support

Introduction

Giving effective help is a key component of the "promotive" interactions described by Johnson and Johnson [1] that lead students to benefit from collaboration. An abundance of evidence suggests that the act of giving help improves learning (see [2]). In explaining to others, students have the opportunity to engage in knowledge-building, where they reflect on their knowledge, identify gaps, and move to repair them [3]. For the receiver to benefit, additional conditions must be met: The help must be needed, elaborated, target the receiver's misconception, and used constructively by the receiver [4]. Most students do not exhibit these positive helping behaviors spontaneously and it appears that they must be assisted in order to give and receive beneficial help [5].

One potential way of supporting student interaction is through intelligent tutoring. Early results have shown adaptive support of student collaboration to be better than nonadaptive support of student collaboration and individual learning [6]. Despite this promise, few adaptive collaborative learning support (ACLS) systems have been implemented, and those that have were generally not evaluated for their impact on student interaction and learning [7]. One of the obstacles to constructing ACLS systems is the difficulty in developing a model of effective dialog that can serve as a basis for providing feedback. These systems generally use collaborative dialogue only to assess simple metrics, such as equal participation in the learning task [8]. However, some model sequences of collaborative interactions in order to detect particular behaviors and provide relevant feedback [6]. We apply this approach to student helping behavior.

Our overall goal is to develop an intelligent tutoring system that supports students in peer tutoring, an activity which provides students with many opportunities to give and receive help. By integrating different theories of good helping behavior, we developed a model of good peer tutoring that can serve as a basis for feedback. The model is relevant for classroom use in that it attempts to facilitate learning for students in both roles. Then, using peer tutoring data that we collected, we examined the validity of the positive and "buggy" behaviors present in the model, demonstrating that the model can indeed be used as part of an intelligent tutor for good helping behaviors.

1. Theoretical Model of Good Peer Tutoring

We modeled good helping skills within the context of a peer tutoring activity, because these helping behaviors are integral to learning from the activity. Peer tutoring has been shown to lead to deep learning in classroom environments for both students involved, as long as students engage in positive interactions [9]. In particular, peer tutors learn by preparing to tutor, reflecting on tutee errors, and providing elaborated help (see [3] for a review), while tutees learn from trying to understand their tutor's explanations and overcoming problem-solving impasses [10]. We have constructed a model of good peer tutoring which focuses on help-related behaviors that should lead to learning for both the help-giver and help-receiver. The model addresses when and how to give help.

The model, depicted in Figure 1, begins when the tutee starts a new step in a given problem. For the tutee behavior component of the model (the dark-shaded area of the diagram), we have adapted a model for good help-seeking developed by Aleven and colleagues [11]. The model encourages tutees to solve problems on their own, but ensures that tutors provide scaffolding when appropriate. In our adaptation of the Aleven model, tutees can perform two behaviors: trying a step or asking for a hint.

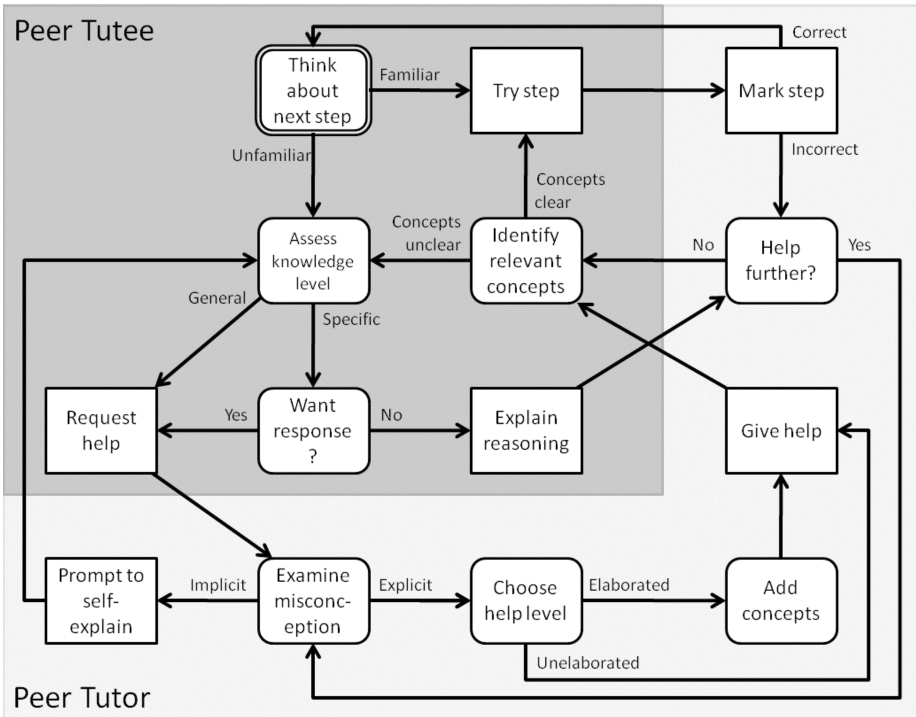

Figure 1. Model of tutor and tutee helping behavior, designed to contribute to the learning of both parties.

Tutees should ask for a hint when they begin an unfamiliar step, after they make an error they do not know how to fix, or after they have received a hint they do not know how to use. They should try a step if it is familiar, if they understand the help given to them, or if they understand the error they just made on the step. We added two further elements to the model. First, tutees can choose to self-explain instead of requesting help. Self-explanations have been shown to be very beneficial for student learning [12], and may also allow tutors to reflect on their content and target explanations toward tutee misconceptions. Further, if students choose to request help instead of self-explain, requests that include specific references to the problem have been shown to be more useful than general requests [4]. Therefore, if tutees have specific knowledge about the appropriate next step, our model suggests that they add specific content to their request.

The peer tutor side of the model (the light-shaded area) is based on a combination of findings [see 4, 10]. Here, the peer tutor's behaviors include giving yes-no feedback, giving help, and prompting the tutee to self-explain. Yes-no feedback is beneficial for tutee learning in that it provides them with feedback on their problem-solving, and potentially beneficial for tutor learning in that tutors reflect on the nature of correct and incorrect problem-solving steps. Peer tutors can then deliver help after an incorrect step, after a help request, or after a self-explanation, and take several cognitive steps in constructing the help. VanLehn and colleagues [10] show that help tailored toward a tutee misconception is beneficial for the tutee. Therefore, when tutees ask for help, if they have recently committed an error, peer tutors should identify the tutee misconception. If they cannot, they should prompt the tutee to self-explain until the misconception becomes clear. The self-explanation benefits the tutee as well [12]. After peer tutors have identified the misconception, they can begin constructing the help, deciding whether the help should be elaborated or unelaborated. Elaborated help, where the tutee elaborates on the content of the help, has been shown to be beneficial for both tutor and tutee learning, as the process of constructing the help leads tutors to reflect on their own knowledge and move to repair gaps [4]. Tutees can then use the elaborated help to build on their knowledge. Augmenting elaborated help with conceptual content further facilitates these processes [13]. However, there may be some cases where unelaborated help is the most appropriate kind of help. If the tutee lacks the relevant knowledge to continue with the problem it may be better for the tutor to give the answer, treating the problem as a worked example.

2. Data Collection

We hypothesized that students do indeed exhibit the good peer tutoring behaviors present in the model, and these behaviors are related to learning. Further, we wished to explore whether certain behaviors not represented in the model might be buggy behaviors that are negatively related to learning. In order to evaluate this hypothesis, we used peer tutoring data that we had collected with an extension to the Cognitive Tutor Algebra (CTA). The data used are drawn from a study that compared adaptive support of peer tutors (cognitive hints and feedback) to fixed support of peer tutors (access to problem solutions) [14]. The peer tutoring activity involved a preparation phase and a collaboration phase. Students were put in same-ability pairs so that they could participate equally in the interaction. In the preparation phase, each member of the pair solved a separate set of problems. In the collaboration phase, students took turns tutoring each other on the problems they had solved during the preparation phase.

As tutees solved the problems, peer tutors watched remotely, and marked steps right or wrong. Tutees and peer tutors could interact with each other in a chat window (see Figure 2). To assist peer tutors in helping tutees, peer tutors could consult a worked-out problem solution. In the adaptive condition, peer tutors received feedback if they made errors in marking tutee steps, and were able to ask for a hint from the cognitive tutor.

In the study, students took a pretest, spent two class periods in the intervention, and took a delayed posttest two weeks later. Participants were randomly assigned to condition. There were 31 students included in the process analysis for the two collaborative conditions (14 fixed, 17 adaptive). Out of these students, 10 did not take the delayed posttest, leaving 21 participants (10 fixed, 11 adaptive). Although all students showed learning gains, there were no significant differences in the delayed gain for the fixed (M = .30, SD = .50) and adaptive conditions (M = .29, SD = .19). We use the 717 lines of chat produced by the students in the study to explore the model.

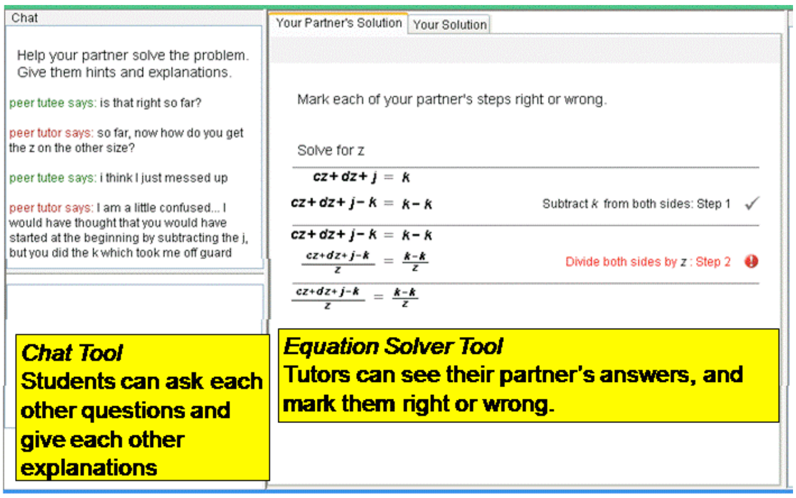

Figure 2. Peer tutor's interface. Tutees solve the problem and can ask tutors questions. Tutors mark steps right or wrong and provide explanations to the tutee.

3. Evaluating the Model Using Dialog Protocols

3.1. Defining the model behaviors

We coded tutee utterances for help requests and self-explanations, and coded tutor utterances for elaborated help, unelaborated help, prompts and feedback. We also coded student dialog for general behaviors not found in the model. Table 1 describes the full coding of tutee and tutor behaviors. Based on the model, we performed two additional classifications. We coded requests for whether they were specific or general, in order to be able to evaluate the outcome of the "assess knowledge level" node. We also coded elaborated help for whether it referred to one of seven specific problem concepts in accordance with the "add concepts" node (e.g., "you have to get both t on the same side" was coded for the "combine variable terms" concept). Two raters independently coded 85% of the dialogs on each dimension (kappa = 0.74 for tutee codes, kappa = .86 for tutor codes), and resolved disagreements through discussion.

Table 1. Coding scheme for tutor and tutee dialogue.

Role	Category	Model nodes	Description	Examples
Tutee	Request (specific or general)	Request Help	Statement relating to the problem that requires a response from the tutor.	"how do I get b by itself", " help"
Tutee	Self-Explanation	Explain Reasoning	Tutee statement s containing specific problem-related content	"so I get w on one side"
Tutee	Problem-related statement	None	Tutee statements containing general problem-related content	"I'm lost"
Tutor	Elaborated help	Choose Help Type, Give Help	Explanation of a step, hint on how to complete a step, describing an error	"now get m by itself"
Tutor	Unelaborated help	Choose Help Type, Give Help	Direct instruction on how to complete all or part of the next step,	"factor out t", "then divide"
Tutor	Feedback	Mark Step	Indication of whether a step was right or wrong	"good", "no"
Tutor	Prompt	Prompt to Self-Explain	Ask about knowledge of next step or actions on previous step	"why did you take that step"
Both	Activity-related statement	None	Coordination and activity-related statements	"what are you doin?"
Both	Off-topic	None	Statements not related to the problem or activity	"Hes datin Nichol"

3.2. Model behaviors and learning

Next, we examined how much students exhibited the behaviors present in the model and how those behaviors related to learning. In this section, we will discuss trends as well as significant results, as the relatively small amount of chat may indicate that students may not have engaged in sufficient dialog over the course of the hour and a half intervention to have a highly visible impact on their learning. First, we looked at tutee help-seeking behaviors (see Table 2). Tutees made a large number of requests (33%), particularly compared to their number of self-explanations (8%). The number of tutee requests for help were indeed marginally correlated with peer tutor learning gains, $r(19) = .4.18$, $p = .07$, suggesting that the requests triggered reflective processes. While not significant, self-explanations trended toward being positively correlated with *tutee* learning, $r(20) = .327$, $p = .14$. Apparently, tutees who learned more better explained their reasoning.

Table 2. Frequencies of help-seeking behaviors per tutee.

Code	Specific requests	General requests	Self-explanation	Problem-related	Activity-related	Off-topic
Mean	1.61	3.16	1.06	2.16	5.97	2.03
Standard deviation	1.56	2.35	1.06	2.05	4.59	3.40
Percent	11%	22%	8%	13%	37%	10%

We then turned to tutor help-giving behaviors (see Table 3). Here, the majority of student talk was unelaborated help (36%). Students did use elaborated help and feedback, but prompts, which were an integral part of our model, were rare (3%). The total amount of feedback and elaborated help given was correlated with tutor learning, $r(20) = 0.436$, $p < .05$, suggesting that tutors who engaged in reflection and knowledge construction learned the most. Interestingly, *percent* unelaborated help was marginally negatively correlated with tutor learning gains, $r(20) = -0.411$, $p = .06$, indicating that unelaborated help was problematic when it crowded out other forms of talk. Students introduced concepts into their elaborated help 35% of the time (SD = 40%). This percentage was marginally correlated with tutee learning, $r(19) = .426$, $p = .06$, supporting the model hypothesis that conceptual elaborated help was best.

Table 3. Frequencies of help-giving behaviors per peer tutor.

Code	Elaborated help	Unelaborated help	Feedback	Prompts	Activity-related	Off-topic
Mean	2.19	8.06	1.97	.68	6.45	2.03
Standard deviation	2.20	8.04	1.91	.98	4.85	3.40
Percent	12%	36%	10%	3%	31%	8%

Next, we examined the relationships between tutor and tutee model behaviors, such as how the need for tutee help interacts with help given. According to the model, tutors should give help after a direct help request, have the option of giving help after an incorrect step or self-explanation, and should not give tutees help after a correct step. Table 4 displays the number of times students gave help in each situation, and the help they gave. Giving help when needed (measured by percent requests answered) was positively correlated with tutee learning (M=61%, SD=33%, $r(19) = 0.481$, $p < 0.05$), while giving help when optional (measured by percent self-explanations and incorrect attempts responded to) was marginally correlated with *tutor* learning (M = 61%, SD = 53%, $r(19) = 0.383$, $p < 0.10$). While the percent help given when not needed only trended to being related to tutee learning, $r(19) = -0.317$, $p = 0.17$, percent unelaborated help when not needed (the worst kind of help in the table), was significantly negatively correlated with tutee learning, $r(19) = -0.553$, $p = .01$.

Table 4. Frequencies of types of help given at particular times.

	Elaborated	Unelaborated	Feedback	None
Requests	M = .65 (SD = 0.80)	M = 2.26 (SD = 2.58)	M = .90 (SD = 1.27)	M=1.90 (SD=2.10)
Incorrect Attempts	M = 0.87 (SD = 1.45)	M = 1.97 (SD = 1.96)	M = 0.32 (SD = 0.94)	M = 7.76 (SD = 5.77)
Self-explanation	M = 0.16 (SD = 0.52)	M = 0.42 (SD = 0.72)	M = 0.10 (SD = 0.30)	M = .35 (SD = 0.55)
Correct Attempts	M = 0.26 (SD = 0.58)	M = 2.26 (SD = 4.02)	M = 0.81 (SD = 1.11)	M = 23.76 (SD = 8.20)

4. Discussion

We have described a model of peer tutoring and validated it with student data. The data complements the model in several ways. First, we operationalized the concepts in the model, distinguishing between the different ways students request help and the different ways they give help. Second, we found several links between model behaviors and tutor learning that corresponded to previous literature on learning from peer tutoring, suggesting our model is indeed a valid representation of good peer tutoring. We further used the student data to identify three categories of suboptimal behaviors: departing from the model (e.g., giving help after a correct step), not engaging in theoretically positive behaviors (e.g., prompts), and over-engaging in certain model-related behaviors (e.g., unelaborated help). It should be noted that, in the model, the type and timing of student help are the main focus, and the model does not explicitly contain cognitive elaboration or reflection processes. As these processes are only visible through student behaviors, we focus on indirectly supporting them rather than explicitly. In general, because our model describes behaviors and their conditions, it can serve as a basis for cognitive tutoring.

However, using intelligent tutoring in a traditional way to limit student behavior to the paths represented in the model may overstructure the student interaction, with negative consequences [15]. Student collaboration is more open-ended than traditional intelligent tutoring domains, and thus a given buggy behavior is not a clear error, but a suboptimal path. Giving unelaborated help when it is not needed becomes a problem when it becomes a pattern, but a single instance of this behavior is not likely to have a negative effect on learning. For this reason, our model is constructed in a way that can facilitate flexibility on the part of the intelligent tutor. The conditions for many of the peer tutor and tutee actions are judgments about their own and their partner's knowledge. As students make help-related decisions based on their assessment of the situation, they can choose to take different model paths (i.e., they can choose to help after an incorrect step or wait until a request). Building this freedom into an intelligent tutoring system also implies that if a peer tutor makes a decision inconsistent with the intelligent tutor estimate of the situation, it may not be appropriate for the intelligent tutor to intervene, as it is possible that students have a better understanding of the context. However, if deviations from the model accumulate, the intelligent tutoring system can pinpoint the student error and act.

We have proposed a theoretical model of peer tutoring that can be used as a basis for feedback, and then validated the model using data that separates positive tutoring behaviors from negative tutoring behaviors. We are currently building an intelligent tutor based on the model that will analyze student interactions as they occur and give relevant feedback. We believe that the model and resulting intelligent tutor will generalize to other collaborative activities involving help exchanges and ultimately increase our understanding of how to support learning from collaboration.

Acknowledgments

This research is supported by the Pittsburgh Science of Learning Center, NSF Grant #0354420. Thanks to Carolyn Rose, Amy Ogan, Ido Roll, Ruth Wylie, & Sean Walker.

References

[1] Johnson, D. W. and Johnson, R. T. (1990). Cooperative learning and achievement. In S. Sharan (Ed.), *Cooperative learning: Theory and research* (pp. 23-37). New York: Praeger.

[2] Ploetzner, R., Dillenbourg, P., Preier, M.,&Traum, D. (1999). Learning by explaining to oneself and to others. In P. Dillenbourg (Ed.), Collaborative learning: Cognitive and computational *approaches* (pp. 103–121). Oxford, England: Pergamon.

[3] Roscoe, R. D. & Chi, M. (2007). Understanding tutor learning: Knowledge-building and knowledge-telling in peer tutors' explanations and questions, *Review of Educational Research*. 77(4), 534-574.

[4] Webb, N. M., & Mastergeorge, A. (2003). Promoting effective helping behavior in peer-directed groups. *International Journal of Educational Research*, 39, 73-97.

[5] Lou, Y., Abrami, P. C., d'Apollonia, S. (2001). Small group and individual learning with technology: A meta-analysis. *Review of Educational Research, 71* (3), 449-521.

[6] Kumar, R., Rosé, C.P., Wang, Y.C., Joshi, M., Robinson, A.: Tutorial dialogue as adaptive collaborative learning support. In: Proceedings of the 13th International Conference on Artificial Intelligence in Education (AIED 2007). IOSPress, Amsterdam (2007).

[7] Walker, E., Rummel, N., & Koedinger, K. R. CTRL: A Research Architecture for Providing Adaptive Collaborative Learning Support. Submitted to *User Modeling and User-Adapted Interaction.*

[8] Baghaei, N., Mitrovic, T., and Irwin, W. (2007). Supporting Collaborative Learning and Problem Solving in a Constraint-based CSCL Environment for UML Class Diagrams. *International Journal of Computer-Supported Collaborative Learning, 2* (2-3), 159-190.

[9] Fantuzzo, J. W., King, J., & Heller, L. (1992). Effects of reciprocal peer tutoring on mathematics and school adjustment: A componential analysis. *Journal of Educational Psychology*, 84, 331-339.

[10] VanLehn, K., Siler, S., Murray, C., Yamauchi, T., & Baggett, W. (2003). Why do only some events cause learning during human tutoring? *Cognition and Instruction*, 21(3), 209-249.

[11] 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 *7th International Conference on Intelligent Tutoring Systems* (pp. 227-239).

[12] Chi, M. T. H., DeLeeuw, N., Chiu, M.-H., & LaVancher, C. (1994). Eliciting self-explanations improves understanding. *Cognitive Science*, 18, 439-477.

[13] Fuchs, L., Fuchs, D., Hamlett, C., Phillips, N., Karns, K., & Dutka, S. (1997). Enhancing students' helping behavior during peer-mediated instruction with conceptual mathematical explanations. *The Elementary School Journal, 97*(3), 223-249.

[14] Walker, E., Rummel, N., & Koedinger, K. (2008). To tutor the tutor: Adaptive domain support for peer tutoring. In B. Woolf, E. Aimeur, R. Nkambou, S. Lajoie (Eds), Proceedings of the *9th International Conference on Intelligent Tutoring Systems.* (pp. 626-635).

[15] Dillenbourg, P. (2002). Over-scripting CSCL: The risk of blending collaborative learning with instructional design. In Kirschner, P. A. (Ed.), Three worlds of CSCL: Can we support CSCL? (61-91). Heerlen: Open Universiteit Nederland.