2006

The Help Tutor: Does Metacognitive Feedback Improve Students' Help-Seeking Actions, Skills, and Learning?

Ido Roll
Carnegie Mellon University

Vincent Aleven
Carnegie Mellon University

Bruce M. McLaren
Carnegie Mellon University

Eunjeong Ryu
Carnegie Mellon University

Ryan S.J.d Baker
Carnegie Mellon University

See next page for additional authors

Authors
Ido Roll, Vincent Aleven, Bruce M. McLaren, Eunjeong Ryu, Ryan S.J.d Baker, and Kenneth R. Koedinger

In the Proceedings of the 8th International Conference on Intelligent Tutoring Systems
Jhongli, Taiwan, June 26-30, 2006

The Help Tutor: Does Metacognitive Feedback Improve Students' Help-Seeking Actions, Skills and Learning?

Ido Roll, Vincent Aleven, Bruce M. McLaren, Eunjeong Ryu,
Ryan S.J.d. Baker, and Kenneth R. Koedinger

Human Computer Interaction Institute, Carnegie Mellon University
5000 Forbes Ave., Pittsburgh PA 15213
{iroll, eryu}@andrew.cmu.edu, {aleven, mclaren, rsbaker, koedinger}@cs.cmu.edu

Abstract. Students often use available help facilities in an unproductive fashion. To improve students' help-seeking behavior we built the Help Tutor – a domain-independent agent that can be added as an adjunct to Cognitive Tutors. Rather than making help-seeking decisions for the students, the Help Tutor teaches better help-seeking skills by tracing students actions on a (meta)cognitive help-seeking model and giving students appropriate feedback. In a classroom evaluation the Help Tutor captured help-seeking errors that were associated with poorer learning and with poorer declarative and procedural knowledge of help seeking. Also, students performed less help-seeking errors while working with the Help Tutor. However, we did not find evidence that they learned the intended help-seeking skills, or learned the domain knowledge better. A new version of the tutor that includes a self-assessment component and explicit help-seeking instruction, complementary to the metacognitive feedback, is now being evaluated.

1 Introduction

Not only that teaching metacognition holds the promise of improving current learning of the domain of interest, but also, or even mainly, it can promote future learning and successful regulation of independent learning. However, considerable evidence shows that metacognitive skills are in need of better support. For example, while working with Intelligent Tutoring Systems (ITS), students try to game the system [6] or do not self-explain enough [1].

Recently, several researchers have explicitly incorporated metacognitive support into ITS. Conati et al. [8] and Aleven et al. [1] scaffold self-explanation; Baker et al. reduce harmful gaming [6]; Bull et al. [7] and Zapata-Rivera et al. [17] encourage reflection using open learner models; and Gama offers a metacognitive suite in the form of scaffolding self-evaluation, planning and reflection [10]. While many of these components indeed improve learning, they do not focus directly on improving the subset of metacognitive skills that relates to help-seeking. Also, as far as we know, so far there was no evaluation of transfer of metacognitive skills from ITS to other learning environments.

1.1 Help seeking behavior

In this paper we focus on supporting metacognitive skills that regulate help-seeking behavior. The need for effective help-seeking strategies is apparent in many aspects of learning, formal or otherwise. The ability to seek help efficiently has been shown to contribute to learning [5; 12], and was correlated with better learning while working with ITS [16].

However, students' help-seeking behavior is often faulty (for an overview, see [4]). Students have a tendency both to overuse and under-use help: they avoid using help when they need it, but when they do seek help, they typically ask for more than is actually required [2].

In the current work we try to improve general help-seeking skills by building the Help Tutor, a domain-independent plug-in agent that can supplement a tutoring system such as a Cognitive Tutor. In this paper we describe a classroom evaluation study we conducted with the Help Tutor, having the dual goals of (a) assessing its effectiveness with respect to improving students' help-seeking behavior, skills and their learning of domain-specific skills and knowledge, and (2) learning about the requirements for and characteristics of a successful metacognitive tutoring system.

1.2 The Cognitive Tutor

The Geometry Cognitive Tutor (see Figure 1) is part of the Cognitive Tutors curriculum, a family of ITS commonly used in high schools around the United States [11]. The main window of the tutor is the Scenario window, which presents the problem and includes the main interaction with the student (on the left). The tutor scaffolds the solution process for the student by outlining the steps that are required to reach the final answer.

In the upper-right-hand corner students can see an estimation

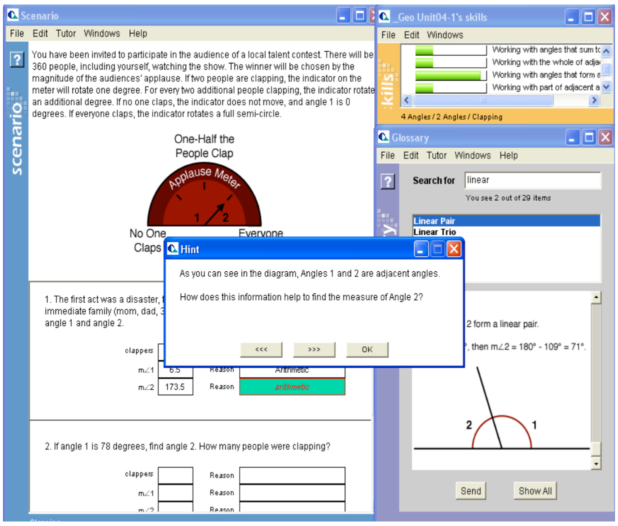

You have been invited to participate in the audience of a local talent contest. There will be 360 people, including yourself, watching the show. The winner will be chosen by the magnitude of the audiences' applause. If two people are clapping, the indicator on the meter will rotate one degree. For every two additional people clapping, the indicator rotate an additional degree. If no one claps, the indicator does not move, and angle 1 is 0 degrees. If everyone claps, the indicator rotates a full semi-circle.

1. The first act was a disaster, t immediate family (mom, dad, 3 angle 1 and angle 2.

2. If angle 1 is 78 degrees, find angle 2. How many people were clapping?

As you can see in the diagram, Angles 1 and 2 are adjacent angles.

How does this information help to find the measure of Angle 2?

Search for linear

You see 2 out of 29 items

Linear Pair

Linear Trio

2 form a linear pair.

, then m∠2 = 180° - 109° = 71°.

Scenario
File Edit Tutor Windows Help
scenario
One-Half the People Clap
Applause Meter
1
2
No One Claps
Everyone
clappers
m∠1
6.5
m∠2
173.5
Reason
Arithmetic
arithmetic
Reason
Hint
<<<
>>>
OK
_Geo Unit04-1's skills
File Edit Windows
skills
Working with angles that sum tc
Working with the whole of adja
Working with angles that form a
Working with part of adjacent a
4 Angles / 2 Angles / Clapping
Glossary
File Edit Tutor Windows Help
2
1
Send
Show All

Figure 1. The Geometry Cognitive Tutor of their knowledge level. The Cognitive Tutor estimates the student's knowledge-level on the target set of cognitive skills using a Bayesian knowledge-tracing algorithm [9].

The Geometry Cognitive Tutor has two main help-seeking mechanisms: on-demand contextual hints, and a glossary. The on-demand contextual hints provide multiple levels of information that students can browse. The first level is typically very general, intended to remind the student of their current goal. Intermediate hints are increasingly more specific. The last (or "bottom-out") hint is very close to the answer. An example of an intermediate-level hint is shown in figure 1. The second help resource is the Glossary, which resembles a searchable information source (such as the internet, or a dictionary). The Glossary contains all relevant definitions and theorems. However, searching it and applying its content require some available cognitive capacity and ability to transfer the general information to the problem, much like using many real-life information sources.

2 The Help Tutor

When designing the Help Tutor, we chose to leave all help-seeking related decisions to the students. While insisting on help or preventing it can be beneficial for learning specific domain knowledge in the short-term, it will likely not improve students' ability to deal with situations of uncertainty or knowledge gaps in the future. Thus, the goal of the Help Tutor is to teach students to be better help-seekers by offering advice, and not by limiting students to only a certain behavior [3].

Similar to the Cognitive Tutor itself, the Help Tutor supports learning-by-doing, i.e., it teaches help-seeking skills by letting students practice them and then giving appropriate feedback. The Help Tutor is a plug-in agent that is added to an existing Cognitive Tutor, and the student interacts with it during the normal course of interaction with the Cognitive Tutor. In the Help Tutor, students' actions are traced using a metacognitive help-seeking model, in addition to the existing domain-level cognitive model [2]. When a student performs a help-seeking error she receives immediate and tailored feedback, in the form of a help-seeking error message.

The Help Tutor is comprised of two conceptual components - detection and intervention [3].

Detection: The help-seeking model, which is used to trace the student's behavior, determines what the preferred action is at each moment, so the assistance-level the student gets should fit her zone of proximal development **[15]**. The model often allows for more than one correct action. For example, a student working on a step for which she is estimated to have a high skill level is expected either to attempt the step with no help or to search the glossary, while the same student, on a step for which she has a low estimated skill level, is expected to ask for an elaborated hint.

The model is implemented using eighty production rules. It marks deviations from the set of recommended actions as help-seeking errors, which can be categorized in five families:

- Help abuse – the use of hints or Glossary in an inappropriate manner (for example, by 'drilling down' hints quickly to the bottom-out hint).
- Try-step avoidance – the use of hints when the student seems sufficiently skilled to solve the step on her own.
- Try-step abuse – trying to solve in an inappropriate manner (e.g., by guessing repeatedly)
- Help-avoidance – trying to solve a step when the student should have asked for help (e.g., after making multiple errors on that step).
- General errors – other errors (e.g., the student exhausted all hints and performed high number of errors, and is still trying instead of consulting with the teacher).

An earlier version of the model was shown to be somewhat domain independent, when compared against two different Cognitive Tutors [13]. However, while it correlated with learning, it produced a much-too-high error rate – of all students' actions, 64% in one dataset and 73% in the other were classified as errors. In order to be effective, the current model should reduce the help-seeking error rate drastically, while maintaining its correlation with learning.

Intervention: The other component of the Help Tutor is the help-seeking error messages, which include only domain-independent metacognitive content for several reasons: to encourage students to focus more on the metacognitive feedback (and not be distracted by the cognitive one), to help students generalize the help-seeking skills, and to make the Help Tutor reusable with different Cognitive Tutors.

The help-seeking messages follow few principles:

- Emphasizing the usefulness of effective help seeking behavior (e.g., "it could be that another hint will do the trick for you.")
- Reinforcing correct use of the tools (e.g., "no need to hurry so much. Take your time and read the hint carefully.")
- Being positive and informal (e.g., "could be a good time to give it a try.")

In order to avoid repetitiveness in the messages displayed to students, each error can elicit several different instances of the same message. The Help Tutor messages use the existing hint-window mechanism, and are distinguished from regular hints in their font (color and type) and timing (proactive vs. reactive).

3 Evaluation

3.1 Experimental design

We evaluated the Help Tutor with 60 students from four classrooms in two high schools in the Pittsburgh area - one urban and one suburban. The students worked with the tutor for six periods. Within each school, the participating classes were taught by the same teacher, and all students were accustomed to the Cognitive Tutor, as they use it regularly in their Geometry classes.

Half of the students in each class worked with a version of the Geometry Cognitive Tutor with the Help Tutor (Help condition), and the other half worked with the Geometry Cognitive Tutor alone (Control condition). Students were counterbalanced between conditions based on their previous achievements in the Cognitive Tutor class. No instruction on help seeking was given in advance.

Students worked with the tutors twice a week. During the other three weekdays students had classroom lectures with their Geometry teachers, which focused on different topics. Due to scheduling considerations at one of the schools, only students in one school completed pre- and post-tests before and after the study (30 students). For the other school we have only the log-files of the students while interacting with the tutors.

3.2 Assessment design

In order to evaluate the Help Tutor appropriately, we defined three objectives for the Help Tutor in this study, each of which depends on the previous one.

- **Capture poor help-seeking actions:** The Help Tutor should identify faulty behaviors, while not interrupting the student's workflow too frequently. Also, the model's conclusions should preferably be generalizable across learning environments.
- **Help students improve their help-seeking behavior:** assuming that the Help Tutor successfully captures help-seeking errors, the intervention should eliminate, or reduce, these errors.
- **Improve learning**: The Help Tutor should improve learning. Fewer help-seeking errors should translate to better performance on the posttest. Overall, students should learn the domain knowledge better, as well as become better help-seekers.

In order to assess how well the Help Tutor met these objectives, we included multiple assessments:

- Students' help-seeking behavior in the tutor was assessed using log files analysis.
- Procedural help-seeking knowledge was assessed also outside the tutor's environment, using a paper test with embedded help-seeking resources. Each problem included three types of hints, counterbalanced between test forms: a problem statement only (No Hint), a problem statement with an open and free hint (Open Hint), and a problem statement with a hint covered by a sticker (Covered Hint). Students were told that removing the sticker costs 10% of their score on that specific item (see figure 2).

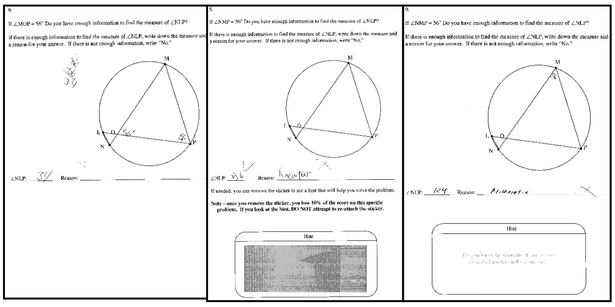

9.

If ∠MOP = 56° Do you have enough information to find the measure of ∠NLP?

If there is enough information to find the measure of ∠NLP, write down the measure and a reason for your answer. If there is not enough information, write "No."

∠NLP: 34 Reason:

5.

If ∠NMP = 56° Do you have enough information to find the measure of ∠NLP?

If there is enough information to find the measure of ∠NLP, write down the measure and a reason for your answer. If there is not enough information, write "No."

∠NLP: 56 Reason: knowfas

If needed, you can remove the sticker to see a hint that will help you solve the problem.

Note – once you remove the sticker, you lose 10% of the score on this specific problem. If you look at the hint, DO NOT attempt to re-attach the sticker.

Hint

9.

If ∠NMP = 56° Do you have enough information to find the measure of ∠NLP?

If there is enough information to find the measure of ∠NLP, write down the measure and a reason for your answer. If there is not enough information, write "No."

∠NLP: 124 Reason: Arithmetic

Hint

Figure 2. Embedded hints in pre- and post-tests. From left to right: No Hint, Covered Hint, and Open Hint conditions

- Students' declarative help-seeking knowledge was assessed using questionnaire items. Students were asked five multiple-choice questions, which described situations to which the students were supposed to respond, e.g.:

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

[ ] First I would review my calculations. Perhaps I can find the mistake myself?
[ ] The Tutor must have made a mistake. I will retype the same answer again.
[ ] I would ask for a hint, to understand my mistake.

3.3 Results

Overall students performed 59,034 actions with the tutor during the study, an average of approximately 1,000 actions, or 350 steps, per student. A typical tutor problem consists of 6-10 steps.

Although there was a significant improvement in scores from pre- to post- test, it was rather small: on average, students improved from 36% on the pretest to 41% on the posttest (F(2,28)=6.4, p=0.015). Also the log-files from the interaction with the tutor reveal rather little learning. On average, students mastered only 6 of all 49 skills that were practiced during that period.

As seen in figure 3, students scored significantly better on test items with embedded hints, compared to the No-Hint condition (t(29)=2.1, p=0.04). Students revealed hints on 24% of the Covered Hints problems.

Table 1. Help-seeking error rate and correlations with posttest scores, controlling for pretest

	Help-seeking errors overall	General Errors	Help Avoid-ance	Help Abuse	Try-Step Avoid-ance	Try-Step Abuse
Error rate	**17%**	1%	5%	6%	6%	<0.5%
Correlation with learning	**-0.42****	-0.34*	-0.41**	-0.17	-0.27	-0.10

* - marginally significant (p<0.1); ** - statistically significant (p<0.05)

Objective 1: Capturing erroneous help-seeking actions. The Help Tutor identified 17% of the actions as help-seeking errors (table 1). Higher frequency of help-seeking errors was negatively correlated with posttest scores (controlling for pretest scores). This is a significant improvement from the old model, which as noted befrore, captured about 70% of the actions and yielded similar correlation with learning.

Students' help-seeking performance as assessed by the Help Tutor also correlated with both declarative and procedural help-seeking knowledge outside the tutor environment. Scoring high on test items with embedded hints (controlling for score on items with no hints) was correlated to performing better help-seeking actions while working with the Help Tutor (r=0.5, F(2,27)=10, p<0.01). In other words, the same students who had better help-seeking skills while working with the tutor (according to the Help Tutor) also had better help-seeking procedural knowledge, as measured by comparing their scores on paper-test items with hints to those with no hints (figure 3).

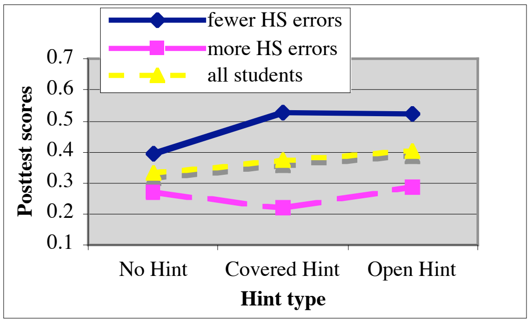

Figure 3. Better use of hints during the posttest correlates with better help-seeking behavior in the tutor

Also, students who made fewer help-seeking errors in the tutor scored significantly higher on the help-seeking multiple choice declarative knowledge questionnaire (78% vs. 59%, t(28)=2.2, p=0.04), showing that they were aware of the correct use of help.

Objective 2: Helping students improve their help-seeking behavior. In order to be effective, the Help Tutor feedback should reduce the rate of help-seeking errors students make. This objective was only partially met: on the face of it, the difference in help-seeking error rate between the Help and Control groups is not significant: 16% vs. 19% (F(2,57)=2.5, p=0.12). However, when examining students' help-seeking errors more closely, we find that the Help Tutor had a different effect on different actions: When asking for hints, students working with the Help Tutor made significantly fewer help-seeking errors (see Table). However, errors can also be made when trying, for example, when students try too fast, or avoid needed help. There was no improvement in the rate of such errors while working with the help-tutor. Perhaps the low rate of errors related to try-step does not leave much room for improvement on these errors (Controlling for school, the interaction between condition and action is significant (F(3,54) =21.0, p<0.0001)).

When looking at the context of the actions, the Help Tutor was effective only after a hint (F(3,53)=7.0, p=0.02).

These analyses show that the Help Tutor influenced students' behavior mainly during or following hint requests, and not as much on other actions. This can be best viewed when looking at the depth of hints students are viewing: The overall number of steps on which students asked to see hints was indifferent to the Help Tutor (14% for the Help group vs. 17% for the control group, not significant). However, the ratio

Table 2. Help seeking error rate per action type and context, and rate of drilling down to bottom-out hint

		Control group	Help group
Action type	Try-step	9%	9%
	Hint	45%	33%**
	(first hint)	(27%)	(17%)*
	(following hints)	(52%)	(37%)**
Context	On first action	9%	8%
	After an error	18%	18%
	After a hint	40%	31%**
% drilling down to bottom-out hint		72%	46%**

* p < 0.01; ** p < 0.001

of bottom-out hints (where students drill-down to the bottom-out hint) to all hints dropped drastically following the use of the Help Tutor: from 72% in the Control group to 46% in the Help group (F(3,53)=35, p<0.0001).

Objective 3: improve learning. Besides improving the help-seeking behavior, the Help Tutor should promote learning in both dimensions: learning of the domain knowledge, and learning of the help-seeking skills. While we observed overall learning from pre- to post-test, we were not able to identify any effect of the Help Tutor on learning (t(28)=0.1, p=0.95). Both groups improved from 36% on pretest to 41% on posttest.

While the error-rate on actions involving hints was lower for the Help group, we did not see any evidence for metacognitive learning with time – that is the help-seeking error rate was lower in the Help group throughout the study and did not have a significant learning effect. This finding suggests that rather than learning the help-seeking skills, students only followed suggestions. To evaluate this we looked at the frequency with which the Help Tutor's recommendations were followed. In this analysis we looked at the subset of actions that were performed after the Help Tutor displayed a message advising the student to act differently. We compared the actions of the Help group students to those of the Control group students in similar situations (i.e., situations in which the Help Tutor, if was used by this group of students, would have recommended them to act differently). The Help group students followed the Help Tutor recommendation when it advised them to ask for a hint (t(44)=2.5, p=0.02), but did not follow Try-Step recommendations any more than they would have done anyhow (as evaluated by the Control group).

There was also no improvement of help-seeking declarative knowledge. The differences between groups were not significant, and changes from pre- to post-test were not significant either: the Help group changed from 60% to 64% (t(16)=0.7, p=0.5); the Control group changed from 64% to 73%. (t(12)=1, p=0.3).

4 Discussion

The Help Tutor was successful in improving behavior - it captures hint usage errors, which are correlated with poorer learning, and reduces their rate significantly. Even more encouraging is the environment-independent nature of the tutor - the erroneous behavior the Help Tutor captures in the Cognitive Tutor environment is negatively correlated to successful hint-usage in the paper-and-pencil test and to declarative help-seeking knowledge. However, the Help Tutor did not yet achieve its broader goal, i.e., improving all help-seeking related actions, including faulty try-steps attempts, helping the students learn transferable help-seeking skills, and improving learning.

It appears that the Help Tutor achieved positive effects mainly because students followed its advice, and not because they assimilated the help-seeking principles. One possible explanation is the timing of the Help-Tutor messages. We did not expect students to be attentive after a successful completion of a step at the domain level, so the Help Tutor does not interfere in these instances. As a result, the Help Tutor interfered when students may have been consumed with problem solving, and thus were less likely to give the messages sufficient attention. The student might have used the Help-Tutor messages in the local scope in which they were given, to assist them in the domain level and did not internalize the rule or principle governing the specific situation. Hence, the student did not learn to evaluate her own needs and to regulate her learning. More reflective feedback at the end of the problem-solving process or before starting to solve might have been helpful.

5 Conclusions and future work

In this line of research we built a model of desired help-seeking behavior and used it to create the Help Tutor, which provides students with feedback on their help-seeking behavior in addition to any other feedback that the tutor provides. The results of the study show that the Help Tutor successfully captured help-seeking errors that were negatively correlated with learning, and furthermore, were correlated to differences in help usage in a paper and pencil test. Students who worked with the Help Tutor reduced their errors in using hints, as compared to students who used the regular tutor, but did not reduce their errors in faulty solution attempts, and did not learn better help-seeking techniques over time. We hypothesize that this might be due to inadequate preparation prior to working on the problems and a lack of a reflective process after the domain-problems were solved.

We have re-designed the Help Tutor based on the findings from this study. First, conceptual instruction on help-seeking is provided to students by the teacher using a short video in advance. The instruction focuses on successful help-seeking principles, and adopting positive dispositions towards help seeking. We have also incorporated self-assessment into the Help Tutor, which encourages students to reflect upon their needs for assistance [14]. In addition, we have attempted to improve the model by allowing it to catch more Try-Step errors. One last change is the scope of the study. To emphasize the domain-independent nature of help-seeking behavior, the current evaluation stretches across two different units of the Geometry Cognitive Tutor class.

Acknowledgements. We would like to thank Ido Jamar, Jo Bodnar, Dale Walters, Mike Konieczki, Kathy Dickensheets and Grant McKinney for their help carrying out this study. The research reported in this paper is supported by NSF Award No. IIS-0308200. The contents of the paper are solely the responsibility of the authors and do not necessarily represent the official views of the NSF.

6 References

1. Aleven, V., & Koedinger, K.R. (2002). An effective metacognitive strategy: learning by doing and explaining with a computer-based Cognitive Tutor. *Cognitive Science* 26(2), 147-79

2. Aleven, V., McLaren, B.M., Roll, I., & Koedinger, K.R. (to appear). Toward meta-cognitive tutoring: A model of help seeking with a Cognitive Tutor. *International Journal of Artificial Intelligence in Education*
3. Aleven, V., Roll, I., McLaren, B.M., Ryu, E.J., & Koedinger, K.R. (2005) An architecture to combine meta-cognitive and cognitive tutoring: Pilot testing the Help Tutor. *in proceedings of 12th International Conference on Artificial Intelligence in Education (AIED 2005),* Amsterdam, The Netherlands: IOS press.
4. Aleven, V., Stahl, E., Schworm, S., Fischer, F., & Wallace, R. (2003). Help seeking and help design in interactive learning environments. *Review of Educational Research* 73(2), 277-320
5. Arbreton, A. (1998). Student goal orientation and help-seeking strategy use, in S.A. Karabenick (Eds.), *Implications for learning and teaching* (pp. 95-116). Mahwah: Erlbaum.
6. Baker, R.S., Corbett, A.T., Koedinger, K.R., & Roll, I. (2005) Detecting When Students Game The System, Across Tutor Subjects and Classroom Cohorts. *in proceedings of User Modeling 2005,* Berlin: Springer-Verlag.
7. Bull, S., Dimitrova, V., & Brna, P. (2002). Enhancing Reflective Modelling through Communicative Interaction in Learning Environments, in P. Brna, M. Baker, K. Stenning & A. Tiberghien (Eds.), *The Role of Communication in Learning to Model* (pp. 183-211). Mahwah NJ: Lawrence Erlbaum Associates.
8. Conati, C., & VanLehn, K. (1999) Teaching meta-cognitive skills: implementation and evaluation of a tutoring system to guide self-explanation while learning from examples. *in proceedings of International Conference on Artificial Intelligence in Education,*
9. Corbett, A., & Anderson, J.R. (1995). Knowledge tracing: Modeling the acquisition of procedural knowledge. *User Modeling and User-Adapted Interaction* 4(4), 253-78
10. Gama, C. (2004) Metacognition in Interactive Learning Environments: The Reflection Assistant Model. *in proceedings of 7th Conference on Intelligent Tutoring Systems,* 668-77. Berlin: Springer-Verlag.
11. Koedinger, K.R., & Anderson, J.R. (1997). Intelligent Tutoring Goes to School in the Big City. *International Journal of Artificial Intelligence in Education* 8, 30-43
12. Renkl, A. (2002). Learning from worked-out examples: Instructional explanations supplement self-explanations. *Learning and Instruction* 12, 529-56
13. Roll, I., Baker, R.S., Aleven, V., McLaren, B.M., & Koedinger, K.R. (2005) Modeling Students' Metacognitive Errors in Two Intelligent Tutoring Systems. *in proceedings of User Modeling 2005,* 379-88. Berlin: Springer-Verlag.
14. Roll, I., Ryu, E., Sewall, J., Leber, B., McLaren, B.M., Aleven, V., & Koedinger, K.R. (submitted). Towards Teaching Metacognition: Supporting Spontaneous Self-Assessment. *in proceedings of the International Conference on Intelligent Tutoring Systems 2006*
15. Vygotsky (1978). *Mind in society* Cambridge, MA: Harvard University Press.
16. Wood, H.A., & Wood, D.J. (1999). Help seeking, learning, and contingent tutoring. *Computers and Education* 33, 153-69
17. Zapata-Rivera, D., & Greer, J.E. (2002). Exploring Various Guidance Mechanisms to Support Interaction with Inspectable Learner Models. *6th International Conference on Intelligent Tutoring Systems*, 442-52