

Do Learning by Teaching Environments with Metacognitive Support Help Students Develop Better Learning Behaviors?

John Wagster¹, Jason Tan¹, Yanna Wu¹, Gautam Biswas¹, and Dan Schwartz²

¹Department of EECS/ISIS

Vanderbilt University

Nashville, TN 37235.

²School of Education

Stanford University

Stanford, CA 94305.

Abstract

We have developed Teachable Agent environments that use learning by teaching with metacognitive support to help middle school students learn about complex science topics. To demonstrate the effectiveness of this approach, we have run studies that compare three systems where (i) students are taught by an agent, (ii) students teach a computer agent, and (iii) students teach a computer agent and receive metacognitive support while teaching. Students' activities on the system, captured in log files, were coded using six primary learning activities. In this paper, we analyze behavior fragments systematically derived from the activity sequences, and identify behaviors that correlate well with high and low student performance. Our results show that students who teach and receive metacognitive support exhibit more of the high performing behaviors than the other two groups.

Keywords: learning-by-teaching; metacognitive support, learning behaviors.

Introduction

We have been using learning by teaching models to create learning environments for middle school students that promote the development of higher-order cognitive skills for problem solving in science and math domains (Biswas, et al., 2001; 2005; Schwartz, et al., to appear). To teach, one must gain a good understanding of the domain knowledge and then structure the knowledge in a form that they can present to others (Bargh, 1980). Preparing to teach is a self-directed and open-ended activity where one explores, integrates, and structures knowledge first for oneself, and then for others. In addition to preparatory activities, teachers answer questions, and provide explanations and demonstrations during teaching and receive feedback from their students. These activities also seem to have significant cognitive consequences. For example, we might expect that the teachers' knowledge structures would become better organized and differentiated through the process of communicating key ideas and relationships to students and reflecting on students' questions and feedback (Chi, 2001). We look upon teaching as a metacognitive, reflective, and iterative process with three main phases: decision-making, performing actions, and monitoring (McAlpine, 1999).

We have designed teachable agents (TA's) that provide important structures to help shape teacher thinking (Biswas, et al, 2005; Blair, et al., 2004). TA's are software programs

where students teach a computer agent using well-structured visual representations. Using their agent's performance (which is a function of how well it was taught) as a motivation, students work to remediate the agent's knowledge, and, in this process, they learn better on their own. We discuss one of our Teachable Agent Systems, Betty's Brain, below. An important property of our TA environments is that students ideally monitor how their agents answer questions and solve problems, and they can correct them when they notice discrepancies between their own knowledge and the agent's. For this reason our *learning-by-teaching* environments are well-suited to helping students become more knowledgeable of and responsible for their own cognition and reasoning. As a result, the students are likely to develop problem solving and monitoring skills that go beyond the learning of specific domain content; rather they provide the much larger framework that guide students on *how to learn* and how to *prepare for future learning* (Schwartz and Martin, 2004).

Previous studies conducted in 5th grade science classrooms showed evidence that learning-by-teaching with metacognitive support helped students develop better learning and self-monitoring strategies, and this prepared them for future learning on related topics, even when this learning happened outside of the support provided by the TA environment (Biswas, et al., 2005). We also conjectured that the metacognitive support produced "learned behaviors" that were indicative of good learning practices. We combined intuition and empirical observations to select behaviors that we believed were indicative of independent learning with understanding. Preliminary analysis demonstrated that students in the learning by teaching condition with metacognitive feedback were more likely to demonstrate these behaviors than students who did not receive this kind of feedback. Students with these behaviors also showed better learning performance (see Tan, Biswas, and Schwartz (2006) and Tan, et al. (2007)).

In this paper we perform a more systematic statistical analysis to link student learning and observed student behaviors. The student learning measure is defined by their performance on the transfer task. As a second step, differences in the use of these behaviors between the three conditions are studied. The rest of the paper provides an overview of Betty's Brain and the metacognitive support in the system, a description of our experimental study, and a summary of our findings and future work.

Learning by Teaching: Betty's Brain

Betty's Brain is an intelligent learning environment based on the learning by teaching paradigm. The interface to the system is illustrated in Figure 1. The teaching process is implemented as three primary activities: (i) *teach*: Students explicitly teach Betty using a concept map representation, (ii) *query*: Students use a template to generate questions to see how much Betty has understood, and (iii) *quiz*: Students observe Betty's performance on a set of predefined questions. Once taught, Betty uses qualitative reasoning methods to reason through chains of links (Forbus, 1984; Biswas, et al., 2005) to answer questions, and, if asked, explain her reasoning using text and animation schemes. Betty also provides feedback that reflects the students' teaching behaviors. The goal is to get the students to adopt more metacognitive strategies in their learning tasks (Tan, Biswas, and Schwartz, 2006). Students reflect on Betty's answers and her explanations, and revise their own knowledge as they make changes to the concept maps to teach Betty better. Details of the Betty's Brain system and experiments that we have conducted with this system are summarized in (Biswas, et al., 2005; Tan, Biswas, and Schwartz, 2006). Next we discuss the metacognitive support provided to students as they learn about river ecosystems.

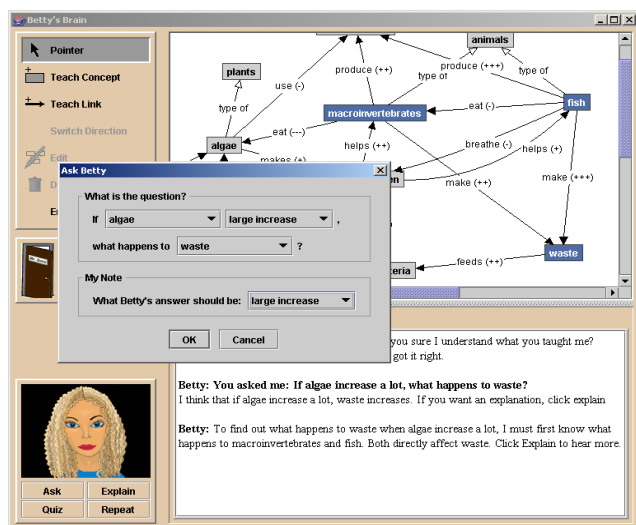


Figure 1: Betty's Brain System with Query Window

Metacognitive Support in Betty's Brain

Cognitive science researchers have established that metacognition and self-regulation are important components in developing effective learners in the classroom and beyond (Bransford, 2000, Brown, and Cocking, 2000; Butler and Winne, 1995; Zimmerman, 1989). Pintrich (2005) differentiates between two aspects of metacognition for learners: (i) *metacognitive knowledge* that includes knowledge of general strategies and when they apply, as well as knowledge of one's own abilities, and (ii) *metacognitive control and self-regulatory processes* that learners use to monitor and regulate their cognition and learning. We believe the TA environments when combined with adequate scaffolding and

feedback can provide appropriate educational opportunities for students to develop both metacognitive knowledge and control, and thereby, improve their subsequent learning.

We adopt a self-regulated learning (SRL) framework that describes a set of comprehensive skills that start with setting goals for learning new materials and applying them to problem solving tasks, deliberating about strategies to enable this learning, monitoring one's learning progress, and then revising one's knowledge, beliefs, and strategies as new materials and strategies are learnt. In conjunction with these higher level cognitive activities, social interactions and motivation also play an important role in the self-regulation process (Zimmerman, 1989). We believe that two interacting factors of our TA implementations are particularly supportive of self regulation. The first is the visual *shared representation* that the students use to teach their agents. The second factor, *shared responsibility*, targets the positive effects of social interactions to learning. This manifests in the form of a joint effort where the student has the responsibility for teaching the TA (the TA knows no more and no less than what the student teaches it), whereas the TA has the responsibility for answering questions and taking tests.

Betty's persona in the SRL version incorporates metacognitive knowledge that she conveys to the students at appropriate times to help the student develop and apply monitoring and self regulation strategies (Tan, Biswas, and Schwartz, 2006). Table 1 provides a summary of some of these self-regulation characteristics, which drive her interactions with the student. For example, when the student is building the concept map, Betty occasionally responds by demonstrating reasoning through chains of events. She may query the user, and sometimes remark (right or wrong) that the answer she is deriving does not seem to make sense. The idea of these spontaneous prompts is to get the student to reflect on what they are teaching, and perhaps, like a good teacher check on their tutee's learning progress. These interactions are directed to help Betty's student-teacher understand the importance of monitoring and being aware of one's own abilities. On other cues, the Mentor (and sometimes Betty herself) provides suggestions on cognitive strategies the students may employ to improve their own learning and understanding of the subject matter under consideration.

Experimental Design

To study the effect of metacognitive and self-regulation strategies on learning behaviors, we designed three version of the TA system. We refer to the system used in the control condition as the intelligent tutoring system (ITS) because this directed learning environment contains some aspects of the traditional ITS (Wenger, 1987). In this condition, the students were taught instead of teaching someone else. Mr. Davis, the Mentor agent, asked the students to construct a concept map to answer three sets of quiz questions. When students submitted their maps for a quiz, Mr. Davis provided corrective feedback that was based on errors in the quiz answers (Biswas, et al., 2005). System 2 was a Learning by Teaching (LBT) environment, where students were asked to teach Betty by creat-

ing a concept map. The students were told that Betty needed help to pass a test so she could join the high school science club. Students using the LBT system could query Betty to see how well she was learning, and they could ask Betty to take quizzes at any time during the teaching process. After Betty took a quiz, Mr. Davis graded the quiz, and provided Betty and her student-teacher with corrective feedback. The text of the feedback was identical to what was provided in the ITS system. System 3 was a learning-by-teaching system with Self Regulated Learning (SRL). Students in this condition also taught Betty but the primary differences between the LBT and SRL systems were in Betty’s behavior and interactions with the student, as well as the feedback that the Mentor provided after Betty took a quiz. Betty’s persona in the SRL version incorporated metacognitive knowledge (Table 1),

Table 1: Self-Regulation Patterns and Feedback

Self-Regulation Feature	Related Task or Activity	Teachable Agent and Mentor feedback
Monitoring Knowledge	Query	Betty and the Mentor encourage student to ask questions. Betty answers questions and provides explanations. Mentor suggests general debugging strategies.
Monitoring Knowledge	Quiz	The Mentor and Betty ask students to reflect on the questions not answered correctly to focus on what to learn. Mentor discourages students from using trial and error methods to get a particular answer right. Mentor advises students to reason using chain of events. Betty may refuse to take the quiz if the student has not checked to see if she has understood the new information that she has been taught.
Formative Self-Assessment	Query and Quiz	Students can ask Betty to explain their answers. Provides a collaborative environment for self-assessment.
Goal Setting	Ask Mentor	When asked, Mentor gives advice on what to study and how to study.
Keeping records and monitoring	Quiz	TA keeps track off and makes student aware of changes in quiz performance.
Seeking Information	Look up on-line resources Ask Mentor	Resources structured to help student access information by topic and by keywords. Mentor provides help when asked, or in response to Betty’s quiz performance.
Social interactions (seeking assistance) from peers	All	TA behaves more like an enthusiastic peer than a passive tutee. May suggest strategies to improve performance
Social interactions (seeking assistance) from Mentors	Mentor	When asked, Mentor volunteers advice on how to be a better learner, a better teacher, and learn from the resources. Mentor also provides situation-specific advice after TA has taken a quiz.

which she communicated to the students to help them develop and apply monitoring and self regulation strategies to aid their own learning (Tan, Biswas, and Schwartz, 2006).

Experimental Study and Results

The study was conducted in two 5th grade science classrooms in a Metro Nashville school. 53 students from the two classrooms were divided into three equal groups using a stratified sampling method based on standard achievement scores in mathematics and language. The three groups, ITS, LBT, and SRL, worked for seven 45-minute sessions over a period of two weeks to create their concept maps on aquatic ecosystems. A PFL (preparation for future learning) study (Tan, et al., 2007) was conducted approximately 8 weeks after the main study. Students were administered pre- and post-tests before and after the main study.

Analysis of Students’ Behaviors

Student activity sequences in each session of the main study were extracted from the system log files. The sequences contained six primary activities: (i) Edit Map (EM), (ii) Ask Query (AQ), (iii) Request Quiz (RQ), (iv) Resource Access (RA), (v) Request Explanation (RE), and (vi) Continue Explanation (CE). Actions where the students were adding, modifying, or deleting concepts and links in their concept map were classified as EM activities. The RQ and RA activity labels are self explanatory. Students in the LBT and SRL groups could ask Betty queries (AQ), and then check Betty’s reasoning by asking for explanations (RE). Betty’s explanations often involved multiple steps that mirrored the multiple steps used by the reasoning process to generate an answer. Betty provided an initial response to a request for an explanation (RE), and then followed it up with more details if the student clicked on the “Continue Explanation” (CE) button. The ITS group also had access to the query and explanation features for debugging their concept maps. Explanations were provided by the Mentor agent. An example activity sequence for a student working on the LBT system in one of the seven sessions appears below.

RA,EM,AQ,EM,AQ,RQ,EM,AQ,RA,EM,AQ,RQ,EM,RA,EM,RQ,RA,EM,RQ,EM,RQ,EM,RQ,RA,AQ

In previous work (Tan, Biswas, and Schwartz, 2006; Tan, et al., 2007) we used intuition and empirical observations to link behavior sequences to manifestations of metacognitive control and self regulation (Zimmerman, 1989; Pintrich, 1995). A primary finding in the earlier studies was that students who frequently exhibited the “Quiz-Edit-Quiz” behavior (defined as RQ_EM_RQ or EM_RQ_EM) were more likely to have concept maps with low scores. The pattern appeared to reflect trial and error (edit map, see if it worked using the quiz, then repeat to fix problems). On the other hand, students who asked queries to check on the changes they had made to their concept map (EM_AQ) and requested explanations after asking queries (EM_AQ_RE) were more likely to produce high scoring concept maps. Preliminary analysis showed that students in the SRL condition used the EM_AQ and

EM_AQ_RE patterns more frequently than the other groups, and the ITS group used the EM_RQ_EM pattern more often than the LBT and SRL groups. We concluded that the meta-cognitive support helped the SRL students learn good monitoring behavior. Furthermore, the SRL group also produced better concept maps than the ITS and LBT groups.

Identifying Behavior Patterns Indicative of High and Low Performing Students

In this study, we decided to adopt a more systematic approach for linking students' behavior patterns and their learning performance. One question we wanted to answer was what types of activity patterns are correlated with learning. Therefore, we correlated activity patterns in the main study phase (a) with learning at transfer, and (b) with learning during the main study phase. As discussed earlier, we used the transfer study concept map score as a measure of PFL. Our first step identifies behaviors in the main study that are indicative of high and low PFL performance. This is reinforced by finding activity patterns that correlate with main study performance, and together they help establish the most important behavioral patterns for learning. A second question we attempt to answer is whether the different instructional regimes led to different behavior patterns (and learning). Although the following analyses are only correlational, they are a preliminary method for identifying how different behavioral patterns lead to different levels of learning. In future work we will attempt to more definitely establish the causal relation between behaviors and student learning.

We define students' learning performance by the quality of their concept map at the end of the transfer (PFL) study. Concept map quality is computed as the sum of the correct concepts and correct links in the student's concept map. Concepts and links were defined to be correct if they appeared in the expert map¹ or if they were graded to be relevant by two coders because they demonstrated a correct understanding of the domain (even if they were not necessary to answer the quiz questions).

For the correlation computations, we restricted the number of considered activity patterns in the main study to lengths of two and three.² Of the 30+150=180 possible patterns of lengths 2 and 3 students used a total of 122 different patterns. The mean correlation value for these patterns with the transfer map was 0.087 (SD = 0.146). The activities with large positive correlations were associated with *high performance*, and the activities with large negative correlations were associated with *low performance*. A cutoff criterion of $M \pm 2.SD$ was used to select the highest and lowest performance patterns. Table 2 lists the activity patterns with correlation values above the high cutoff of 0.379 and Table 3 list the

activity patterns whose correlation values were below the low cut off of -0.205.

Table 2: Activity Patterns with high correlation values with Transfer Study Concept Map Score

Activity Pattern	Correlation Value
AQ_RA_EM	0.460
EM_AQ_RA	0.419
AQ_RA	0.414

The three activity patterns that correlated well with high performance included two activities: (i) RA, resource access, for seeking more information about the domain, and (ii) AQ, asking queries to check on answers generated by their concept map. Our interpretation is that students used the AQ_RA_EM and EM_AQ_RA activity patterns to check the correctness of their concept maps by asking queries and then looking up the resources to see if the answers were correct. AQ_RA_EM would imply that the students then went on to make changes in their concept maps, and EM_AQ_RA would imply that students were checking on the changes they had just made to their concept maps. We should clarify that the answers to queries were not directly available in the resources. The online resources were organized like a textbook with added hyperlink structures and keyword search features. Students had to read relevant portions of the text and infer the relations between entities that they then used to construct the concept map.

Table 3: Activity Patterns with low correlation values with Transfer Study Concept Map Score

Activity Pattern	Correlation Value
RQ_EM	-0.31
RQ_EM_RQ	-0.280
EM_RQ_EM	-0.214
AQ_EM	-0.207

Three of the four patterns that showed strong correlations with low performance, i.e., EM_RQ_EM, RQ_EM_RQ, and RQ_EM were linked to the suboptimal *Quiz-Edit-Quiz* strategy that we have discussed before (Biswas, et al., 2005; Tan, et al., 2006). AQ in the fourth pattern AQ_EM may be considered a good activity, however, the fact that students went on to directly make changes in their concept maps instead of RA (resource access) or RE (request explanation), which would have implied monitoring activities, led us to believe that these students were not using the AQ feature in a very useful way.

In previous studies, we had conjectured and demonstrated qualitatively that significant use of activity patterns that included the query and explanation mechanisms (AQ, RE, CE) was indicative of high performance. The pattern AQ_RE is the 4th highest ranked activity pattern (correlation value = 0.35) was a little below the high cutoff level. The high rank for the AQ_RE activity pattern is encouraging, but from this analysis one may conclude that the students who perform well in the PFL study use a *balanced strategy*

¹ The expert map was used by the mentor agent, Mr. Davis, to grade the students' concept maps and provide feedback. However, the students did not have access to the expert map.

² A maximum length of 3 was chosen to reduce computation time. In future work, we will look at longer behavior patterns.

of initiating their monitoring processes by asking queries and then following them up by asking for explanations (to check on the reasoning mechanisms) or reading the resources further (to check on the correctness of the answer). We will study this issue further by examining longer strings of behavior to get a more definitive answer on how good learners approach learning in new domains.

Activity Patterns from Main Study Scores

As a next step, we computed the scores for the final concept maps that the students generated in the main study. We used the same scheme as before for coding the concept map scores. Tables 4 and 5 summarize the activity patterns that showed strong positive and negative correlations with the concept maps. The mean correlation value for the main study scores was 0.097 (SD = 0.209). The high performance cut off value was 0.514, and the corresponding low performance value was -0.321.

EM_AQ_RA appears as a high performing pattern in both the PFL and main study analysis. The second significant activity pattern EM_AQ implies that students often followed their edit map activities by asking questions, but did not always follow them up with a resource access activity. Further investigation of the main study correlations showed that the EM_AQ_RE activity pattern also had a high correlation value (0.44), which confirms the *balanced strategy* approach that we discussed in the last section. Table 4 also shows that, AQ_RA had a high correlation value, 0.416, but the value was slightly below the cutoff. The other activity pattern that correlated highly with PFL scores, AQ_RA_EM had a smaller correlation (0.304) with the main study score. The related activity patterns with positive correlations were AQ_RA_RQ (0.228) and AQ_RA_AQ (0.113). These patterns are harder to explain in the context of good metacognitive strategies for learning.

Table 4: Activity Patterns highly correlated with Main Study concept map score

Activity Pattern	Correlation Value
EM_AQ_RA	0.524
EM_AQ	0.524
....
<i>AQ_RA</i>	<i>0.416</i>
<i>AQ_RA_EM</i>	<i>0.304</i>

Table 5: Activity Patterns with low correlation values with Main Study Concept Map Score

Activity Pattern	Correlation Value
RQ_EM_RQ	-0.44
EM_RQ_EM	-0.434
EM_RQ	-0.386
RQ_EM	-0.384
RA_EM	-0.332
....
<i>AQ_EM</i>	<i>-0.204</i>

Behavior Patterns by Group

Like before, we hypothesized that the metacognitive support provided to the SRL group during the study would result in the students in this group using the activity patterns indicative of high performance more frequently than the ITS and LBT groups. On the other hand, the ITS group would show more frequent use of the low performance activity patterns, which were directed toward getting the quiz answers right with minimum learning effort (see Biswas, et al., 2005; Tan, et al., 2007). We used an ANOVA to check for significant differences behaviors between the groups (see Table 6). The ANOVA was followed by post-hoc analysis using Tukey's HSD to establish pairwise differences between groups. Table 7 summarizes the results of the post-hoc analysis. Pairwise differences at the $p < 0.05$ level are marked in bold, and those significant at the $p < 0.1$ level are marked in italics.

Table 6: ANOVA Results – Behavior Differences Between Groups

Behavior	F(2, 51)	Sig
<i>AQ_RA_EM</i>	2.554	0.088
EM_AQ_RA	16.925	< 0.001
AQ_RA	3.490	0.038
AQ_EM	1.829	0.171
EM_RQ_EM	8.345	0.001
RQ_EM_RQ	8.656	0.001
RQ_EM	7.111	0.002

Table 7: Post Hoc Analysis of Pairwise Differences Between Groups Based on Behavior

Behaviors	Compared Groups	Sig.
<i>AQ_RA</i>	<i>ITS-SRL^a</i>	<i>0.070</i>
	<i>ITS-LBT^a</i>	<i>0.064</i>
	LBT-SRL	0.994
<i>AQ_RA_EM</i>	ITS-SRL	0.404
	<i>ITS-LBT^a</i>	<i>0.072</i>
	LBT-SRL	0.578
EM_AQ_RA	ITS-SRL^a	< 0.001
	ITS-LBT^a	0.003
	<i>LBT-SRL^a</i>	<i>0.088</i>
EM_RQ_EM	ITS-SRL^b	< 0.001
	<i>ITS-LBT^b</i>	<i>0.092</i>
	LBT-SRL	0.162
RQ_EM_RQ	ITS-SRL^b	< 0.001
	<i>ITS-LBT^b</i>	<i>0.075</i>
	LBT-SRL	0.169
RQ_EM	ITS-SRL	0.001
	ITS-LBT	0.225
	LBT-SRL	0.120

a - Second group performed behavior significantly more than first group
 b - First group performed behavior significantly more than second group

The results show significant differences between the SRL and ITS groups for three of the behaviors (one high performing behavior: EM_AQ_RA, and two low performing: RQ_EM_RQ and EM_RQ_EM). The only significant dif-

ference between ITS and LBT is the EM_AQ_RA pattern. If one relaxes the significance level, shown italicized, to $p < 0.1$, five patterns show significant differences between the SRL and ITS groups, five of the behavior patterns are different between the ITS and LBT groups, and there is one behavior difference between the SRL and LBT groups (EM_AQ_RA). This analysis tends to support the fact that the SRL group with metacognitive support used more high performing behavior patterns to support learning than the other two groups, and the ITS group used more of the low performing behavior patterns than the other two groups. The LBT group was in between. However, the results are not as definitive (statistically) significant as we had hoped for. The important question was whether these differences translated to better learning (i.e., generation of better concept maps).

Table 8 shows the concept map scores for each group in the main study. It is clear that the SRL students produced better concept maps (correct concepts + links) than the ITS and LBT groups. The differences in concept map quality are statistically significant.

Table 8: Concept Map Quality: Main study

Group	Main Study	
	Correct Concepts mean (sd)	Correct Links mean (sd)
ITS	9.78(2.5)	13.06(3.8)
SRL	13.68(3.1) ^{a,b}	17.89(5.0) ^a
LBT	10.71(2.6)	14.94(4.7)

a – significantly greater than ITS ($p < 0.05$)

b – significantly greater than LBT ($p < 0.05$)

Conclusions

The results of this study establish that metacognitive support does aid in more effective learning of domain content. This was reflected in the concept map quality measure, where the students who taught and received metacognitive support performed better than the students who taught and received no support. We noted that high-performing students developed a *balanced strategy* incorporating information seeking and self-monitoring, and low-performing students used the classic *Quiz-Edit-Quiz* strategy. Similarly, students who taught had better quality concept maps than students who were not taught.

Our results show that the SRL group tended to use behaviors indicative of high performance more than the ITS and LBT groups, and the ITS group used more of the behaviors that were indicative of poor performance. However, the behavior results were not as conclusive as the performance results (concept map quality). Part of the reason for this may be that the behavior sequences may need to be analyzed more thoroughly such as analyzing larger patterns. We believe a more in-depth analysis of both student behaviors and additional performance metrics or assessments we have yet to analyze will more clearly reveal the underlying differences. Also, examining the formation of these behaviors over time may lead to a better understanding of the differences between groups and learners. We, also, will continue to fo-

cus our attention on the emergence of novel behaviors used by learners.

Acknowledgments

This work has been supported by a Dept. of ED IES grant #R305H060089 and NSF REESE Award #0633856.

References

- Bargh, J.A. and Y. Schul, (1980). On the cognitive benefits of teaching. *Journal of Educational Psychology*, 72(5): 593-604.
- Biswas, G., Schwartz, D., and Bransford, J.D. (2001) Technology Support for Complex Problem Solving: From SAD Environments to AI, in *Smart Machines in Education*, Forbus, and Felto- vich, Editors. AAAI Press: Menlo Park, CA. 71-98.
- Biswas, G., Schwartz, D., Leelawong, K., and Vye, N. (2005), Learning by Teaching: A New Agent Paradigm for Educational Software, *Applied Artificial Intelligence*, 19(3), 363-392.
- Blair, K. P., Schwartz, D., Milo, L. and Mole, J. (2004), Comput- ers as Constructivist Teachable Agents. (Eds.). *Proc. 6th Intl. Conf. of the Learning Sciences*, 588-593.
- Bransford, J.D., Brown, A.L, and Cocking, R.R., (2000) *How People Learn*. National Academy Press: Washington, D.C.
- Butler, D.L. and Winne, P.H. (1995), Feedback and Self-Regulated Learning: A Theoretical Synthesis, *Review of Educational Re- search*, 65(3): 245-281.
- Chi, M.T.H., Siler, S. A., Jeong, H., Yamauchi, T., & Hausmann, R. G. (2001), Learning from Human Tutoring. *Cognitive Science*. 25(4): 471-533.
- Forbus, K. (1984), Qualitative Process Theory, *Artificial Intelli- gence*, 24(1-3): 85-168.
- McAlpine, L., Weston, C., Beauchamp, J., Wiseman, C., & Beau- champ, C. (1999), Building a Metacognitive Model of Reflec- tion. *Higher Education*, 37(2): 105-131.
- Novak, J.D. (1996), Concept mapping as a tool for improving science teaching and learning, in *Improving teaching and learn- ing in science and mathematics*, Treagust, Duit, & Fraser, eds. Teachers College Press: London. 32-43.
- Pintrich, P. R. (2005), The role of metacognitive knowledge in learning, teaching, and assessing. *Theory into Practice*, 41(4): 219-225.
- Schwartz, D. L. and Martin, T. (2004), Inventing to prepare for learning: The hidden efficiency of original student production in statistics instruction. *Cognition & Instruction*, 22: 129-184.
- Schwartz, D. L., Pilner, K. B., Biswas, G., Leelawong, K., and Da- vis, J. P. (to appear), Animation of Thought: Interactivity in the Teachable Agent Paradigm – Learning with Animation: Re- search and Implications for Design, Lowe, R. and Schnotz, W., eds., Cambridge Univ. Press.
- Tan, J., Biswas, G., and Schwartz, D. (2006), Feedback for Meta- cognitive Support in Learning by Teaching Environments. Feedback for Metacognitive Support in Learning by Teaching Environments, The 28th Annual Meeting of the Cognitive Science Society, Vancouver, Canada, 828-833.
- Tan, J., Wu, Y., Wagster, J. and Biswas, G. (2007), Effect of Me- tacognitive Support on Student Behaviors in Learning by Teach- ing Environments, *AI in Education*.
- Wenger, E. (1987) *Artificial Intelligence and Tutoring Systems*. Los Altos, California: Morgan Kaufmann Publishers.
- Zimmerman, B. J. (1989), A Social Cognitive View of Self- Regulated Academic Learning. *Journal of Educational Psychol- ogy*, 81(3): 329-339.