Incorporating Self Regulated Learning Techniques into Learning by Teaching Environments

Gautam Biswas (gautam.biswas@vanderbilt.edu), Krittaya Leelawong, Kadira Belynne, Karun Viswanath, and Nancy Vye
Department of EECS & ISIS, Box 1824 Sta B, Vanderbilt University
Nashville, TN 37235 USA

Daniel Schwartz (daniel.schwartz@stanford.edu), Joan Davis
School of Education, Stanford University
Stanford, CA 94305 USA

Abstract

This paper discusses Betty's Brain, a teachable agent in the domain of ecosystems that combines learning by teaching with self-regulation mentoring to promote deep learning and understanding. Two studies demonstrate the effectiveness of this system. The first study focused on components that define student-teacher interactions in the learning by teaching task. The second study examined the value of adding meta-cognitive strategies that governed Betty's behavior and self-regulation hints provided by a mentor agent. The study compared three versions: an intelligent tutoring version, a learning by teaching version, and a learning by teaching plus self-regulation strategies. Results indicate that the addition of the self-regulation mentor better prepared students to learn new concepts later, even when they no longer had access to the self-regulation environment.

Introduction

The recent proliferation in computer-based learning environments has produced a number of tutoring systems (Wenger, 1987) and pedagogical agents (Johnson, et al., 2000). The typical intelligent tutoring system curriculum is problem-driven. The system selects problems for the user to solve, and provides feedback on the solutions generated. The tutoring paradigm has been very successful. At the same time, it often emphasizes localized feedback, and does not always help students practice higher-order cognitive skills especially in complex domains (e.g., picking what questions to ask or how to examine resources for learning). Problem solving in complex domains requires active decision-making by learners in terms of setting learning goals and applying strategies for achieving these goals. The current paper examines ways to address these latter goals using an "intelligent" learning environment.

Our goal has been to introduce effective learning paradigms that advance the state of the art in computer-based learning systems and support students' abilities to learn, even after they leave the computer environment. Our approach has been to create environments where students teach computer agents. This paper reports the results of two studies. One study explored different features of a specific learning by teaching environment, Betty's Brain. The second study manipulated the metacognitive support students received when teaching "Betty" and measured its effects on the students' abilities to subsequently learn new content several weeks later.

The cognitive science and education research literature supports the idea that teaching others is a powerful way to learn. Research in reciprocal teaching, peer-assisted tutoring, programming, small-group interaction, and self-explanation hint at the potential of *learning by teaching* (Palinscar & Brown, 1984; Cohen, et al. 1982; Papert, 1993; Chi, et al., 1994). Bargh and Schul (1980) found that people who prepared to teach others to take a quiz on a passage learned better than those who prepared to take the quiz themselves. The literature on tutoring has shown that tutors benefit as much from tutoring as their tutees (Chi, et al., 2001; Graesser, et al., 1995). Biswas et al. (2001) report that students preparing to teach made statements about how the responsibility to teach forced them to gain deeper understanding of the materials. Other students focused on the importance of having a clear conceptual organization of the materials. Additionally, teachers can provide explanations and demonstrations during teaching and receive questions and feedback from students. These activities seem significant from the standpoint of their cognitive consequences in improving understanding of complex concepts.

A key benefit of the learning by teaching process focuses on the need to structure knowledge in a compact and communicable format. This requires a level of abstraction that may help the teacher develop important explanatory structures for the domain. For example, many people find that preparing a conference presentation helps them decide which concepts deserve the "high level" status of introductory framing. The need to structure ideas not only occurs in preparation for teaching, but can also occur when teaching. Good learners bring structure to a domain by asking the right questions to develop a systematic flow for their reasoning. Good teachers build on the learners' knowledge to organize information, and in the process, they find new knowledge organizations, and better ways for interpreting and using these organizations in problem solving tasks.

Despite its potential benefits, learning-by-teaching can initially seem inefficient. For example, students may need to learn the right way to teach, which can slow down their learning of the subject matter in the short run. At the same time, learning-by-teaching may have long-term benefits in that it helps students appreciate what a complete and communicable answer needs to look like, and they may learn how to consult resources to understand deeply enough that they can teach well. In this case, it seems important to evaluate not only how well students learn the target knowledge of the teaching episode, but also how well they are prepared to learn in the future as a result of learning-by-teaching (Bransford & Schwartz, 1999).

We have adopted a new approach to designing learning-by-teaching environments that ideally supports the learning outcomes described above, provide tools that enable users to visually organize and reason about their domain knowledge as they teach a computer agent, and include feedback to promote better self-regulation during the learning and teaching processes. A key challenge to the learning-by-teaching approach is that students are usually novices with regard to domain content and teaching tasks. To help with the domain content, our design includes content-integrated instruction that encourages students to access and think about resources, and check their reasoning during the teaching (and learning) process by interacting with the teachable agent and assessing its performance. To help with the teaching and learning aspects, we have made the computer agent more participatory in the learning process, and developed a Mentor agent that acts as a "meta-cognitive" coach, and provides strategy and content feedback about teaching with understanding, while avoiding the very specific localized feedback that is characteristic of many tutoring systems. Ideally the combination of the two can help students not only learn the content of a specific lesson, but also prepare students to learn in the future when they no longer have access to the system.

Implementing Learning by Teaching Systems

Our teachable agents (TAs) provide important structures to help shape the thinking of the *learner-as-teacher*. Each agent manifests a visual structure that is tailored to a specific form of knowledge organization and inference. In general, our agents try to embody four principles of design:

- Teach through visual representations that organize the reasoning structures of the domain (e.g., directed graphs and matrices).
- Build on well-known teaching interactions to organize student activity (e.g., teaching by "laying out," teaching by example, teaching by telling, teaching by modeling).
- Ensure the agents have independent performances that provide feedback on how well they have been taught (each agent uses a distinct AI reasoning technique, such as qualitative reasoning, logic, and genetic algorithms).
- Keep the start-up costs of teaching the agent very low (as compared to programming). This occurs by only implementing one modeling structure with its associated reasoning mechanisms.

Betty's Brain makes her qualitative reasoning visible through a dynamic, directed graph called a *concept map* (Novak, 1996). The fact that TAs represent knowledge structures rather than the referent domain is a departure from many simulation-based learning environments. Simulations often show the behavior of a physical system, for example, how an algal bloom increases the death of fish. On the other hand, TAs simulate the behavior of a person's thoughts about a system. Learning empirical facts is important, but learning to use the expert structure that organizes those facts is equally important. Therefore, we have structured the agents to simulate particular forms of thought that may help teacher-students structure their thinking about a domain.

Fig. 1 illustrates the interface of Betty's Brain. Students explicitly teach Betty using a graphical drag and drop interface to create and modify their concept maps in the top pane of the window. They use the *Teach Concept* button to create new concepts, and the *Teach Link* button to create relations between concepts. When teaching the agent about relationships, students use a popup template to specify the name (e.g., *breathe*, *produce*, *helps*) and type of relationship *(causal*, *type of*, and *descriptive*). For causal relations, students indicate whether the relation implies an *increase* (++) or *decrease* (—). For example, in the map in Fig.1, the concept map implies an increase in *fish* will result in a decrease in *dissolved oxygen*. Note that the student generates all concept and relationship names. They are not chosen from a menu.

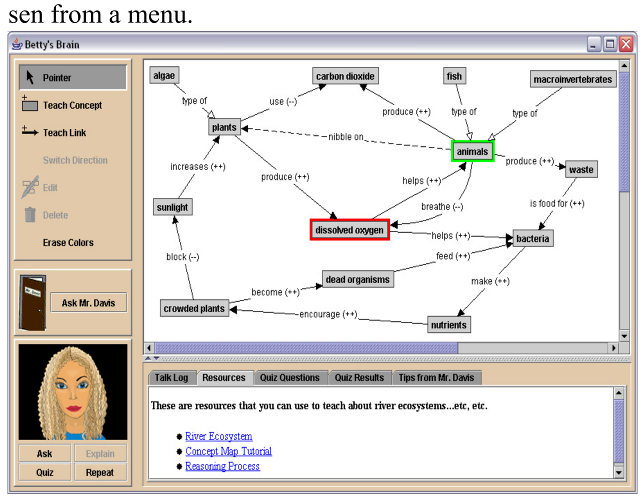

Once taught, Betty reasons with her knowledge and answers questions. Users can formulate their own queries using the *Ask* button, and observe the effects of their teaching by analyzing Betty's responses. Templates are provided to ask Betty two kinds of questions: (i) *If <concept A> increases (decreases) what happens to <concept B>?* and (ii) *Tell me all you know about <concept A>*. For the latter question, Betty enumerates all the concepts that are directly linked to *<concept A>*. For the former question, Betty uses qualitative reasoning methods to derive her answers to question through a chain of causal inferences. For example, using the concept map in Fig. 1, Betty can conclude that an increase in *algae* will cause *fish* to increase.

Betty also provides explanations for how she derives her answers by depicting the derivation process using multiple modalities: text, animation, and speech. Details of the reasoning and explanation mechanisms in Betty's Brain are presented elsewhere (Leelawong, et al., 2001).

We should clarify that Betty does not use machine learning algorithms to achieve automated learning from examples, explanations, and induction. Our focus is on the well-defined schemes associated with teaching that support a process of instruction, assessment, and remediation. These schemas help organize student interaction with the computer, much as people's well-defined schemas for spatial organizations helped to create the desktop metaphor for windows-based computer systems.

The system also includes sets of teacher-generated quiz questions. Betty can take the quiz, and students see how she performs and receive the correct answer. The quiz questions are structured to provide students cues on concepts and relations that are important in the domain of study. Examples of some quiz questions are shown in Fig. 2.

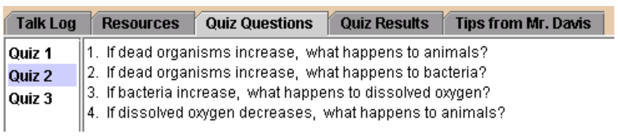

Figure 2: Quiz Questions

A Prior Study without a Self-Regulation Mentor

To study the effectiveness of Betty's Brain we conducted an experiment on 50 high-achieving fifth grade students from a science class in an urban public school located in a southeastern US city. The students were asked to teach Betty about river ecosystems. We examined the effects of the interactive features of the teachable agent environment using a 2×2 between-subjects design. One group of students could submit their agent to take a Quiz (and receive feedback on the correct answer). A second group could Query their agent by generating their own questions and seeing how Betty chains through her map to reach the answer (there was no expert feedback on the answer). The third condition, which could neither Query nor Quiz the agent, was basically using a graphing package. Students who had both Query and Quiz features could ask Betty questions and see her performance on the quiz questions. Students were given instructions on how to use the system, and then they used the software for 3 one-hour sessions. To help students learn what to teach, reference materials were made available during and in between their teaching sessions with Betty.

We hypothesized that having the query feature would help students debug their own thinking and reasoning in the problem domain, and this would result in maps with more inter-linked concepts. Betty's answers and her explanations would make explicit the process of reasoning across chains of links in a concept map. For the Quiz condition, we expected that students would map backward from the quiz questions and use the feedback they received about her answers to produce more accurate concept maps.

Analysis of the scope of students' maps and the types and accuracy of links contained therein are presented in (Leelawong, et al., 2002). On the positive side, students who used the Query and/or Quiz mechanisms understood causal relations better than the students who did not. This was reflected in their concept maps, which had a larger proportion of causal links than the No Quiz and No Query group. As predicted, students who had access to the Query feature had the most inter-linked maps and most elaborate reasoning chains. The Quiz feature was effective in helping students decide the important domain concepts and types of relationships to teach Betty.

We also noted some negative aspects to our system. Our observations of students during the study suggested that students who had the quiz feature were too focused on *"getting the quiz questions correct"* rather than *"making sure that Betty (and they themselves) understood the information"* (Davis, et al., 2003). The activity logs of the students who used the quiz showed a pattern of quick one-link corrections followed by a retake of the quiz. The query mechanism and resources were used sparsely, and it is unlikely they gained a deep understanding of causal structures. On the other hand, the Query-only group spent more time with Betty's explanations and reading resources. Surprisingly, students who had the query feature without the benefit of quiz feedback produced as many valid relevant causal links as the conditions with the quiz and quiz and query feature. This demonstrated the value of explicitly illustrating the reasoning process (by having Betty explain her answers) so that students understand causal structures.

Reflections on these results made us rethink our design and implementation of TA environments. A primary concern was the student's focus on getting quiz questions right without trying to gain an understanding of interdependence and balance in river ecosystems. We realized that interactions between the student-as-teacher, Betty, and the quiz feature had to be improved to facilitate better learning. Further, in exit interviews, students emphasized that they would have liked Betty to be more active and exhibit characteristics of a good student during the teaching phase (Davis, et al., 2003). Several students suggested that we should "*do some sort of game or something and make the system more interactive,*" and "*Betty should react to what she was being taught, and take the initiative and ask more questions on her own.*" Consistent with this feedback, we noted that the first version of Betty was passive and only responded when asked questions. We believed that to create a true learning by teaching environment, Betty needed to better demonstrate qualities of human students. A tutor gains deeper understanding from interactions with a tutee (Chi, et al., 2001) that includes answering the tutee's questions, explaining materials, and discovering misconceptions. Betty should be designed to benefit her users in the same way.

Self-Regulated Learning and Betty's Brain

As mentioned earlier, an important realization from this first study was that we were dealing with young children who were novices in teaching practice and in domain knowledge content. To accommodate this, the learning environment was redesigned to provide appropriate scaffolds and proper feedback mechanisms to help students overcome their initial difficulties in learning and teaching about a complex domain. The scaffolds took on three primary forms. First, we made improvements in the online resources available for learning about river ecosystems. We reorganized the resources to emphasize the concepts of interdependence and balance. This changed the partitioning of the resources to the three primary cycles that govern ecosystem behavior: (i) the oxygen/carbon dioxide cycle, (ii) the food chain, and (iii) the decomposition cycle. A hypertext implementation allowed direct access to sections and subsections. An advanced keyword search technique provided access to information using keywords. (Students in the study below found the resources to be much more useful, and used them extensively while teaching Betty.)

The second change is that we redesigned the quiz so that the questions would support users in systematically building their knowledge about river eco-systems. The questions were no longer randomly sampled from the full domain, but they gradually introduced more complex questions. Furthermore, the first item in each quiz was a comprehensive question that covered all of the domain concepts and relations associated with a particular cycle. This prevented students from taking a sequential approach of building the concept map to answer one question at a time. We also improved the feedback the students received.

These two changes were important, but we doubted they would be sufficient in supporting users in becoming better learners and teachers, nor did they address our users requests for a more *"life like"* Betty. Therefore, our third change, and most relevant to the study below, was to make Betty more reactive to what she was being taught, as well as to use self-regulation strategies in her interactions with her student-teacher. Along with this, we added a mentor agent to the system to help users observe and develop metacognitive and self-regulation strategies to support active and independent learning. Self-regulated learning should be an effective framework for providing feedback because it promotes the development of higher-order cognitive skills (Corno & Mandinach, 1983), and it is critical to the development of problem solving ability (Pintrich & DeGroot, 1990).

Our new design adopted some aspects of the framework of *self-regulated learning,* described by Zimmerman (1989) as situations where students are "*metacognitively, motivationally, and behaviorally participants in their own learning process.*" Self-regulated learning strategies involve actions and processes that can help one to acquire knowledge and develop problem solving skills (Pintrich & DeGroot, 1990). Zimmerman describes a number of self-regulated learning skills that include goal setting and planning, seeking and organizing information, keeping records and monitoring, and self-evaluation. We developed mechanisms by which Betty forced the student to conform to the self-regulation strategies. In parallel, the Mentor agent included resources that helped students develop these skills during their learning and teaching.

This resulted in a number of changes to Betty's Brain. For example, when a student begins the teach phase by constructing the initial concept map, both the Mentor and Betty make suggestions that the student *set goals* about what to teach, and make efforts to gain the relevant knowledge by studying the river ecosystem resources. The Mentor continues to emphasize the reading and understanding of resources, whenever the student has questions on *how to improve their learning*. The user is given the opportunity to *evaluate her knowledge* while studying. If she is not satisfied with her understanding, she may *seek further information* by asking the Mentor for additional help. While teaching, the student as teacher can interact with Betty in many ways, such as asking her questions (*querying*), and getting her to take *quizzes* to evaluate her performance. Users are given a chance to predict how Betty will answer a question so they can check what Betty learned against what they were trying to teach.

Some of the self-regulation strategies manifest through Betty's persona. These strategies make Betty more involved during the teach phase, and drive her interactions and dialog with the student. For example, during concept map creation, Betty spontaneously tries to demonstrate *chains of reasoning*, and the conclusions she draws from this reasoning process. She may query the user, and sometimes remark (right or wrong) that an answer she is deriving does not seem to make sense. This is likely to make users reflect on what they are teaching, and perhaps, like good teachers they will assess Betty's learning progress more often. At other times, Betty will prompt the user to *formulate queries* to check if her reasoning with the concept map produces correct results. There are situations when Betty emphatically refuses to take a quiz because she feels that she has *not been taught enough*, or that the student has not given her *sufficient practice by asking queries* before making her take a quiz.

After Betty takes a quiz offered by the Mentor agent, she discusses the results with the user. Betty reports: (i) her view of her performance on the particular quiz, and if her performance has improved or deteriorated from the last time she took the quiz, and (ii) the Mentor's comments on Betty's performance in the quiz, such as: "*Hi, I'm back. I'm feeling bad because I could not answer some questions in the quiz. Mr. Davis said that you can ask him if you need more information about river ecosystems.*") The Mentor agent's initial comments are general, but they become more specific if errors persist, or if the student seeks further help (*"You may want to study the role of bacteria in the river"*).

In addition to self-regulation advice that included information on how to be a better learner and better teacher, the domain content feedback from the Mentor agent was directed to make the student think more about interdependence among concepts. Students seeking specific help were first directed to relevant sections in the resources for further study and reflection, rather than being told what was wrong in their concept maps. When the Mentor provided specific feedback, it was about *chains of events* to help students better understand *chains of events* to help students better understand Betty's reasoning processes.

Overall, we believe that the introduction of self-regulation strategies provides useful scaffolds to help students learn about a complex domain, while also developing metacognitive strategies that promote deep understanding and abilities to learn in the future. One of the achievements of the new system is that students retain control rather than being told what to do (e.g., they need to request help from the mentor and they teach Betty). Only when the student seems to be hopelessly stuck, does the Mentor spontaneously intervene to help students advance in their learning (and teaching) task.

A Study of the Added-Value of Self-Regulation

A new experiment with fifth graders was designed to compare the Teachable Agent system with the self regulation mentor (SRL) against two other approaches: (i) A learning by teaching (LBT) version that was similar to the Query & Quiz version before, and (ii) An externally-guided learning system (ITS) designed with a pedagogical agent. In the ITS version, the pedagogical agent asked students to create concept maps that could answer a set of quiz questions (therefore, there was no *teaching* component), and the agent would provide feedback on how to correct their map when their quiz answers had errors. All three groups had access to identical resources on river ecosystems and the same query and quiz features. To evaluate student learning, we examined pre-posttest scores, how they used the system, the quality of their final maps, and their ability to reproduce the maps subsequently. Importantly, several weeks later, we asked the students to learn about the Nitrogen cycle, which had not been covered during the initial instruction. This permitted us to determine which group had been better prepared to learn, once they no longer could rely on the scaffolds of their respective version. Our expectation was that the SRL students would do better on this latter measure of preparation for future learning (Bransford & Schwartz, 1999), because they had learned how to "take charge" of their own learning.

Experimental Procedure

A fifth grade classroom was divided into three equal groups of 15 students each using a stratified sampling method based on standard achievement scores in mathematics and language. The students worked on a pretest with twelve questions before they were separately introduced to their particular versions of the system. The three groups worked for six 45-minute sessions over a period of three weeks to create their concept maps. All groups had access to the online resources while they worked on the system.

All three conditions had the same quiz questions while working with the system, and they had access to the query feature and Mentor agent (Mr. Davis), though he appeared with different capacities. The task given to the ITS group was to create concept maps that correctly answered the 16 questions that were divided up into three quizzes. They had the same interface to create and modify their concept maps as the other groups, but Betty did not exist in the ITS system. The ITS feedback came from the Mentor, who told students if their map held the correct answers to the quiz questions and provided hints on how the students could correct their maps. The two other groups, LBT and SRL, were told to teach Betty and help her pass a test so she could become a member of the school Science club. Both of these groups had access to the three quizzes. The LBT group only received mentor feedback about the quality of Betty's specific answers to the quiz. The SRL group received more extensive feedback from the Mentor, but only when they queried him. Coupled with the Mentor, the SRL Betty was also endowed with self-regulation strategies that governed her behavior. Therefore, the SRL condition was set up to develop more active learners by promoting the use of self-regulation strategies.

At the end of the six sessions, every student took a post-test that was identical to the pretest. Two other delayed post-tests were conducted about seven weeks after the initial experiment: (i) a *memory test*, where students were asked to recreate their ecosystem concept maps from memory (there was no help or intervention when performing this task), and (ii) a *preparation for future learning transfer test*, where they were asked to construct a concept map using on-line resources and answer questions about the land-based nitrogen cycle. Students had not been taught about the nitrogen cycle, so they would have to learn from resources during the transfer phase. (All three conditions simply used the concept mapping interface, resources, and "correct/incorrect" feedback from the mentor on several quiz questions.)

For learning about river ecosystems, students in all conditions improved from pre- to posttest on their

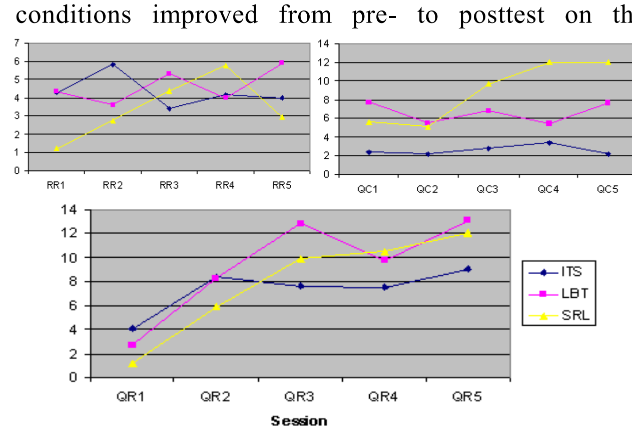

Figure 3: Resource Requests (RR), Queries Composed (QC), & Quizzes Requested QR) per session.

knowledge of interdependence (*p's*<.01, paired T-tests), but not ecosystem balance. There were few differences between conditions in terms of the quality of their maps. However, there were notable differences in their use of the system during the initial learning phase. Fig. 3 shows the average number of resource, query, and quiz requests per session by the three groups. It is clear from the plots that the SRL group made a slow start as compared to the other two groups. This can primarily be attributed to the nature of the feedback; i.e., the ITS and LBT groups received specific content feedback after a quiz, whereas the SRL group tended to receive more generic feedback that focused on self-regulation strategies. Moreover, in the SRL condition, Betty would refuse to take a quiz unless she felt the user had taught her enough, and prepared her for the quiz by asking questions. After a couple of sessions the SRL group showed a surge in map creation and map analysis activities, and their final concept maps and quiz performance were comparable to the other groups. It seems the SRL group spent their first few sessions in learning self-regulation strategies, but once they learned them their performance improved significantly.

For the delayed memory test, the table below presents the mean number of expert causal links and concepts in the student maps. Results of ANOVAs using Tukey's LSD to make pairwise comparisons showed that the SRL group recalled significantly more links that were also in the expert map (which nobody actually saw).

Student Map Included:	SRL Mean (se)	LBT Mean (se)	ITS Mean (se)
Expert Concepts	6.7 (.6)	6.4 (.5)	5.8 (.6)
Expert Causal Links	3.3[a] (.6)	1.7 (.6)	2.0 (.6)

[a] Significantly greater than LBT, $p < .05$

We thought that the effect of SRL would not be to improve memory, but rather to provide students with more skills for learning subsequently. When one looks at the results of the test of preparation for future learning, the differences between the SRL group and the other two groups are significant. The table below summarizes the results of the transfer test, where students read resources and created a concept map for the land-based nitrogen cycle. There are significant differences in the number of expert concepts in the SRL versus ITS group maps, and the SRL group had significantly more expert causal links than the LBT and ITS groups. When learning about the river ecology, the SRL students had received some guidance in how to use resources productively and how to think about the quality of their map. This guidance transferred to learning about the nitrogen cycle.

Student Map Included:	SRL Mean (sd)	LBT Mean (sd)	ITS Mean (sd)
Expert Concepts	6.1[a] (.6)	5.2 (.5)	4.1 (.6)
Expert Causal Links	1.1[ab] (.3)	0.1 (.3)	0.2 (.3)

[a] Significantly greater than ITS, $p < .05$;
[b] Significantly greater than LBT, $p < .05$

Conclusions

The results demonstrate the significant positive effects of SRL strategies in understanding and transfer in a learning by teaching environment. Students in all three groups demonstrated the same learning performance in traditional learning tasks, but the SRL group outperformed the other two in the far transfer test. We believe that the differences between the SRL and the other two groups would have been more pronounced if the transfer test study had been conducted over a longer period of time. Lastly, we believe that the concept map and reasoning schemes have to be extended to include temporal reasoning and cycles of behavior to facilitate students' learning about the concept of balance in ecosystems.

References

Bargh, J. A., & Schul, Y. (1980). On the cognitive benefits of teaching. *Journal of Educational Psychology, 72*(5), 593-604.

Biswas, G., Schwartz, D., Bransford, J., & TAG-V. (2001). Technology Support for Complex Problem Solving: From SAD Environments to AI. In Forbus & Feltovich (eds.), *Smart Machines in Education* (pp. 71-98). Menlo Park, CA: AAAI Press.

Bransford, J. D., & Schwartz, D. L. (1999). Rethinking transfer: A simple proposal with multiple implications. In A. Iran-Nejad & P. D. Pearson (Eds.), *Review of Research in Education* , **24**, 61-101. Washington DC: American Educational Research Association.

Chi, M. T. H., De Leeuw, N., Mei-Hung, C., & Levancher, C. (1994). Eliciting self explanations. *Cognitive Science, 18*, 439-477.

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.

Cohen, P. A., Kulik, J. A., & Kulik, C.-L. C. (1982). Educational outcomes of peer tutoring: A meta-analysis of findings. *American Educational Research Journal, 19*(2), 237-248.

Corno, L., & Mandinach, E. B. (1983). The role of cognitive engagement in classroom learning and motivation. *Educational Psychology, 19*(2), 88-108.

Davis, J. M., Leelawong, K., Belynne, K., Bodenheimer, R., Biswas, G., Vye, N., et al. (2003). Intelligent User Interface Design for Teachable Agent Systems. *Proc. Intl. Conf. on Intelligent User Interfaces*, Miami, Florida, 26-34.

Graesser, A. C., Person, N., & Magliano, J. (1995). Collaborative dialog patterns in naturalistic one-on-one tutoring. *Applied Cognitive Psychologist, 9*, 359-387.

Leelawong, K., Davis, J., Vye, N., Biswas, G., Schwartz, D., Belynne, K., et al. (2002). The Effects of Feedback in Supporting Learning by Teaching in a Teachable Agent Environment. *Proc. Fifth Intl. Conference of the Learning Sciences*, Seattle, Washington, 245-252.

Leelawong, K., Wang, Y., Biswas, G., Vye, N., & Bransford, J. (2001). Qualitative reasoning techniques to support learning by teaching: The Teachable Agents project. *Proc. Fifteenth Intl. Workshop on Qualitative Reasoning*, San Antonio, Texas, 73-80.

Novak, J. D. (1996). Concept Mapping as a tool for improving science teaching and learning. In D. F. Treagust, R. Duit & B. J. Fraser (eds.), *Improving Teaching and Learning in Science and Mathematics* (pp. 32-43). London: Teachers College Press.

Palincsar, A. S. & Brown, A. L. (1984). Reciprocal teaching of comprehension-fostering and comprehension -monitoring activities. *Cognition and instruction, 1*, 117-175.

Papert, S. (1993). *The Children's Machine: Rethinking school in the age of the computer*. New York, NY: Basic Books.

Pintrich, P. R., & DeGroot, E. V. (1990). Motivational and self-regulated learning components of classroom academic performance. *Journal of Educational Psychology, 82*(1), 33-40.

Zimmerman, B. J. (1989). A Social Cognitive View of Self-Regulated Academic Learning. *Journal of Educational Psychology, 81*(3), 329-339.

Acknowledgements

This work has been supported by a NSF ROLE 0231771 & 0231946. The assistance provided by John Bransford and John Kilby is gratefully acknowledged.