The Effects of Feedback in Supporting Learning by Teaching in a Teachable Agent Environment

Krittaya Leelawong, Joan Davis, Nancy Vye, Gautam Biswas
Box 1679-Station B, Vanderbilt University, Nashville, TN 37235
Tel: 615-322-7763, Fax: 615-343-5459
kiktaya@vuse.vanderbilt.edu, joan.davis@vanderbilt.edu, nancy.vye@vanderbilt.edu, biswas@vuse.vanderbilt.edu

Daniel Schwartz
Stanford University
daniel.schwartz@stanford.edu

Kadira Belynne, Thomas Katzlberger & John Bransford
Vanderbilt University
kadira.belynne@vanderbilt.edu, katzlbt@vuse.vanderbilt.edu, john.bransford@vanderbilt.edu

Abstract: The idea that teaching others is a powerful way to learn is both intuitively compelling, and one that has garnered support in the research literature. The present study investigates aspects of the "learning by teaching" process that contribute to enhanced learning outcomes for students. We developed a computer-based teachable "agent" that students explicitly teach using concept maps. Results indicate that providing students with opportunities to quiz their agent decreases the amount of irrelevant information and increases the proportion of causal information in students' maps, whereas having opportunities to query their agent increases the interconnectedness of concepts in students' maps. The results point to the importance of including various forms of feedback in designing teachable agent environments that promote learning.

Introduction

The idea that teaching others is a powerful way to learn is both intuitively compelling, and one that has garnered support in the research literature. For example, Bargh and Schul (1980) found that people who prepared to teach others to take a quiz on a passage learned the passage better that those who prepared to take the quiz themselves. The literature on tutoring suggests a similar conclusion in that tutors have been shown to benefit as much from tutoring as their tutees (Graesser, Person, & Magliano, 1995; Chi, Siler, Jeong, Yamauchi, & Hausmann, 2001). Biswas and colleagues (Biswas, Schwartz, Bransford, & TAG-V, 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.

Reflection on these studies and others lead us to conjecture that the creation of a computer-based system, where students can assume the role of "teacher," may provide an effective and motivating environment for learning. We designed an environment that lets students explicitly teach a computer agent. Once taught, the agent reasons with its knowledge and answers questions. Students observe the effects of their teaching by analyzing these responses and by getting additional feedback from a teaching expert.

Although the research literature (e.g., Bargh & Schul, 1980) suggests that learning benefits accrue from preparing to teach, it is not clear if and how other aspects of the learning by teaching process contribute to enhanced outcomes. In addition to preparatory activities, teachers provide explanations and demonstrations during teaching and receive questions and feedback from students. These activities also seem significant from the standpoint of their cognitive consequences. For example, we might expect that 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, et al., 2001).

The purpose of the present study was to examine the learning benefits of different activities associated with "learning by teaching" in our teachable agent environment. This was done by constructing computer-based agents that students could teach domain knowledge. In particular, we created an agent environment called Betty's Brain which can operate in three modes: (i) the TEACH mode, where students impart knowledge to the agent Betty by means of a dynamic concept map interface, and access content materials as needed to learn information for teaching, (ii) the QUERY mode, where students ask Betty questions (using question templates) which she answers by reasoning with information that the student has taught her, and (iii) the QUIZ mode, where students evaluate how well they have taught Betty by observing her performance on a quiz. At times, an expert teacher agent intervenes to make suggestions that may help Betty (and the student) correct her answers.

In the present study we examined the effects of the interactive features of the teachable agent environment that emulate the feedback that instructors receive from students during teaching. All students had the opportunity to TEACH their agent, and we manipulated whether students could QUERY Betty and observe her QUIZ performance following their teaching efforts. Crossing these variables created four versions of the teachable agent environment: 1. TEACH Only version (No QUERY or QUIZ), 2. QUERY version, 3. QUIZ version and 4. FULL version (QUERY & QUIZ).

We hypothesized that having opportunities to query and/or quiz Betty would positively, but differentially, impact students' learning. The query feature helps students debug their own thinking and reasoning in the problem domain. If Betty answers questions in unexpected ways, students know that they need to add to or modify their concept maps. In addition, and perhaps more important, when Betty explains her answers, she makes explicit the process of reasoning across links in a concept map (i.e., infer the effect of one concept on another through a chain of relations). Therefore, we might expect that students who use the QUERY versions of the software would create maps containing more inter-linked concepts. With respect to the quiz condition, we expected that students would become better at identifying important concepts and links to include in their maps because they could map backward from the quiz questions. We also expected that overall they would produce more accurate concept maps because they had access to feedback on Betty's quiz performance.

Methods

Description of Betty's Brain

Figure 1 illustrates the Betty's Brain interface. The system possesses multimedia capabilities. Students use a graphical drag and drop interface to create and modify their concept maps. When queried, Betty can provide explanations for how she derives her answers, and simultaneously depict the derivation process on the concept map by animation. In the sections below, we describe the software's 3 modes: TEACH, QUERY and QUIZ.

TEACH Betty

Students teach Betty by means of a concept map interface. A concept map is a collection of concepts and relations between these concepts (Novak, 1996). A relation is a unidirectional link connecting two entities. Concept maps help to categorize groups of objects and express interactions among them. They also provide a mechanism for representing knowledge hierarchies and cause-effect relations (Stoyanov and Kommers 1999). This makes the concept-mapping technique very amenable to applications in scientific domains, in particular, for modeling dynamic systems. In this study, we asked students to teach Betty about rivers--the living and non-living things in a river and how living things survive. In this way, the focus is on creating a model of a river ecosystem and on ideas of balance and interdependence.

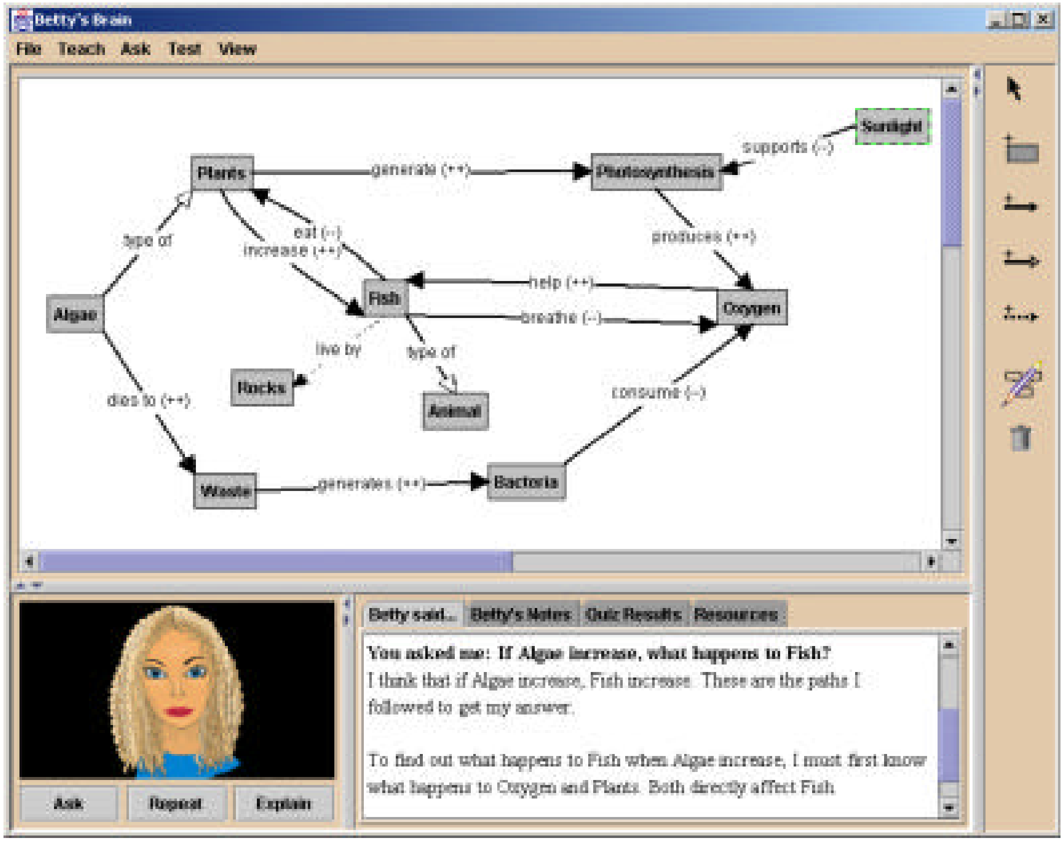

Figure 1. Betty's Brain Interface

Figure 1 displays an example of a concept map that student created in Betty's Brain—the map represents what the student has taught Betty. Note that this map is not totally accurate or complete. The labeled boxes correspond to concepts (the labels are concept names), and the labeled links correspond to relations. Students can use three kinds of links, (i) causal, (ii) hierarchical, and (iii) descriptive. Students use descriptive links to embed notes or interesting characteristics of an object in their concept map (e.g., "Fish live by Rocks"). Hierarchical links let students establish class structures to organize domain knowledge (e.g., "Fish is a type of Animal").

A causal link specifies an active relationship on how a change in the originating concept affects the destination concept. Two examples of this type of relation are "Fish eat Plants" and "Photosynthesis produces Oxygen". The causal relations are further qualified by increase ('++') and decrease ('--') labels. For example, "eat" implies a decrease relation, and "produce" an increase. Therefore, an introduction of more fish into the ecosystem causes a decrease in the number of plants, but an increase in the number of plants causes an increase in oxygen.

QUERY Betty

Students are able to query Betty about what they have taught her. The query mode consists of two mechanisms: (i) a reasoning mechanism and (ii) an explanation mechanism. The reasoning mechanism enables Betty to analyze the knowledge that the student has taught her so that she can answer questions. The explanation mechanism enables Betty to produce a detailed explanation of how she generated her answer. Currently, Betty's Brain has templates for four question types:

Type 1: What will happen to Concept A when we increase/decrease Concept B?
Type 2: What will happen when we increase/decrease Concept A?
Type 3: What can cause Concept A to increase/decrease?
Type 4: Tell me what you know about Concept A.

In this paper, we discuss the reasoning and the explanation mechanisms for the first type of question. The mechanisms of the second and the third types are similar (see Leelawong, Wang, Biswas, Vye, Bransford, & Schwartz 2001). The last type of question is descriptive; no reasoning mechanism is involved. The explanation mechanism uses a simple process to convert concept-link-concept pairs to a specific text form.

The reasoning mechanism is based on a simple chaining procedure to deduce the relationship between a set of connected concepts. If the teachable agent is asked to answer the question about the effects of an addition of Algae on Fish, she searches the concept map and deduces that Fish increase when more Algae are added to a river.

As mentioned earlier, Betty employs animation and speech to explain her thinking to the students. A written explanation is also available. The structure of Betty's explanations is closely tied to the reasoning algorithm. To avoid information overload, the explanation is broken down into segments. If users ask for more explanation, Betty works backward, and links the concept back to the closest nodes. When asked "What happens to fish when algae increase?", Betty's initial response is: "I think that when Algae increase, Fish increase." Students can then ask Betty for a more detailed explanation. The following is the sequence of explanation that Betty generates:

1. "You asked me: If Algae increase, what happens to Fish? I think that if Algae increase, Fish increase."
2. "To find out what happens to Fish when Algae increase, I must first know what happens to Oxygen and Plants. Both directly affect Fish."
3. "An increase in Algae causes Plants to increase, which causes Fish to increase."
4. "Now let me explain how Algae affect Fish through Oxygen. An increase in Algae causes Oxygen to remain the same which causes Fish to remain the same."
5. "So, no change in Fish from Oxygen and an increase in Fish from Plants cause an increase in Fish."

After each step in the explanation, students have the option to: (i) ask Betty to repeat the explanation or (ii) look at the explanation in text form.

QUIZ Betty

During the quiz phase, the student observes Betty's responses to a set of pre-scripted questions. The teaching expert informs Betty (and the student) if Betty's answers are right or wrong. The teaching expert also gives hints to help the student debug the concept map.

The teaching expert employs a simple mechanism for generating feedback. The system is provided with an expert concept map (built by a teacher or other expert) in the domain of study. The student's concept map structure is superimposed on the expert's, and the teaching expert searches for a missing concept (first) or relation that is considered essential for the right answer, and uses this to generate a hint for the student. A hint is given, if necessary, for each quiz question. Currently, the system implements three levels of hinting. First, the expert suggests that the student read a section of the resource materials that relates to the concept or link. The second hint for the same question explicitly mentions the name of the missing concept or relation, and asks the student to look for more information on that topic. The third hint tells how to correct the concept or relation in the map.

Procedures

Study participants were 50 high-achieving fifth grade students from a science class in an urban public school located in a southeastern city. Students were randomly assigned to one of four versions of the software: TEACH only, QUERY, QUIZ, or FULL.

The software was used in 3 sessions of one hour each. At the beginning of session 1, students were introduced to features of the software. They were asked to teach Betty about river ecosystems—to model things contained in rivers and how living things in a river meet their survival needs. In between sessions with Betty, students engaged in independent study to prepare themselves to teach Betty. Reference materials were also available for students to access as needed when preparing to teach and when teaching Betty.

Results

Analysis of the scope of students' maps and the types and accuracy of links contained therein suggest several conclusions. Figure 2 shows the mean number of concepts contained in students' maps at the end of each session with Betty's Brain. A 2X2 ANOVA with repeated measures was used to analyze the data. The factors were: Quiz (Quiz or No Quiz), Query (Query or No Query) & Session (Sessions 1, 2 and 3). The Quiz ($F(1,46)=8.47$, $p<.01$), Session ($F(2,92)=77.3$, $p<.001$) and Quiz by Session ($F(2,92)=20.13$, $p<.001$) effects were significant in the

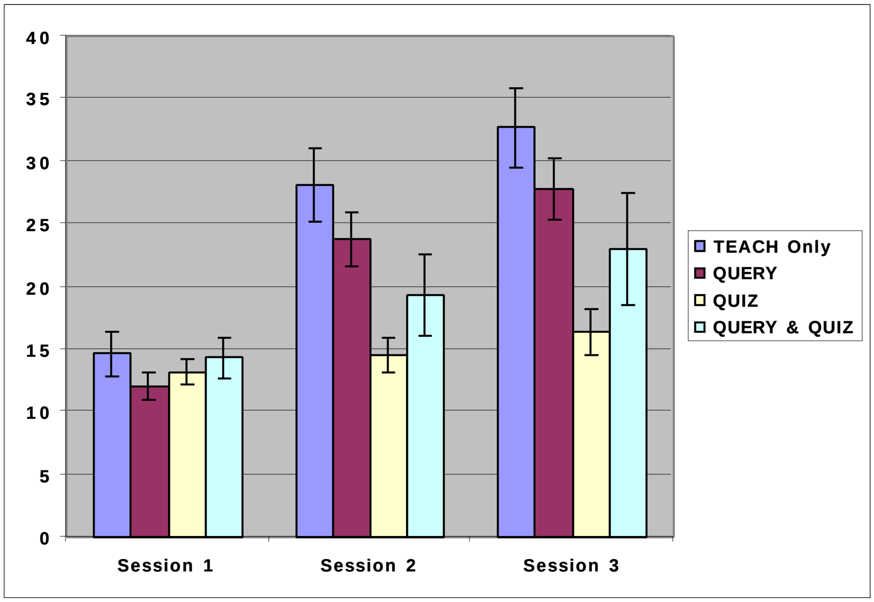

Figure 2. Number of Concepts in Students' Concept Maps

ANOVA. Comparisons of the means (Tukey HSD) indicate that at the end of Session 1 groups did not differ in the number of concepts in their maps, but by the end Session 3, the TEACH only group had significantly more concepts than the FULL and QUIZ groups. The QUERY group had an intermediate number of concepts and did not differ from the other groups. Examination of the maps suggests that maps in the TEACH only group contain many extraneous/irrelevant concepts, whereas the QUIZ group maps contain fewer, but as predicted, more relevant concepts (i.e., more of the concepts contained in the quiz questions). The quiz served to limit the amount of information in students' maps to those concepts that are most important for modeling ecosystem interdependence (e.g., plants, fish, macroinvertebrates, dissolved oxygen, carbon dioxide, waste, bacteria, and sunlight). The intermediate position of the QUERY group may imply that querying, which made students more aware of the causal reasoning process, also helped reduce irrelevant concepts in the concept map.

Figure 3 displays the data on the proportion of causal links in students' maps. The main effect of Quiz and the Quiz by Session interaction were significant ($F(1,46)=4.82$, $p<.05$ and $F(2,92)=18.52$, $p<.001$, respectively). None of the groups differed at the end of Session 1, but by Session 3 students in the QUIZ and FULL version had a significantly higher proportion of causal links in their maps than students in the TEACH only version. In Session 3 the QUERY group did not differ from and was intermediate to the other groups on this measure. These results suggest that quizzing, along with the teacher feedback, focuses students on modeling causal relationships among concepts. It is precisely these relationships that enable students to understand the more global concepts of balance and interdependence related to ecosystems.

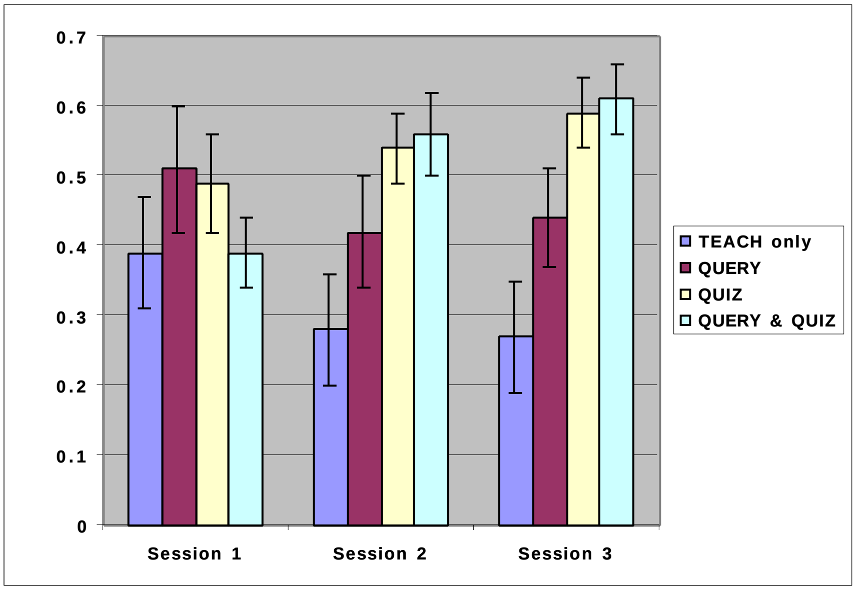

Figure 3. Proportion of Causal Links in Students' Concept Maps

Figure 4 shows the ratio of links to concepts in students' maps, and is a measure of the interconnectedness or density of their maps. The main effects of Query was significant (F(1,46)=4.87, p<.05). Overall, QUERY and FULL students had significantly denser maps in than other students. Evidently, having the opportunity to query Betty, which made the reasoning process more explicit, helped students interrelate concepts in their maps. The effect of Quiz (F(2,92)=20.10, p<.001) and the Quiz by Session interaction (F(2,92)=3.36, p<.05) were also significant. Quiz students' maps became increasingly dense over sessions.

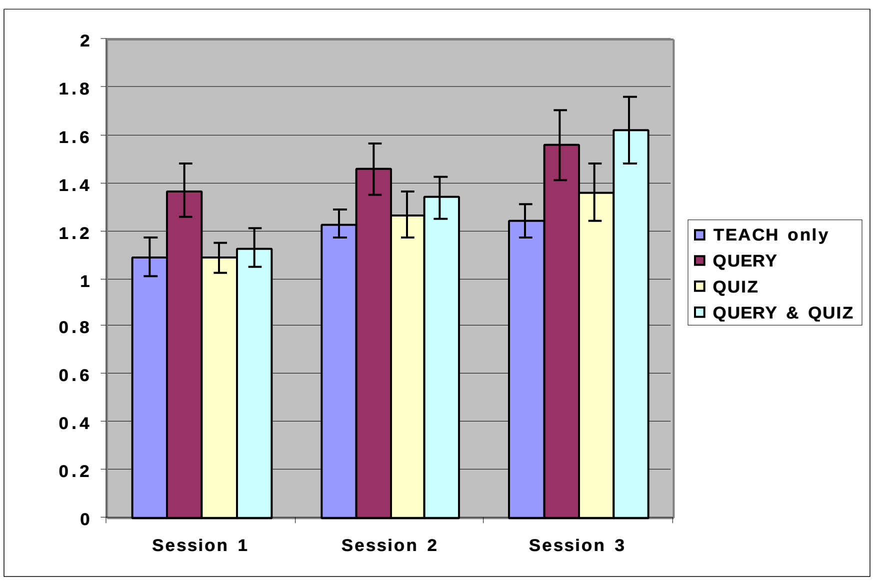

Figure 4. Ratio of Links to Concepts in Students' Concepts Maps

Figure 5 shows the number of valid causal links contained in students' maps. The main effects of Session (F(2, 90)=45.39, p<.001) and Query (F(1,45)=4.92, p<.05) were significant. The Query by Quiz (F(1,46)=4.06, p<.05), Query by Session ($F(2,90)=3.2$, $p<.05$), and Quiz by Session ($F(2,90)=4.69$, $p<.05$) interactions were also significant. Comparisons of the means indicate that by Session 3, QUERY students had significantly more valid links in their maps than students in the TEACH only group. QUIZ and FULL students were intermediate and did not differ from the QUERY and TEACH only groups.

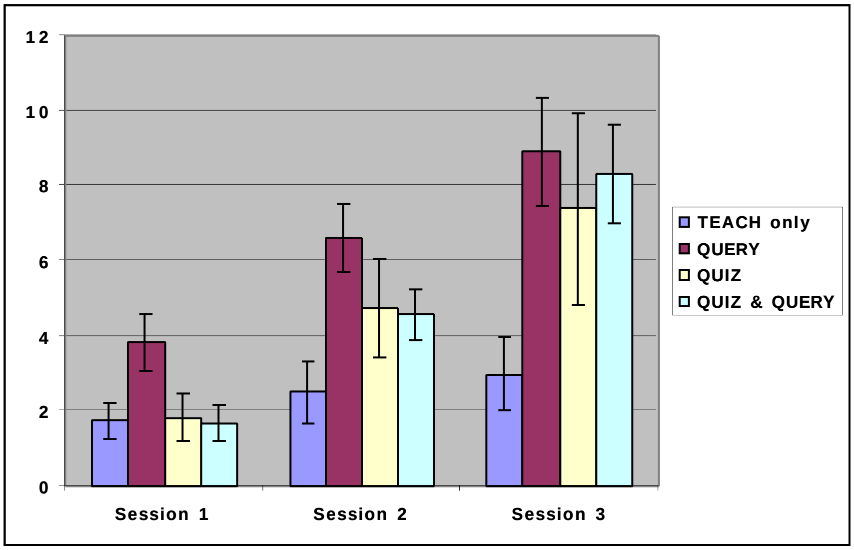

Figure 5. Number of Valid Causal Links in Students' Concepts Maps

When coding the validity of the links in students' maps credit was given for correct links comprising the quiz questions (i.e., links comprising the teaching expert's map), as well as for other relevant links related to water ecosystems (e.g., "Fungi makes CO_2"). Although the QUERY group had more valid links (expert and relevant combined) than the TEACH only group, the QUIZ and FULL groups had more links from the teaching expert's map than students in the TEACH only group (the QUERY group was intermediate). The main effects of Quiz and Session, and the interaction of Quiz by Session were significant ($F(1,45)=15.78$, $p<.001$), $F(2,90)=40.8$, $p<.001$, $F(2,90)=16.99$, $p<.001$, respectively) in the analysis of the number of expert links comprising students' maps. This shows that students in the quiz conditions were guided by the quiz in determining concepts and relations to teach Betty. However, it was not clear how much global understanding the QUIZ only group had of their concept maps.

Discussion

Results from the study indicate that both the query and quiz features had beneficial effects on students' learning about ecosystems. In the query feature, Betty models cause-effect reasoning by animating her map and explaining links in her causal reasoning chain. Students who had access to this feature had a significantly higher ratio of links to concepts, indicating that concepts in their maps were more inter-linked. In this way, the query appears to be effective in helping students develop an understanding of the interrelationships of things—living and non-living—in an ecosystem. [It is interesting to speculate as to whether these students would produce longer causal reasoning chains if queried about the consequences of specific changes to an ecosystem.]

Results indicated that providing students with opportunities to quiz their agent decreased the number of irrelevant concepts, increased the proportion of causal information, and increased the number of expert causal links in students' maps. In these ways, the quiz feature was effective in helping students decide the important domain concepts and types of relationships to teach Betty. Students inferred—and reasonably so--that if a concept or relationship was in the quiz, it was important for Betty to know. This notwithstanding, our observations of students during the study suggest that quiz students may have been overly-focused on "getting the quiz questions correct" rather than "making sure that Betty (and themselves) understood the information." We believe that some of this could be attributed to the nature of the suggestions provided by the teacher agent, which led students to focus on making local changes to their maps, and not paying attention to consequences at the level of the (eco)system.

Surprisingly, students in the QUERY condition produced as many valid relevant causal links as the conditions with the quiz feature, and without the benefit of quiz feedback. This demonstrates the value of explicitly illustrating the reasoning process (by having Betty explain her answers) so that students understand causal structures.

The FULL group did not generate significantly higher-quality maps than the QUIZ and the QUERY groups. An investigation of the activity logs revealed a pattern where students' primary focus was to get the quiz questions correct. After getting Betty to take the quiz, they used the teacher agent's hints to make corrections to their maps, and used the query feature only to check if Betty now would answer the questions correctly. They then quickly returned to the quiz mode to see how well Betty performed. In other words, the query mechanism was not used to reflect on the reasoning mechanisms and to gain a deeper understanding of the causal structures they had created, before corrections were attempted on the concept map. As noted above, the feedback we designed for the teaching agent may have inadvertently focused students on making local changes to their maps instead of reasoning more globally in their maps.

These findings suggest that we need to modify the quiz feature so as to focus students on the interrelationships between concepts and the consequences of these relationships on ecosystem. The quiz feature needs to promote more reflective learning by students. Furthermore, in exit interviews a number of students indicated that while they found the overall environment to be quite interesting and easy to work with, they would like Betty to be more active and participatory in the learning process. Betty is passive and only responds when asked questions. We believe that to create a true learning by teaching environment, Betty needs to be more interactive and demonstrate more human-like qualities.

In summary, our goal is to develop Betty's Brain as a generic teachable agent that can be applied to a variety of scientific domains, where reasoning with cause-effect structures helps in learning about the domain. Results indicate that providing students with opportunities to quiz their agent decreases the amount of irrelevant information and increases the proportion of causal information in students' maps; whereas having opportunities to query their agent helps students develop an understanding of the interrelationships of things—living and non-living—in an ecosystem. The results point to the importance of various forms of feedback when designing teachable agent environments that promote learning.

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 and Feltovich, (Eds.), *Smart Machines in Education* (pp. 71-98). Menlo Park, CA: AAAI Press.

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.

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., Wang, Y., Biswas, G., Vye, N., Bransford, J., & Schwartz, D. (2001). Qualitative reasoning techniques to support learning by teaching: The Teachable Agents project. In *Proceedings of the Fifteenth International Workshop on Qualitative Reasoning* (pp. 109–116). San Antonio, Texas: AAAI Press.

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

Stoyanov, S., & Kommers, P. (1999). Agent-support for problem solving through concept-mapping. *Journal of Interactive Learning Research, 10(3/4*), 401-42.

Acknowledgements

This research was supported by National Science Foundation Grant NSF-9873520. The opinions expressed here are those of the authors and should not be attributed to the funding agency. We would like to thank Babs Rector, Suzanne Cassel and Cynthia Mayfield for their invaluable help on this project.