Intelligent User Interface Design for Teachable Agent Systems

J. Davis, K. Leelawong, K. Belynne, B. Bodenheimer, G. Biswas, N. Vye, and J. Bransford
Dept. of EECS and Learning Sciences Institute
Vanderbilt University, Nashville, TN 37235
Tel: 615-343-6204, Fax: 615-343-6702
biswas@vuse.vanderbilt.edu

ABSTRACT

This paper describes the interface components for a system called Betty's Brain, an intelligent agent we have developed for studying the *learning by teaching* paradigm. Our previous studies have shown that students gain better understanding of domain knowledge when they prepare to teach others versus when they prepare to take an exam. This finding has motivated us to develop computer agents that students teach using concept map representations with a visual interface. Betty is intelligent not because she learns on her own, but because she can apply qualitative-reasoning techniques to answer questions that are directly related to what she has been taught through the concept map. We evaluate the agent's interfaces in terms of how well they support learning activities, using examples of their use by fifth grade students in an extensive study that we performed in a Nashville public school. A critical analysis of the outcome of our studies has led us to propose the next generation interfaces in a multi-agent paradigm that should be more effective in promoting constructivist learning and self-regulation in the learning by teaching framework.

Categories & Subject Descriptors: [I.2.11] Distributed Artificial Intelligence - *Intelligent agents*

General Terms: Design, Experimentation

Keywords: Teachable agents, learning by teaching, feedback, active learning

INTRODUCTION

A critical part of an interactive, intelligent system is an interface that allows knowledge to be added to the system. This paper describes the interface components of a system that enables a user to organize and input problem solving knowledge about a domain for *instructing* an intelligent agent. The agent then uses this knowledge to answer questions and solve problems in this domain. We call such an agent a *teachable agent.* Our motivation for building teachable agents is based on the belief that by the act of teaching, a person more effectively learns and assimilates domain knowledge than through standardized instruction, which may include studying for quizzes and tests. We have deployed teachable agents into K-12 science classrooms, and present, in this paper, an assessment of the effectiveness of our interface.

We begin with the motivation of our work and place it in the context of related efforts. Next, we provide a description of the interface and capabilities of the teachable agent, Betty's Brain. Then we demonstrate by walkthrough examples typical user activities and behaviors in their interaction with the agent, and analyze how these activities impact their learning of the domain material. An assessment of how the system's interface supports the learning activities has led us to propose a new design for the intelligent user interfaces for teachable agents that we believe will be a more effective implementation of the *learning by teaching* paradigm.

The Context for Our Work

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 [1] found that people who prepared to teach others to take a quiz on a passage learned the passage better than those who prepared to take the quiz themselves. The literature on tutoring has also shown that tutors benefit as much from tutoring as their tutees [5, 8]. Biswas and colleagues [3] 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. Beyond the 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 in improving learning and understanding. 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 [5].

IUI'03, January 12–15, 2003, Miami, Florida, USA.

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 have designed an environment that lets students explicitly teach an agent. Once taught, the agent reasons with its knowledge and answers questions. Students observe the effects of their teaching by analyzing these responses. A second agent, the teaching expert, provides additional feedback.

Teachable agents are related to, but different from, the pedagogical agents such as those of Cole [6], Baylor [2], Graesser [7], or Johnson [9]. Pedagogical agents in those systems are primarily demonstrative, as noted by Johnson [9]: *"...agents can demonstrate complex tasks, employ locomotion and gesture to focus students' attention .... and convey emotional responses to the tutorial situation."* Similarly, Cole's Baldi [6] is a demonstrative medium in which Baldi teaches by demonstrating facial motions. A teachable agent, on the other hand, has little or no *a priori* knowledge to demonstrate. The interface between the agent and the user serves two purposes, that of allowing the user to teach the agent, and that of engaging the user's attention through the agent's personality and interactivity. Since our user base is children, this latter purpose is critically important in motivating them to struggle with learning what is often challenging material and communicating it to the agent. So, while these pedagogical agents have informed our work, the different requirements of our application may require that the modalities we use to anthropomorphize our agent be different from those in the previous applications.

Our studies with Betty's Brain have examined the learning benefits of different activities associated with learning by teaching. Betty's Brain 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, accessing content materials as needed to learn information for teaching, (ii) the QUERY mode, where students ask Betty questions (using question templates) that 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. The expert teacher agent makes suggestions that may help Betty (and the student) correct their answers.

Our goal in this work is primarily to examine the effects of the interactive features of the teachable agent interface that emulate the feedback that instructors receive from students during teaching. Our hypothesis was that having opportunities to query and/or quiz Betty would positively, but differentially, impact students' learning. The query feature would help students debug their own thinking and reasoning in the problem domain. We expected the quiz feature would help students identify important concepts and links to include in their concept maps. We also expected that overall they would produce more accurate concept maps because they had access to feedback on Betty's quiz performance.

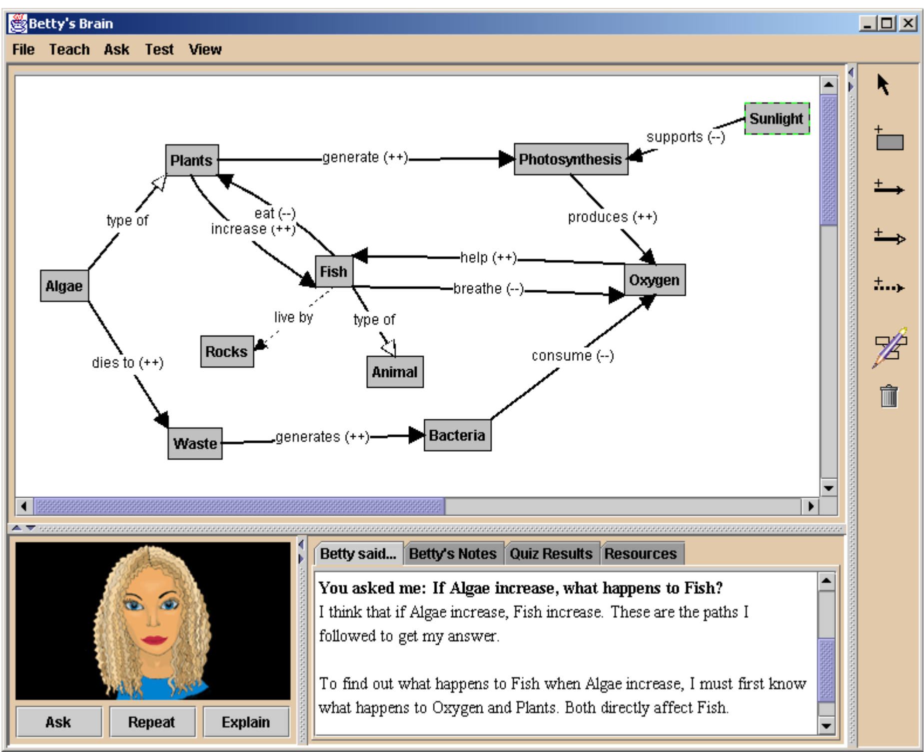

Figure 1. Betty's Brain Interface

DESCRIPTION OF BETTY'S BRAIN

Figure 1 illustrates the interface of our teachable agent, Betty's Brain. Students use a graphical drag and drop interface to create and modify their concept maps in the top pane of the window. Students can query Betty using the *Ask* button, and she provides an explanation for how she derives her answers by depicting the derivation process using multiple modalities: text, animation, and speech. The visual display of the face with animation in the lower left is one way in which the user interface attempts to provide engagement for the user. This modality of face and voice is designed to increase the social interaction of the user with the system. As Reeves and Nass [13] have noted, the inclusion of these features can indicate to the user the existence of a social presence in the system, hopefully engaging and motivating them.

The system is implemented using a generic agent architecture that is illustrated in Figure 2. The primary component of the agent is its decision maker that incorporates the qualitative reasoning mechanism for generating answers to queries from the concept map structure, and decision-making schemes that implement strategies that govern the dialog process with the user. The executive controls the dialog mechanisms, and Betty's speech and animation engines. These are primarily used to explain how Betty derives her answer to a question. In the sections below, we describe the software's three 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 [12]. A relation is a unidirectional link connecting two entities. Concept maps help to categorize groups of objects and express interactions among them as cause-effect relations. They also provide a mechanism for representing knowledge hierarchies [15], which makes the technique very amenable to representation of phenomena in scientific domains, in particular, for modeling dynamic systems. In this study, we asked students to teach Betty about the living and non-living things in a river and how living things survive by interacting with one another. The focus is on creating a river ecosystem model that incorporates the general principles of balance and interdependence.

Figure 1 displays an example of a concept map that a student created in Betty's Brain—the map represents what the student has taught Betty. This map is not a complete representation of all the knowledge in the domain, but merely an example. 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 to 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 two question types:

Type 1: *What will happen to Concept A when we increase/decrease Concept B?*

Type 2: *Tell me what you know about Concept A.*

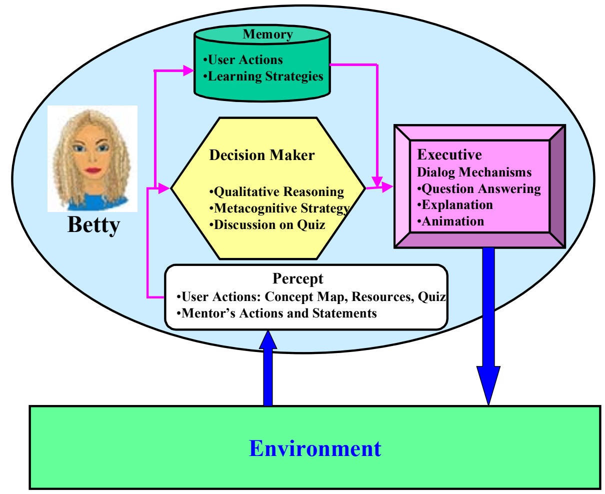

Figure 2: Betty's Agent Architecture

We briefly explain the reasoning and the explanation mechanisms for the first type of question. More details can be found in Leelawong, et al. [10, 11]. The second 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 on Fish when of Algae are added, 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. Betty's response then takes the form, *"To find out what happens to Fish when Algae increase, I must first know what happens to Oxygen and Plants. Both directly affect Fish."*

"An increase in Algae causes Plants to increase, which causes Fish to increase."

Through further interaction, Betty reveals the complete explanation.

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, implemented as a second agent in the environment, also gives hints to help the student debug the concept map. This agent employs a simple mechanism for generating feedback. The system is provided with an expert concept map (built by the classroom teacher or other expert) in the domain of study. The student's concept map structure is overlaid 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. The first hint points the student to resource materials, both on-line and text-based, that relate to the concept or link. As the second hint for the same question, the expert agent explicitly mentions the name of the missing concept or relation. The third hint is very direct. It names a missing concept or tells students how to correct a causal relation in their current map.

Research Findings

During the Spring semester of 2002, we ran an extensive study on a class of high-achieving fifth grade students (eleven and twelve year olds) from a science class in an urban public school located in Nashville [11]. Their goal was to model a river ecosystem using concept maps. This study was linked to material they were studying as part of

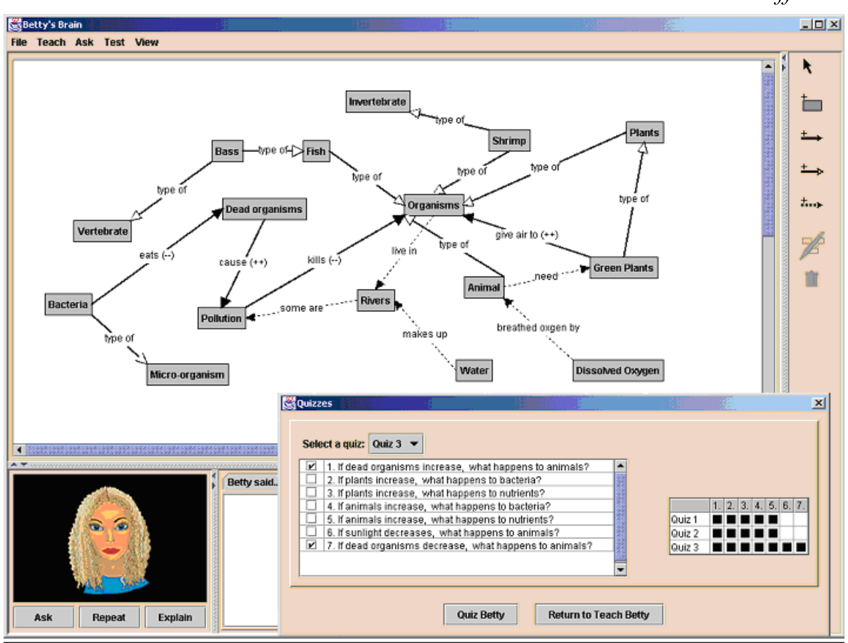

Figure 3. The initial concept map and the first quiz

their regular curriculum. Twelve of these students worked with the full version of Betty's Brain. In this paper, our focus is on how our agents' interactions and interfaces influenced the students' learning activities. The software was used in three sessions, each session lasting about one hour. 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 how living things in a river meet their survival needs. In between sessions, 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 from this study indicate that both the query and quiz features had beneficial effects on students' learning about ecosystems. The *query* feature helped students understand the role of causal relations in defining the interdependencies among entities. This feature thus resulted in more interlinked concept maps. The results also indicated that providing students with opportunities to make their agent take *quizzes* generated by the teacher or a domain expert decreased the number of irrelevant concepts, increased the proportion of causal information, and increased the number of expert causal links in students' maps. Overall, the quiz feature was effective in helping students determine relevant domain concepts and relations to teach Betty so she could improve her performance. Students reasonably inferred that if a concept or relationship was in the quiz, it was important for Betty to know. This inference notwithstanding, our observations of students during the study suggest that students were more focused on "*getting the quiz questions correct*" rather than" *making sure that Betty (and they themselves) understood the interdependence relations involved, and their effects on the overall system*." We believe this behavior occurred in part because the type of suggestions provided by the teacher agent led students to focus on making local changes to their maps, and not consider consequences at the level of the (eco)system. We were disappointed that the overall quality of the student's maps was not very high. We analyze the reasons for these findings by doing a detailed walkthrough of some of the student activities below.

AN EXAMPLE WALKTHROUGH

Once students had created their initial concept maps, their focus shifted to getting individual quiz questions correct, rather than to understand and model the global consequences of interdependence and balance. The resultant behavior can be summarized by an iterative cycle of actions that corresponded to a *"test, modify, and retest"* behavior. A study of the log files indicated that after students had asked Betty to take the quiz, they first used the results and the teacher's feedback to make local changes to the concept map, then they used the query mechanism to see if the answer to the particular question they were focused on was being answered correctly, and then they tried the quiz again. They did not seem to reflect on why an answer was wrong, and what consequences would a change in one part of the map have on other parts of the map. Thus, students often found themselves in situations where changes they made to correct one answer resulted in a number of other answers becoming incorrect. This pattern of behavior clearly does not support the learning of interdependence and balance, nor does it demonstrate a system-level understanding of the domain.

On a few occasions, some students demonstrated a more constructive learning strategy, which involved more reflection on why a question had not been answered correctly. Students consulted the resources to gain more knowledge of the entities and relations involved, and used debugging methods that used the query mechanism to get a more detailed understanding of Betty's reasoning mechanisms. We demonstrate through examples these two different kinds of behavior.

Figure 3 shows the initial concept map a student created before he got Betty to take her first quiz. The quiz interface, which is accessed by a pull-down menu, is also shown in the figure. In the current version of the system, the student has a choice of three quizzes. Each quiz has between five and seven questions of varying difficulty, and the student has the choice of picking one or more questions for Betty to answer. In our example, the student selected from Quiz 3 two similar questions that had the same source and target concepts but different trends (increase/decrease) for the source concept. Once Betty answered the questions, the students saw how she had performed on the quiz. Betty's performance on quiz questions that she had taken earlier is also shown on this interface to the right of quiz pane in Figure 3. A green check implies a correct answer, and a red cross implies that Betty's answer to that question was incorrect. Thus the student receives comprehensive feedback of how well Betty has performed thus far, which is a direct indication of the accuracy and completeness of his concept map. The student also saw the teacher agent's feedback to Betty. This is shown in Figure 4. The teacher often commended Betty for getting her answers right, or she provided hints to Betty on how she could improve on her incorrect answers. The hints are in the three-level form that was described earlier.

In our particular example, the student decided to go back to the teach mode and make corrections to the concept map. Ideally, the student would have used the query mechanism and Betty's explanations to reflect on where knowledge in the concept map was incorrect or incomplete, and then consult the resources to determine what changes needed to be made. We make two observations here. First, the student chose to ignore an explicit suggestion from the teacher that he should look up resources about how *animals* are affected by *dead organisms*. On the other hand, the student misinterpreted the feedback, and went on to create a direct causal link between animals and dead organisms. This type of observation suggests that the teacher's feedback should be made clearer.

In accordance with the "test, modify, and retest" behavior pattern, the student immediately went back to the quiz mode, and asked Betty to answer the same questions as in the previous cycle. This time Betty did answer the question, but her answers were incorrect. At this point, the teaching expert came back with more explicit feedback that indicated to the student that the relation between *dead organisms* and *animals* was not direct, but it was mediated through another concept, such as *Nutrients* (see Figure 5). This behavior pattern, indicative of shallow learning (or no learning), occurred for a number of cycles.

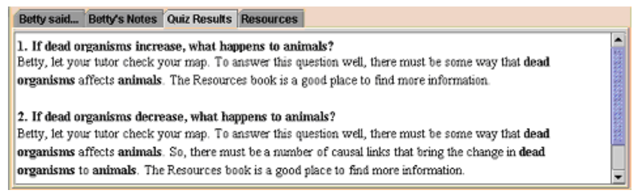

Figure 4. The teacher agent's feedback on the first quiz

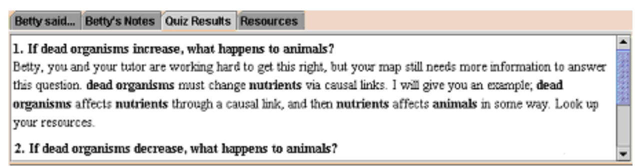

Figure 5. A more specific hint from the teacher agent

At some point during session 2, perhaps through discussion with one of the experimenters, this student realized that a better way of improving Betty's performance would be to study Betty's explanations for the answers she generated using the query mechanism. This change in behavior is shown in Figures 6-10. At this point, the feedback from the teacher agent was like the one for the previous case. The teacher agent suggested that Betty's (i.e., the student's) concept map needed to include causal links between *Animals* and *Dissolved Oxygen,* (Figure 7). This time the student first queried Betty on one of the quiz questions (see Figure 8), then reflected on Betty's explanations (shown in Figure 9), and then used the online resources to study the section on Animals (shown in Figure 10). In spite of these efforts, the student failed to generate the correct concept map in this session. Also, when we looked through the student's logs in a later session, he had reverted back to his "test, modify, and retest" behavior.

We believe that one of the primary shortcomings of this system is that it does not promote the use of strategies for effective learning. Part of this may be attributed to weaknesses in the teacher agent feedback, which seems to promote making local changes rather than think globally about system-level effects. Also, too much reliance on the quiz mechanisms, caused the students to focus solely on performance, i.e., how do I get these quiz questions right, rather than reason about the global structure of the concept map, and its implications as a model of ecosystem behavior.

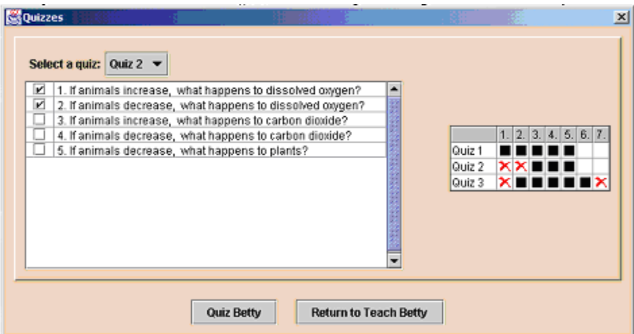

Figure 6. The second set of selected quiz questions

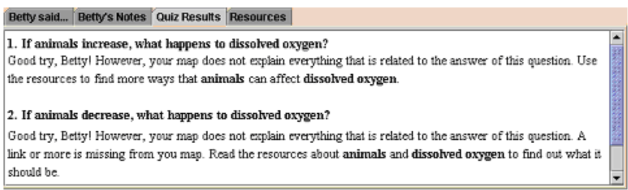

Figure 7. The feedback from the teacher agent

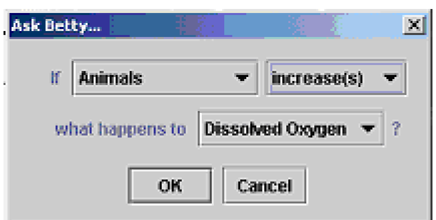

Figure 8. The Causal Query Dialog

Figure 9. The Betty's Notes panel

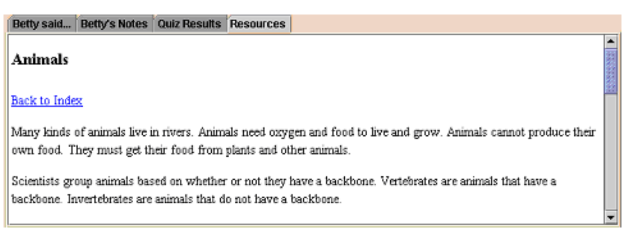

Figure 10. The Resources panel

To confirm the observed behavior patterns described above, we analyzed the log files of all twelve students. The results are summarized in Table 1. On average, students used the query feature 0.53 times between two quizzes. This implies that they did not exploit Betty's explanations to study the structure of their causal map, and reflect on where problems may have occurred. A number of students did use the online and text resources to find information. But the explanation mechanism would have helped them understand if they had translated this information into the right causal structure, and how one causal link may affect the relation between different sets of concepts.

The number of modifications made to a concept map between two quiz events had an average value of 2, a clear indication that students were focused on local changes, i.e., changing one link at a time to generate the right answer for a particular quiz question. The result was that most students did not end up with very accurate concept maps for the domain. Very often, when they modified their concept maps to correct an answer to a quiz question, other questions that Betty had answered correctly in the past, became incorrect. A number of students were very frustrated by this, but only a few began to realize that they needed to look at their concept map as a global structure, as opposed to focusing only on individual links. We observed that students who used the query mechanism more developed this realization, but most of the others concentrated more on the quiz. It is interesting that students tended to ask about two queries each time Betty took the quiz, but very often they were looking for the effects of opposite trends for the same pair of concepts (see the walkthrough example above). We would have liked the students to focus on multiple queries because that would have enabled them to realize the global effects of local dependencies.

Session	Number of times Quiz Taken	Number of Questions attempted per Quiz	Number of Queries made between two quizzes	Number of modifications to concept map between two Quizzes
1	2.9	1.4	0.4	3.0
2	8.8	2.3	0.8	3.5
3	13.6	2.2	0.4	2.1

Table 1: Analysis of the debugging activities

After the third session with the system, we conducted exit interviews, which indicated that the students liked to work with Betty in the environment that we had created. After the first session, most students had no trouble with the system's interfaces. At times, they were puzzled by the fact that Betty could not answer a query or quiz question. On a number of occasions, this was because of missing concepts and links, or incorrect links. There were some situations where the students could not figure out that they had modeled a causal link backwards, i.e., they had modeled B is affected by A (B→A) instead of A affects B (A→B). In these situations, the experimenters had to intervene, and help the student.

A very important piece of feedback that we received from a number of students was that they would have liked Betty to participate more actively during the teaching phase, i.e., they wanted Betty to exhibit characteristics of a good student and be a more *active* learner. Our studies also indicate that the feedback from the system (both from the teacher agent and Betty) must be improved to facilitate better learning among students during the teaching phase. All these issues have prompted us to take an in depth look into the design of a new generation of teachable agents, which we describe in the next section.

In designing a new multi-agent architecture for a teachable agents environment, we have adopted ideas from a National Research Council study on "How People Learn"[4], which states that an effective learning environment must consider four issues for effective learning environments: (i) be *learner-centered*, i.e., focus on relating subject matter to students' prior experiences, understanding, and preferred style of learning, (ii) be *knowledge-centered*, that is include the knowledge and skills necessary to gain problem solving expertise, (iii) be *assessment-centered*, which emphasizes that learners receive feedback both during (formative assessment) and after (summative assessment) the teaching process to help them stay on track in terms of meeting their learning goals, and (iv) be *community-centered*, which recognizes that learning can occur outside of classroom environments and encourages learning by collaboration.

A NEW TEACHABLE AGENT INTERFACE DESIGN

Our analyses of student activities and the accuracy of the final concept maps indicate that the students had become proficient in using the concept map structure and the query and the quiz mechanisms. However, it was not clear as to how much their understanding of domain knowledge had improved.

A previous study [11] indicated that Betty's Brain helped college students become more aware of the interdependence between sets of concepts. Students working with pencil and paper tended to create single causal links, whereas students working with Betty's Brain created longer chains of causal structures. However, the study reported in this paper demonstrates that younger students have difficulty comprehending the global consequences of interdependence among entities, and how that affects balance in the river ecosystem.

We indicated that number of students would like Betty to be more active and participate in the learning process. On average, we found that students using the teachable agent with voice spent 2.5 times more time with the agent than students using the agent without voice. This result is consistent with that of previous researchers, such as Reeves and Nass [13], but it allows us to hypothesize that the presence of animation may further increase the time spent with the agent. Additionally, several students commented that the agent needed a better and more realistic voice. These comments indicate that the users tend to demand high quality modalities in their agent. Supporting this, several students also suggested that the teachable agent, Betty, should be more interactive, e.g., *"react to what she was being taught, and take the initiative and ask more questions on her own,"* and *"do some sort of game or something and make it more interactive."*

Consistent with this idea, we note that the current version of Betty is passive and only responds when asked questions. We believe that to create a true learning by teaching environment, Betty needs to better demonstrate qualities of human students. A tutor gains deeper understanding from interactions with a tutee [5] that includes answering questions, explaining materials, and discovering misconceptions. Betty should be designed to benefit her users in the same fashion.

From the knowledge-centered standpoint, we need to modify the environment's features to make students more cognizant of the global interdependence issues rather than local interdependence, particularly with regard to balance. For example, a river ecosystem can sustain its balance because of the interactions among the systems of carbon dioxide and oxygen cycle, food chain, and decomposition. Unless students understand how interdependence produces balance, it is difficult to understand the effects of external phenomena that may affect the system. In addition, students should have a greater understanding of the idea that local changes affect global behavior.

Consistent with this need, the user resources are being updated to emphasize the processes and cycles that describe domain phenomena, as opposed to the individual entities that make up the domain. The on-line resources will be reformatted in hypertext format to enable keyword access, and the overall structure of the document will explicitly reflect the phenomena of interdependence and balance through the primary cycles, such as the food chain and the oxygen cycle. The role of the different entities in the river ecosystem will be developed in these contexts. Therefore, entities, such as plants and bacteria will appear in different roles in different cycles and processes. This realization should make the student realize and reason about the global implications of interdependence in the concept map.

In addition, other aspects of effective learning environments must be incorporated in conjunction with the knowledge-centered aspects. As seen in the walkthrough example, the students struggle to learn while teaching. They attempt to use all features—the query mechanisms, the quiz mechanism, the teacher agent's feedback and paper and the on-line resources. However, our environment was not sufficiently designed to cope with the learner-centered aspect when the user is both a domain and teaching novice.

To improve the learner-centered aspect, our environment should more actively assist students in learning and mastering domain knowledge. The teacher agent can guide and help students to adopt constructive learning techniques, and once they do, encourage them to continue on this path. Betty can also encourage users by asking if they would like to discuss the results of the quiz with her, allowing her to demonstrate learning strategies through her dialog and actions.

The Time for Telling strategy [14] will also be integrated into Betty's Brain. Betty can construct questions that involve long chains and multiple-path events in the current concept map, and ask the user to analyze the answers she generates. Discussion of these answers will lead to better understanding of the concept map structure and Betty's reasoning mechanisms. The teacher agent can suggest reading materials that are related to the discussion.

We would also like to associate with Betty behaviors that are linked to good learning practice. Betty can be endowed with self-regulation strategies that drive her interactions with the user. For example, to counter the student's sole focus on performance in the quiz, Betty may refuse to take the quiz if the student repeatedly ignores the teacher agent's feedback. Along the same lines, Betty can express skepticism if the student does not look up resources before attempting to make corrections in the concept map.

The assessment-centered aspect of our design framework can be improved if the interface allowed a different quiz structure that emphasizes global system behavior instead of local changes. For example, the feedback could better emphasize the global nature of changes at the expense of a local change made to get one particular quiz question correct. The quiz feedback could also promote reflective learning to a higher degree if Betty incorporated reflective activities, potentially encouraged by the teacher agent. Our interactive Betty can also enhance the community-centered aspect. By collaborating on the quiz results and discussing the causal effects locally and globally in the concept map, Betty shares a common space in learning with the user. Thus, different instantiations of Betty could be contrasted among users. Collaborations between students teaching Betty are also being developed at Stanford University.

This set of design changes is being incorporated into a new multi-agent architecture. This architecture includes enhanced versions of the Betty and teacher agent that have a computational structure similar to what was shown in Figure 2. A third agent, called the Tracker, follows activities and changes that are made in the concept map environment, and uses an automata-based mechanism to derive useful patterns from the low level actions and changes. For example, the tracker is trained to catch the local test-modify-retest behavior and report this to Betty and the teacher agent. This agent also keeps track of time between actions, so that Betty can react to long periods of inactivity by the user. The tracker agent avoids repetition, and makes it easier to implement the percept mechanisms and the dialog strategies of the Betty and teacher agent. In summary, we believe that the next version of Betty's Brain will result in a generic teachable agent architecture that can be applied in a variety of scientific domains, where reasoning with cause and effect structures helps in learning about the domain.

REFERENCES

1. Bargh, J.A. and Schul, Y. On the cognitive benefits of teaching. *Journal of Educational Psychology*, *72* (5). 593-604.
2. Baylor, A. and Ryu, J. Effects of image and animation on agent persona. *Journal of Educational Computing Research*.
3. Biswas, G., Schwartz, D., Bransford, J. and The Teachable Agents Group at Vanderbilt University. Technology Support for Complex Problem Solving: From SAD Environments to AI. in Forbus and Feltovich eds. *Smart Machines in Education*, AAAI Press, Menlo Park, CA, 2001, 71-98.
4. Chi, M.T.H., Siler, S.A., Jeong, H., Yamauchi, T. and Hausmann, R.G. Learning from Human Tutoring. *Cognitive Science*, *25* (4). 471-533.
5. Cole, R., Massaro, D.W., Rundle, B., Shobaki, K., Wouters, J., Cohen, M., Beskow, J., Stone, P., Connors, P., Tarachow, A. and Solcher, D., New tools for interactive speech and language training: Using animated conversational agents in the classrooms of profoundly deaf children. in *Proceedings of ESCA/SOCRATES Workshop on Method and Tool Innovations for Speech Science Education*, (April 1999).
6. Graesser, A., Wiemer-Hastings, K., Wiemer-Hastings, P. and Kreuz, R. Autotutor: A simulation of a human tutor. *Journal of Cognitive Systems Research*, *1*. 35-51.
7. Graesser, A.C., Person, N. and Magliano, J. Collaborative dialog patterns in naturalistic one-on-one tutoring. *Applied Cognitive Psychologist*, *9*. 359-387.
8. Johnson, W.L., Rickel, J.W. and Lester, J.C. Animated pedagogical agents: Face-to-face interaction in interactive learning environmen. *International Journal of Artificial Intelligence in Education*, *11*. 47-78.
9. Leelawong, K., Davis, J., Vye, N., Biswas, G., Schwartz, D., Belynne, K., Klatzberger, T. and Bransford, J., The Effects of Feedback in Supporting Learning by Teaching in a Teachable Agent Environment. in *The Fifth International Conference of the Learning Sciences*, (Seattle, Washington, 2002).
10. Leelawong, K., Wang, Y., Biswas, G., Vye, N. and Bransford, J., Qualitative reasoning techniques to support learning by teaching: The Teachable Agents project. in *The Fifteenth International Workshop on Qualitative Reasoning*, (San Antonio, Texas, 2001), AAAI Press, 109-116.
11. National Research Council *How People Learn*. National Academy Press, Washington, D.C., 2000.
12. Novak, J.D. Concept Mapping as a tool for improving science teaching and learning. in Treagust, D.F., Duit, R. and Fraser, B.J. eds. *Improving Teaching and Learning in Science and Mathematics*, Teachers College Press, London, 1996, 32-43.
13. Reeves, B. and Nass, C. *The Media Equation: How people treat computers, televisons and new media like real people and places*. Cambridge University Press, Cambridge, 1996.
14. Schwartz, D.L. and Bransford, J.D. A Time for Telling. *Cognition and Instruction*, *16* (4). 475-522.
15. Stoyanov, S. and Kommers, P. Agent-Support for Problem Solving Through Concept-Mapping. *Journal of Interactive Learning Research*, *10* (3/4). 401-442.

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 thank Babs Rector, Suzanne Cassel and Cynthia Mayfield for their invaluable help on this project. We acknowledge, Dan Schwartz at Stanford University, who has been a collaborator on the Teachable Agents project. We also acknowledge feedback from Thomas Katzlberger, Janet Truitt, and Karun Viswanath, who are the other members of the Teachable Agents project.