

Teachable Agents

Learning by Teaching Environments for Science Domains

Krittaya Leelawong, Karun Viswanath, Joan Davis, Gautam Biswas, Nancy Vye, Kadira Belynnne, and John Bransford

Department of EECS and Learning Sciences Institute (LSI)

Vanderbilt University

Nashville TN 37235.

Tel: (615)-343-6204; e-mail: biswas@vuse.vanderbilt.edu

Abstract

The crisis in science education and the need for innovative computer-based learning environments has prompted us to develop a multi-agent system, Betty's Brain that implements the learning by teaching paradigm. The design and implementation of the system based on cognitive science and education research in constructivist, inquiry-based learning, involves an intelligent software agent, Betty, 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. The results of an extensive study in a fifth grade classroom of a Nashville public school has demonstrated impressive results in terms of improved motivation and learning gains. Reflection on the results has prompted us to develop a new version of this system that focuses on formative assessment and the teaching of self-regulated strategies to improve students' learning, and promote better understanding and transfer.

1 Introduction

A research review by Ponder and Kelly [1] determined that the science education crisis in U. S. schools has been present for over four decades. Billions of dollars have been invested in research and reforms to resolve this crisis but several problems remain. Science curricula still need to work on increasing student literacy, encourage conceptual understanding, motivate students, and develop concrete problem solving skills [1, 2]. Clearly, the need for high level of science literacy is critical given the technology push of contemporary society and the complexity of modern life.

Unfortunately, current school curricula tend to emphasize memorization, which provides students with limited opportunities and little motivation to develop "usable knowledge". The ability to apply knowledge to solve real-world tasks cannot be equated to remembering a mere list of disconnected facts. Studies of expertise have shown that knowledge needs to be connected and organized around important concepts, and should

support transfer to other contexts. Other studies have established that improved learning and understanding happens when the students take control of their own learning, and develop metacognitive strategies to assess what they know, and acquire more knowledge, when they need to. Thus the learning process must help students build new knowledge from existing knowledge (*constructivist learning*), guide them to discover learning opportunities while problem solving (*exploratory learning*), and help them to define learning goals and monitor their progress in achieving them (*metacognitive strategies*).

A recent National Research Council study has outlined factors that define the effectiveness of learning environments [2]. Effective environments must be: (i) *learner-centered*, i.e., focus on relating subject matter to students' prior experiences, understanding, and preferred style of learning, (ii) *knowledge-centered*, i.e., include the knowledge and skills necessary to gain problem solving expertise, (iii) *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) *community-centered*, which recognizes that learning can occur outside of classroom environments and encourages learning by collaboration.

Advances in computer technology have led to the proliferation of efforts in building computer-based learning environments. Most of the efforts to date, may be grouped into two primary categories: (i) Intelligent Tutoring Systems (ITS) [3] and (ii) Cognitive Tools [4]. Though they have achieved a fair degree of success, each, on its own, fails to meet all the necessary requirements of good learning environments. ITS systems are constructivist because they focus on learning by problem solving, but their primary focus is mostly on modeling the student, and tailoring feedback to address the student's immediate needs. As a result students do not have opportunities to study the global structure of a domain, and to learn by exploration. Cognitive tools on the other hand, have focused more on open-ended exploratory activities, but they have often failed to help students conditionalize their knowledge to problem solving tasks. Their lack of focused feedback has often resulted in students getting

stuck at *plateaus of performance*. In past work, we have developed Intelligent Learning Environments (e.g., AdventurePlayer [5]) that combine exploratory and constructivist learning with coaching feedback [6] to help students progress, when they reach plateaus of performance.

To advance computer-based learning systems, we began to investigate additional features, such as motivation to improve student's learning. The cognitive science and education research literature supports the idea that teaching others is a powerful way to learn. Bargh and Schul [7] 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 shown that tutors benefit as much from tutoring as their tutees [8, 9]. Biswas and colleagues [10] 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. A second directed study with an older version of Betty's Brain showed that students who taught Betty had a deeper understanding of the domain, and could express their ideas better than those who studied the same material and wrote a summary [24].

Reflection on these studies and others led us to design an environment that lets students explicitly teach a software agent, called Betty's Brain. Once taught, the agent reasons with its knowledge and answers questions. Students observe the effects of their teaching by analyzing these responses. Other agents work in tandem to assist in the student's own discovery learning. We should clarify that our agent does not conform to the traditional intelligent agent architecture, i.e., it does not possess machine learning algorithms that enable it to learn from examples, explanations, or by induction.

It is to be noted that our Teachable Agents (TA) environment have important characteristics that are often not available in tutoring environments. For example, students must accept the responsibility of teaching in a way that improves their tutees performance, deal with questions asked by their tutee, and be able to assess their own understanding of material by looking at how their tutee solves problems and answers questions. These kinds of activities should facilitate an approach to lifelong learning that is valuable and ultimately measurable.

2 Background

Repenning and Sumner [27] have developed visual programming environments that reduce the overhead of learning to program agents. Smith et al. [11] Cocoa program (previously KidSim) allows young users to program their agents by example. Once created, they become alive in the environment and act according to its pre-programmed behavior. Other work, such as the Persona project [12] has focused on sophisticated user interactions, communication and social skills.

Research on helping agents to learn, has focused on agents that can learn from examples, advice, and explanations ([13], [14]). In Huffman and Laird's system [13], agents learn tasks through tutorial instructions in natural language. Users have some domain knowledge which they refine by looking at the agents behaviors. Lieberman and Maulsby [15] focus on teaching "instructible agents" by example and by providing advice. Agents learn by observing user actions, sometimes by being told what is relevant, and sometimes by identifying relevant information, applying it, and learning through mistake correction.

Michie et al. [16] developed the Math Concept Learning System for solving linear equations. Users supplied the strategies for solving problems by entering examples, and the system learned via an inductive machine-learning algorithm, ID3 [17]. Obayashi et al.'s study [18] reported significant learning gains in subjects using their learning-by-teaching system compared to a traditional Computer-assisted Instructor (CAI). Chan & Chou's [19] study concluded that learning by teaching is better than studying alone.

These studies basically employed an approach where a traditional ITS or CAI framework was employed, but the roles of the human and the computer program were reversed. In ITS, students are the recipients of knowledge and the locus of control is mainly with the student modeling component and the pedagogical agent. On the other hand, in the learning by teaching environments, students taught, and probably learnt from the experience. Learning was more indirect, and not likely to show immediate gains. The studies conducted were not sufficiently rigorous or long term to demonstrate that this approach might provide deeper understanding and transfer in more complex domains. Therefore, we decided to take a new approach to designing a learning by teaching environment that would support constructivist learning, and also one that would include feedback to promote better self-regulation and develop understanding. Unlike previous studies on learning by teaching, we took on the challenge of teaching students who were novices in the domain, and also novice teachers.

3 Betty's Brain: A Learning-by-Teaching Environment

As discussed, we have built an environment where students explicitly teach and directly receive feedback

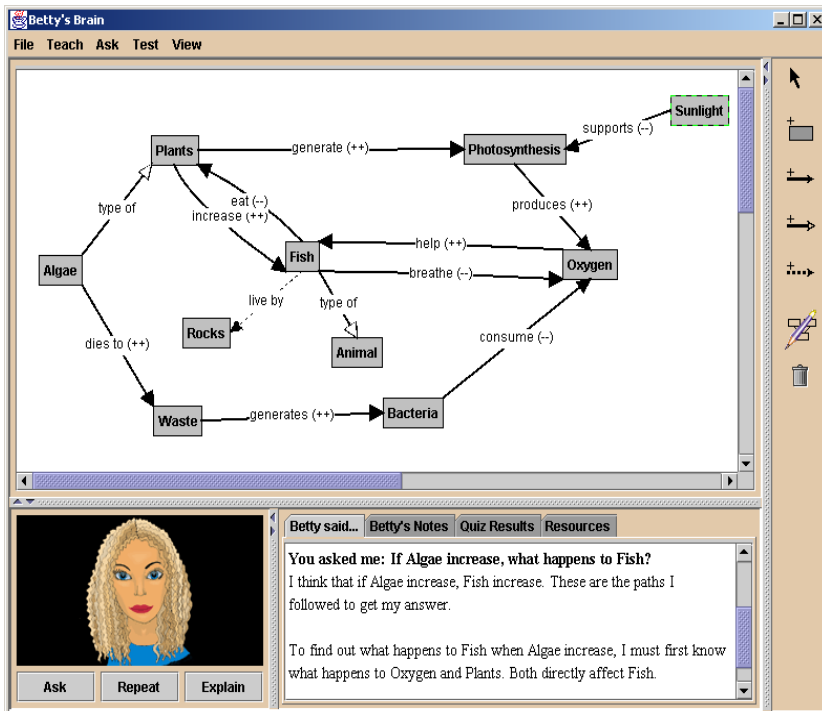


Figure 1. Betty's Brain Interface

about their teaching through interactions with a computer agent, named Betty's Brain. The system has been used to teach middle school students about interdependence and balance among entities that exist in a river ecosystem.

A learning by teaching system requires a representation scheme for students to create their knowledge structures as a part of the teaching process. Since the primary users are middle-school students solving complex problems, this representation has to be intuitive but sufficiently expressive to help these students create, organize, and analyze their problem solving ideas. A widely accepted technique for constructing knowledge is the concept map [20]. Concept maps provide a mechanism for structuring and organizing knowledge into hierarchies, and allow the analysis of phenomena in the form of cause-effect relations [21, 22]. This makes them amenable to modeling scientific domains, in particular dynamic systems. Moreover, an intelligent software agent based on concept maps can employ reasoning and explanation mechanisms that students can easily relate to. Thus the concept map provides an excellent representation that serves as the interface between the student and the teachable agent.

Fig 1 illustrates the interface of 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, and motivate students by increasing the social interaction with the system [23].

The system is implemented as a generic agent architecture illustrated in Fig 2. The primary component of the agent is its decision maker that incorporates the qualitative reasoning mechanisms for generating answers to queries from the concept map structure, and 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 of Operation: TEACH, QUERY and QUIZ. The system also has a second agent, the mentor agent, who provides feedback to Betty after she takes a quiz.

3.1 TEACH Betty

Students teach Betty by means of a concept map interface. Fig 1 displays an example of a concept map that 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.

3.2 QUERY Betty

Students can 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. The system

provides templates using which students can query Betty, e.g., *What will happen to Concept A when we increase/decrease Concept B?*

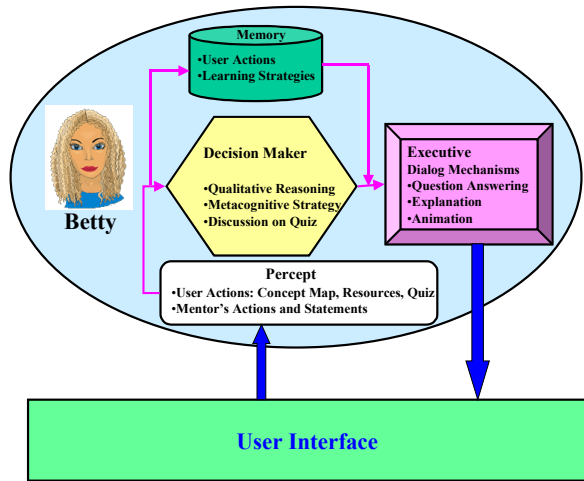


Figure 2: Betty's Agent Architecture

We briefly explain the reasoning and the explanation mechanisms for this type of question. The reasoning mechanism is based on a simple chaining procedure to deduce the relationship between a set of connected concepts. To derive the effect of a change (either an *increase* or a *decrease*) in concept *A* on Concept *B*, Betty performs the following steps:

- Starting from Concept *A*, propagate the effect of its change through all outgoing casual links by pairwise propagation (i.e., follow the link from Entity *A* to all its effects) using the relations described in Table 1.
- At each concept, if the number of incoming casual links is more than one, the forward propagation stops until all incoming links are resolved. To derive the result from two incoming links, use Table 2. A “?” in Table 2 implies an inconclusive change (attributed to the ambiguity of qualitative arithmetic). If the number of incoming links is three or more:
- Count the number of changes that fall into the six categories: large, normal, and small decreases and small, normal, and large increases. Combine the corresponding (i.e., small, medium, and large) changes; always subtract the smaller number from the larger. For example, if there is one small decrease and two small increase incoming arcs, the result is a small increase. To compute the overall effect, if the resultant value set has all increases or all decreases, select the largest change. Otherwise, start at the smallest level of change and combine with the next higher level in succession using Table 2.

To illustrate the reasoning process, assume that the teachable agent is asked to answer the effects of an addition of Algae on Fish using the partial ecosystem concept map from Fig 1.

Betty starts with the initiating concept and computes the result on the end entity by sequential propagation along individual paths. For example,

Algae (+) is a type of (+) Plants (+) generate (++) Photosynthesis (+) produces (++) Oxygen (+): (*Not a final result – there are more incoming links into oxygen*)

Algae (+) die to (++) Waste (+) generates (++) Bacteria (+) consumes (-) Oxygen (-): (*All “Algae -> Oxygen” effects are now traversed.*)

An increase (step 1) and a decrease (step 2) changes in Oxygen result in inconclusive change (“?”) in Oxygen: *Continue propagation.*

Oxygen (?) helps (+) Fish (?): *Wait*

Algae (+) is a type of (+) Plants (+) increase (+) Fish (+): (*All “Algae -> Fish” paths now traversed.*)

An *increase* (step 5) and an *inconclusive* (step 3) change for Fish combine to produce an overall *increase* in Fish.

The answer to this question is that Fish increase when we add more Algae. More details can be found in [24]. The overall qualitative reasoning mechanism is not novel but a simplified implementation of Qualitative Process Theory [25]. However, the focus here is on observing the effects of teaching and feedback.

		Change in Relation					
		+ _L	+	+ _S	- _S	-	- _L
Change in Entity	+ _L	+ _L	+ _L	+	-	- _L	- _L
	+	+ _L	+	+ _S	- _S	-	- _L
	+ _S	+	+ _S	+ _S	- _S	- _S	-
	- _S	-	- _S	- _S	+ _S	+ _S	+
	-	- _L	-	- _S	+ _S	+	+ _L
	- _L	- _L	- _L	-	+	+ _L	+ _L

Table 1: The pair-wise effects

		+ _L	+	+ _S	- _S	-	- _L
		+ _L	+ _L	+ _L	+ _L	+	+ _S
+	+ _L	+ _L	+	+ _S	?	- _S	
+ _S	+ _L	+	+	?	- _S	-	
- _S	+	+ _S	?	-	-	- _L	
-	+ _S	?	- _S	-	- _L	- _L	
- _L	?	- _S	-	- _L	- _L	- _L	

Table 2: Integrating results from two paths

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. Using the example, “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.

3.3 QUIZ Betty

During the quiz phase, the student observes Betty’s responses to a set of pre-scripted questions. The mentor agent informs Betty (and the student) if Betty’s answers are right or wrong. The mentor also gives hints to help the student debug the concept map. This agent employs a simple mechanism for generating feedback using an expert concept map (built by the classroom teacher) in the domain of study. The student’s concept map structure is overlaid on the expert’s, and the mentor agent 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.

4 Experimental Results

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 city. 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: (i) TEACH only version (No QUERY or QUIZ), (ii). QUERY version, (iii). QUIZ version and (iv). 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 impor-

tant, when Betty explains her answers, she makes explicit the process of reasoning across links along multiple paths in a concept map. Therefore, we might expect that students who use the QUERY version 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.

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. 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.

The data collected was analyzed using two-way Anova. Analysis of the scope of students’ maps and the types and accuracy of links contained therein suggest several conclusions. It was clear that the students who used the query and 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 teach only group.

Fig 3 shows the ratio of links to concepts in students’ maps, a measure of the interconnectedness or density of their maps. Overall, QUERY and FULL students had significantly denser maps than other students. Evidently, having the opportunity to query Betty, which made the reasoning process more explicit, helped students understand the importance of interrelations among concepts in their maps. Another observation was that the Quiz students’ maps became increasingly dense over sessions.

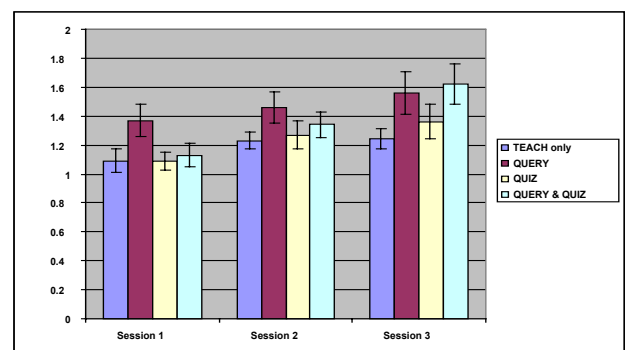


Figure 3. Ratio of Links to Concepts in Students’ Concepts Maps

Fig 4 shows the number of valid causal links contained in students’ maps. 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 much from each other.

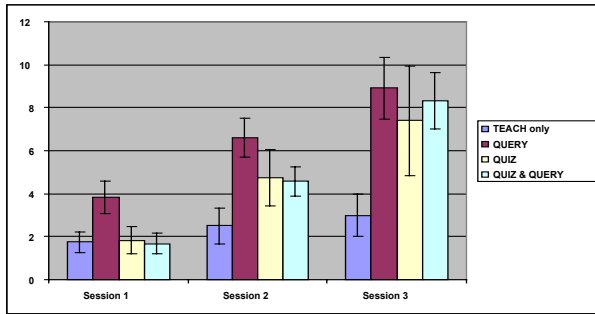


Figure 4. Number of Valid Causal Links in Students' Concepts Maps

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 (determined by our expert raters). Although the QUERY group had the most valid links (expert and relevant combined), the QUIZ and FULL groups had more links from the teaching expert's map than students in the QUERY group. The data indicates that students in the quiz conditions were guided by the quiz and the teacher agent feedback in determining concepts and relations to teach Betty. However, it was not clear how much of a global understanding the QUIZ only group had of their overall concept maps.

4.1 Discussion

Results from the study indicate that both the Query and Quiz features had beneficial effects on students' learning about ecosystems. Students who had access to the Query feature had the most inter-linked maps. The Query mechanism appears to be effective in helping students develop an understanding of the interrelationships between entities in an ecosystem. Also, the opportunities to quiz their agent helped students to decrease the number of irrelevant concepts, increased the proportion of causal information, and increased the number of expert causal links in their maps. 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 they themselves) understood the information.” We believe that some of this could be attributed to the nature of the suggestions provided by the mentor 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 would now 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. As noted above, the feedback we designed for the mentor agent may have inadvertently focused students on making local changes to their maps instead of reasoning more globally in their maps.

5 Agent Architecture

As reported in Davis et al. [26], in exit interviews, the students emphasized that they would have liked Betty to have been an active participant in the teaching phase, i.e., the students wanted Betty to exhibit characteristics of a good student and be a more active learner. Studies also indicate that feedback from the system (both from the Mentor and Betty) must be improved to facilitate better learning. These issues prompted us to design a new Multi-Agent Teachable Agents system.

The Multi-Agent architecture enabled us to overcome drawbacks of the old system as well and

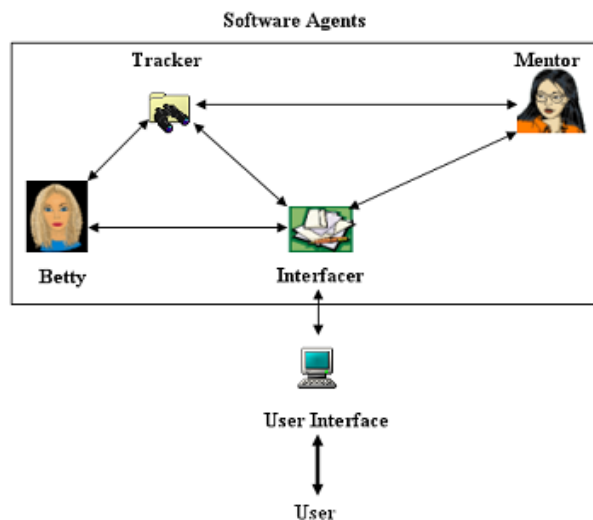


Figure 5: Agent Architecture

introduced new features to promote inquiry based self regulated learning. From a software design viewpoint each agent handles data relevant to its own functionality. Agents subscribe to events generated by the other agents. This allows them to intervene with teaching opportunities and assist the students' own discovery learning. Besides learning opportunities, these agents also test students understanding, and gently guide them to improve their knowledge of the domain [9].

5.1 Betty

To create an effective learning by teaching environment, Betty needs to demonstrate qualities of a good student. Tutors gain deeper understanding from interactions with tutees [6, 7] that includes answering questions, explaining material, and discovering misconceptions. Betty's new communicative "personality" is intended to support such interactions. This personality is governed by a set of self-regulation strategies to promote better learning. For example, after taking a quiz, Betty can encourage users by asking if they would like to discuss results of the quiz with her, allowing her to demonstrate learning strategies through her dialog and actions. She can also exhibit behaviors linked to good learning practice. Betty may refuse to take the quiz if the student repeatedly ignores the mentor agent's feedback. She can also express reservations if the student does not look up the resources before attempting to make corrections in the concept map. Betty, interacting with the tracking agent, (transparent to the user), can also react to long periods of inactivity by the user, and suggest when they may get help.

5.2 Mentor Agent

The mentor is the domain knowledge expert. Currently this knowledge is encoded as an "expert" concept map (hidden from Betty and the student) and structured online resource documents that everyone has access to. The mentor is available for help on demand. The mentor agent usually directs students to the online resources. This document has been updated to emphasize the processes and cycles that describe domain phenomena, as opposed to individual entities that make up the domain. The resources have been reformatted in hypertext format to enable keyword access. The overall structure of the document explicitly reflects the phenomena of balance and interdependence through the primary cycles, such as the food chain and the oxygen cycle.

The mentor can also intervene when the student does not seem to be making progress. She combines the use of a set of metacognitive and pedagogical strategies to make decisions on when to help students. Sometimes she points students to the resources, at other times she may suggest that the student can get better feedback by querying Betty.

5.3 Pattern-Tracker

User actions should be linked to learning, teaching and self-assessment objectives. The pattern tracker uses an automata to determine whether an action or a series of actions match a predefined pattern. Other agents can subscribe to receive notifications on the occurrence of specific patterns. The pattern tracker can act independently for detecting simple patterns. For complex patterns, it needs to interact with the other agents. As an example, when the tracker detects that the user has asked Betty to take a Quiz immediately after making changes in the concept map, it interacts with the other agents to check if this pattern represents a local test-modify-retest behavior. Upon detecting such a behavior, it publishes this as an event.

5.4 Interfacer

The interface agent acts as a middleman between the Graphical User Interface (GUI) and the underlying structures and reasoning mechanism. The user interacts with the interface and the interfacer displays the updated view in the GUI. When the other agents (Betty and Mentor) need to interact with the user, they compose a formatted message, which the interfacer displays to the user. Thus the GUI components can be changed with no effect on the underlying mechanism.

6 Conclusions

Our studies with Betty's Brain demonstrate its effectiveness in promoting deep understanding and self-assessment among students. We have also shown that students require very little instruction in using the various components of the system. Feedback from the fifth grade science class and their teachers indicate that this environment was successful in motivating students and getting them to spend a lot more time in learning about complex domains. More extensive studies are now being conducted with a focus on studying the effects of feedback and self regulated learning in the Multi-agent architecture environment.

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, our collaborator on this project from Stanford University. We also acknowledge feedback from Thomas Katzlberger and Janet Truitt, who are the other members of the Teachable Agents project.

References

- [1] Ponder, G.K., Kelly, J., Evolution, chaos, or perpetual motion? A prospective trend analysis of secondary science curriculum advocacy, 1955-1994. Jour-

- nal of Curriculum and Supervision, 1997. 12(3): pp. 238-245.
- [2] Bransford, J.D., A.L. Brown, and R.R. Cocking, eds. *How People Learn*. expanded ed. 2000, National Academy Press: Washington, D.C.
- [3] Wenger, E., *Artificial Intelligence and Tutoring Systems*. 1987, Los Altos, California: Morgan Kaufmann Publishers.
- [4] Lajoie, S.P. and S.J. Derry, *Computer Environments as Cognitive Tools for Enhancing Learning*, in *Computer as Cognitive Tools*, S.P. Lajoie and S.J. Derry, Editors. 1993, Lawrence Erlbaum Associates: Hillsdale, NJ.
- [5] Crews, T.R., et al., *Adventureplayer: A microworld anchored in a macrocontext*, Intl. Jour. of AI in Educ., vol. 8, pp. 142-178, 1997.
- [6] Burton, R.B., Brown, J. S., *An investigation of computer coaching for informal learning activities*. Intelligent tutoring systems, London: Academic Press, 1982: p. 79-98.
- [7] Bargh, J.A. and Y. Schul, *On the cognitive benefits of teaching*. *Journal of Educational Psychology*, 1980. 72(5): p. 593-604.
- [8] Graesser, A.C., N. Person, and J. Magliano, *Collaborative dialog patterns in naturalistic one-on-one tutoring*. *Applied Cognitive Psychologist*, 1995. 9: p. 359-387.
- [9] Chi, M.T.H., et al., *Learning from Human Tutoring*. *Cognitive Science*, 2001. 25(4): p. 471-533.
- [10] Biswas, G., et al., *Technology Support for Complex Problem Solving: From SAD Environments to AI*, in *Smart Machines in Education*, Forbus and Feltovich, Editors. 2001, AAAI Press: Menlo Park, CA. p. 71-98.
- [11] Smith, D.C., Cypher, A., & Spohrer, J, *Programming Agents without a Programming Language*, in J. M Bradshaw (Ed.) *Software Agents*. 1997, Menlo Park, CA: AAAI/MIT Press. p. 165-190.
- [12] Ball, G., Ling, D., Kurlander, D., Miller, J., Pugh, D., Skelly, T., Stankosky, A., Theil, D., Van Dantzich, M., & Wax, T., *Lifelike computer characters: The persona project at Microsoft research*, in J. M. Bradshaw (Ed.), *Software Agents*. 1997, Menlo Park, CA: AAAI/MIT Press. p. 191-222.
- [13] Huffman, S.B. and J.E. Laird, *Flexibly Instructable Agents*. *Journal of Artificial Intelligence Research*, 1995. 3: p. 271-324.
- [14] Paltheu, S., J.E. Greer, and G.I. McCalla. *Learning by teaching*. in *The International Conference on Learning Sciences*. 1991. Illinois, USA.
- [15] Lieberman, H., and Mulsby, D., *Instructible agents: Software that just keeps getting better*. *IBM Systems Journal*, 35, 3/4., 1996.
- [16] Michie, D., A. Paterson, and J. Hayes-Michie. *Learning by Teaching*. in *Second Scandinavian Conference on Artificial Intelligence (SCAI) 1989*. 1989. Tampere, Finland: IOS Press.
- [17] Quinlan, J.R., *Induction of decision trees*. *Machine Learning*, 1986. 1(1): p. 81-106.
- [18] Obayashi, F., H. Shimoda, and H. Yoshikawa. *Construction and Evaluation of CAI System Based on 'Learning by Teaching to Virtual Student'*. in *World Multiconference on Systemics, Cybernetics and Informatics*. 2000. Orlando, Florida.
- [19] Chan, T.-W. and C.-Y. Chou, *Exploring the Design of Computer Supports for Reciprocal Tutoring*. *International Journal of Artificial Intelligence in Education*, 1997. 8: p. 1-29.
- [20] Novak, J.D., *Concept Mapping as a tool for improving science teaching and learning*, in *Improving Teaching and Learning in Science and Mathematics*, D.F. Treagust, R. Duit, and B.J. Fraser, Editors. 1996, Teachers College Press: London. p. 32-43.
- [21] Kinchin, I.M. and D.B. Hay, *How a qualitative approach to concept map analysis can be used to aid learning by illustrating patterns of conceptual development*. *Educational Research*, 2000. 42(1): p. 43-57.
- [22] Stoyanov, S. and P. Kommers, *Agent-Support for Problem Solving Through Concept-Mapping*. *Journal of Interactive Learning Research*, 1999. 10(3/4): p. 401-42.
- [23] Reeves, B. and Nass C., *The Media Equation: How people treat computers, televisions and new media like real people and places*. 1996: Cambridge University Press, Cambridge.
- [24] Leelawong, K., et al. *The Effects of Feedback in Supporting Learning by Teaching in a Teachable Agent Environment*. in *The Fifth International Conference of the Learning Sciences*. 2002. Seattle, Washington.
- [25] Forbus, K., *Qualitative Process Theory*, *Artificial Intelligence*, vol. 24, pp. 85-168, 1984.
- [26] Davis, J.M., et al. *Intelligent User Interface Design for Teachable Agent Systems*. in *International Conference on Intelligent User Interfaces*. 2003. Miami, Florida: The Association for Computing Machinery.
- [27] Repenning, A., & Sumner, T. (1995). *Agentsheets: A medium for creating domain oriented visual languages*. *Computer*, 28, 17-25.