

Bringing CBLEs into Classrooms: Experiences with the Betty’s Brain System

John Wagster¹, Henry Kwong², James Segedy¹, Gautam Biswas¹, and Dan Schwartz²

¹Vanderbilt University & ²Stanford University

Nashville, TN 37235. USA.

john.wagster@vanderbilt.edu

Abstract

This paper discusses the Betty’s Brain system and our ongoing work on developing a suite of tools that assist students and teachers in classroom learning in science domains. We describe the design and implementation of the system using a client/server architecture, and the initial responses of the teachers to the tools we have developed. Future enhancements and additions to the suite of tools are discussed.

1. Introduction

The Teachable Agents group at Vanderbilt University (www.teachableagents.org) and the AAA Lab at Stanford University (aaalab.stanford.edu) have developed exploratory learning environments that use the learning-by-teaching paradigm to help middle school students learn science content with understanding and the ability to transfer to other learning situations [1][2]. Research has shown that teaching is a powerful way to learn [3][4][5][6]. Preparing to teach is a self-directed and open-ended activity where one explores, integrates, and structures knowledge first for oneself, and then for others.

To demonstrate the positive effects of learning by teaching, we have developed a CBLE (computer-based learning environment) called Betty’s Brain, where students learn by teaching a computer agent named Betty using structured visual representations called concept maps. Our system is geared toward middle school science classrooms, where the students are not domain experts and they have little experience in teaching practices. Therefore, we have built a variety of scaffolds to help students learn as they teach Betty, and guide their interactions with her in a way that emphasizes the cognitive and metacognitive aspects of the teaching and learning processes. For example, the visual concept map structure is a purported representation of Betty’s knowledge, and students teach her by building and revising this repre-

sentation [2]. This creates a “shared representation,” where the student’s own knowledge and understanding of the domain is projected through the visual structure that Betty uses for reasoning and answering questions in the domain. Students can monitor, assess, and reflect on their own learning, by asking Betty questions and observing her performance on quizzes that are provided with the system. The belief that they are helping Betty learn motivates the students’ and guides them toward metacognitive activities that aid their learning.

In the last five years, we have conducted a number of studies in 5th grade science classrooms that measure student learning by the quality of the concept maps the students produce while teaching Betty. We found that learning-by-teaching with metacognitive support helped students learn about river ecosystems, and better prepared them for future learning on related topics [1][2]. We have compared the learning by teaching environment with a non-teaching version. Students who taught Betty developed more complete and interconnected concept maps than students who created maps for themselves (i.e., these students made concept maps and received feedback from the system on the quality of the map, but there was no cover story of teaching an agent). Learning outcomes were strongest for students who also received metacognitive feedback from Betty. These differences persisted during a transfer phase in which students learned about a new domain and taught Betty in the absence of most feedback and prompts.

More recently, we have studied students’ activity patterns as they teach Betty and create their concept maps. Preliminary analysis shows that the quality of students’ concept maps is paralleled by patterns in their behaviors [1]. It seems that self-regulated learning prompts from Betty helped the student teachers engage in productive learning interactions.

Our past studies were related to the classroom science curriculum, but they were conducted in a pull out format. 10-12 students worked on the system at one time, and multiple sessions were run every day to

cover all students from a set of classrooms. The data collected allowed us to demonstrate the positive learning gains of our learning by teaching systems. This increased the teachers' interest in our system and they wanted to be more involved in instructional design and subsequent studies. The natural next step was to expand the scope of the system and our studies in a way that they would become an integral part of the science curriculum and classroom instruction. In this paper, we discuss how the system was extended to include six classrooms of 25-30 students each. The teachers took the lead in the instruction and learning activities related to Betty's Brain.

The next section discusses the system architecture with the added components to make it amenable for classroom instruction. The original Betty's Brain system was expanded to include two important components. The first is a system called "Front-of-the-Class" (FOC) Betty, which allows teachers to project a number of student generated maps in front of the class and ask questions of all of the maps. The second is a set of teacher tools that provides automated methods for grading the student maps, compute statistics on student performance, and observe students' interactive patterns as they learn and teach Betty. Section 3 discusses teachers' responses to some of these tools. Section 4 presents conclusions and future work.

2. The Betty's Brain System

2.1. Client/Server Architecture

For use in classroom environments, the system needed a number of additional components to assist teachers in their instructional tasks and evaluation of student learning. The teacher tools combined with the student system was implemented as a client/server architecture shown in Fig. 1.

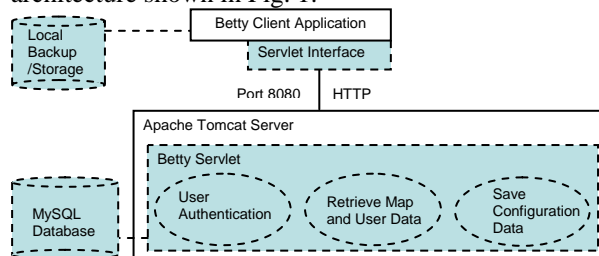


Figure 1. Client/Server Architecture

This architecture allows the system to be used by multiple users simultaneously, and at the same time it provides facilities for collecting data on user performance that can be used by both the teachers and

the researchers. Our original attempts to build a custom server was quickly shelved when we realized that we could utilize Java Web Start to pull programs and data from a Tomcat server to the individual clients and also upload user data back to the server when required. This allowed us to easily deploy the application in schools, while sidestepping some security issues related to network access in setting up the server.

All of the student learning and teacher tools described below are implemented as clients and they can be accessed remotely through a standard internet browser. User authentication back to the servlet on the Tomcat server is handled by a MySQL backend. This allows easy data exchange between the student and teacher tools. Data on the client application is typically cached for speed, and to create a backup to avoid situations when a connection is lost during data transmission. The use of standard HTTP protocols together with a public domain HTTP communications package permit us to retry dropped connections repeatedly. The result is a cheap but reliable solution in potentially unreliable classroom environments.

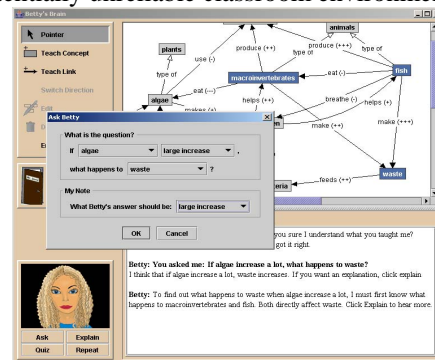


Figure 2. Betty's Brain with Query Window

2.2. Betty's Brain: The Student system

The students' interface to the Betty's Brain system is illustrated in Fig. 2. The teaching process is implemented as three primary activities: (i) *teach*: students explicitly teach Betty by creating a concept map representation, (ii) *query*: students use a template to generate questions to see how much Betty has understood, and (iii) *quiz*: students observe Betty's performance on a set of predefined questions. Once taught, Betty uses qualitative reasoning methods to reason through chains of links [2] to answer questions, and, if asked, explain her reasoning using text and animation schemes. Betty also provides feedback that reflects the students' teaching behaviors. The goal is to get the students to adopt meta-cognitive strategies in their learning tasks [2]. Stu-

dents reflect on Betty's answers and her explanations, and revise their own knowledge as they make changes to the concept maps to teach Betty better.

2.2.1 Metacognitive Support in Betty's Brain. We have adopted a self-regulated learning (SRL) framework that includes a set of comprehensive skills that start with setting goals for learning new materials and applying them to problem solving tasks, deliberating about strategies to enable this learning, monitoring one's learning progress, and then revising one's knowledge, beliefs, and strategies as new materials and strategies are learnt [7]. In conjunction with these higher level cognitive activities, social interactions and motivation also play an important role in the self-regulation process.

Betty's persona in the SRL version incorporates metacognitive knowledge that she conveys to the students at appropriate times to help them develop and apply monitoring and self regulation strategies [1][8]. For example, when the student is building the concept map, Betty occasionally responds by demonstrating reasoning through chains of events. She may query the user, and sometimes remark (right or wrong) that the answer she is deriving does not seem to make sense. The idea of these spontaneous prompts is to get the student to reflect on what they are teaching, and perhaps, like a good teacher check on their tutee's learning progress. These interactions are directed to help Betty's student-teacher understand the importance of monitoring and being aware of one's own abilities.

2.3. Betty's Brain Teacher/Researcher Tools




Figure 3. FOC: Enlarged map for discussion

We have added tools to help teachers and researchers monitor student performance and guide classroom interactions and instructions. We have also added editing facilities for making quick modifications to some of the content within the system.

The first tool called *Front-of-the-Class (FOC)* Betty seen in Fig. 3 is designed to help teachers discuss student-generated concept maps as part of a class activity. This tool allows teachers to bring up a number of student maps on a screen, ask questions of the maps and compare answers. The teacher can also show the paths of reasoning used by a concept map to generate answers, and get their students to discuss differences in the maps and the answers. This provides the students with additional opportunities to develop monitoring skills. Teachers can also use tools to show classroom performance statistics so students get a feel for how well they are performing in comparison with the rest of the class.

Our initial version of the tool had logistic problems attributed to the granularity and limited viewable screen size of the projection technology. We solved this problem by allowing users to create multiple views for displaying the maps, from a bird's eye view with many maps (but very little detail observable) to a single fully enlarged map (so students at the back of the class could see the teacher trace the individual steps of the reasoning mechanism).

The second tool is *Betty's Brain Bank* seen in Fig. 4. It provides teachers and researchers administration, grading, and monitoring functions that are accessible from a common interface.

Administrative tasks include the ability to add users, setup classroom rosters, and modify content in the resources, expert map, and quiz questions. Researchers use their interface to tailor content based on teacher requests without having to modify configuration or data files on the server. There are tools for extracting performance (concept map score, quiz scores) and behavior (students' activity sequences) data for analysis through a web interface. This allows us to present data to the teachers quickly.




Figure 4. Betty's Brain Bank, grading interface

Teachers can use three tools to grade student maps by: (i) comparing answers to all possible causal ques-

tions (APQ) asked of the student and expert maps, (ii) computing scores for a specific set of questions generated by the teacher (these could be the quiz questions from the system), and (iii) checking for specific concepts and links in their maps. Teachers can export these assessments into a tool such as Excel to create plots or to compute weighted scores for students and classroom averages. Teachers can also view individual student maps to check student's work in more detail.

The behavior analysis tools provide teachers with information on students' interactive patterns as they learn for themselves and teach Betty [1][8]. Our tools provide counts for single actions (e.g., how many queries did a student or a set of students ask), predefined action sequences (e.g., how many times was a question followed by a request for an explanation). These patterns are then visualized using JFreeChart. An example can be seen in Figure 5. In recent work, we have used hidden Markov models to extract student behavior models from their activity sequences [8]. More details and teachers reactions to these tools are described in section 3.

2.4. Lessons Learned

Lessons learned fall into three categories: deployment, experimental scaling to include more subjects, and developing tool interfaces that made it easier for teachers to run the system, observe data, and run FOC discussions.

To handle this we setup a client/server architecture with a low-cost custom server and a wireless network in each classroom. Two residential-grade wireless routers played the role of bridge and repeater. This setup allowed us to regulate the bandwidth and access to the Betty's Brain suite of tools. Teachers could look at and grade maps whenever they chose, and use this information to develop lesson plans for the next day. We also ran a backup and test server at Vanderbilt, which allowed us to make changes and quickly deploy the software from a browser to any location with internet access. This allowed teachers and researchers to easily work together when running the study.

Although the system was developed in Java, a platform-independent language, the realities in the classroom were somewhat different. In the school environment, we were limited to using Java 1.5. The school did not have enough laptops for a full class; so we used the school Macs and some of our PC laptops to run the studies. The teachers had access to both Macs and PCs, but their web and wireless access were restricted by school policies. Java Web Start

provided solutions to some of our problems since we could avoid software installation on the school computers with files written only to the user directory of the client. Another, tough technical challenge was porting the text-to-speech functions of the Betty system between the two operating systems. We used a combination of the FreeTTS engine and MBROLA voices, which after exhaustive search was the only TTS API that was free, cross-platform, and somewhat maintained. File Handling, compression, image creation, and a consistent GUI all became unexpected hurdles as well. Overall, Java seems to provide an excellent framework for doing cross-platform work, but with the caveat that the details of our tools had to be tailored to the individual operating system to maintain consistent experimental conditions.

3. Teachers' Response and Discussion

As a part of integrating within the classrooms we have been working specifically with three 5th grade teachers who are participating in our Betty's Brain studies. One of the teachers had this to say about the FOC system: *"I have found that students that struggle in science pay greater attention when their concept maps appear in the Front of the Class system... I think this is an invaluable tool to show students and let them see what others have done. It helps students who struggle [to] see that they are not alone, and they learn from each other."*

Previous CBLE systems and our teachers have used performance data to evaluate student learning, however; we saw the behavior data we collected provided the teachers with a novel window to observe students' learning processes. So, we performed a quick evaluation study with the teachers to determine what student behaviors they considered to be useful, and the form in which they would like these behaviors displayed [9]. We used a combination of Likert scale questions (see example in Table 1 below), open ended questions, and interviews to assess teachers' confidence levels in interpreting the behavior data and its usefulness in understanding students' learning. This preliminary analysis provided us with some insights to the needs of the teachers. Teachers saw four different visualizations: a pie chart, a line chart, a histogram, and a stacked area plot of randomly

Table 1. Example Likert scale questions as part of the evaluations of the behavioral interfaces

Rate your overall confidence in accessing the behavioral analysis and using the behavior interfaces:
How confident are you in using the behavioral interfaces to augment the performance analysis:




Figure 5: Histogram interface with summary analysis at top, behavioral visualization at center

selected behavioral data for a student, and were then asked to analyze student performance using the different visualizations. The visualizations were presented to them in random order.

Teachers' initial concerns about being able to understand the meaning of the behaviors were ascertained from interviews and a pretest, which contained seven point Likert scale and free response questions that asked to record their confidence in using and understanding the content they could extract from the teacher tools. They were then shown the visualizations, which were followed by free response questions that asked them to record their likes and dislikes for each interface. The teachers were then given the posttest with the same questions as in the pretest. Teachers' initial confidence in the ease of use was above average with a mean of 6 ± 1 . After the task teachers confidence decreased a little to 5.5 ± 1 . The difference was not significant so we concluded that the visualizations and access methods had met the teachers' expectations. However, the small decrease in confidence could be attributed to the teachers concerns about the role of the behavior data in their overall assessment of the students. They felt that the new approach would be difficult to justify to administrators and parents.

Setting aside these concerns, the teachers showed interest in behavior patterns that they considered were related to the students looking up and reading the online resources on the system. This was not surprising because these behaviors indicated the students were attempting to learn new material on their own. The teachers also gave credit to students who combined resource access, asking queries, checking explanations, and quizzing Betty in a balanced way as they used the system. These conclusions are in agreement with our analyses of good student performance [1].

4. Conclusions and Future Work

Our deployment of the Betty's Brain system in the classroom with the FOC tools has been successful, and we will analyze the results of our study in more detail this summer. We will also study how to better involve teachers in CBLE deployment so they may be better integrated with classroom instruction. We will continue to refine the Teacher tools to allow the teachers to customize the feedback and prompts provided by the system. We are working on a set of lesson planning tools that allow for the customization of rubrics, and the computer agent feedback. Our goal is to build tools that support the full cycle of curriculum development from the initial planning to feedback about the progress of students and feedback related to the next curricular plan. For administrators we plan to extend this work to allow for quick response to teachers needs during a study.

5. References

- [1]Wagster, J., J. Tan, G. Biswas, and D. Schwartz, "How Metacognitive Feedback Affects Behavior in Learning and Transfer", The 13th International Conference on AI in Education, Marina del Rey, California, 2007.
- [2]Biswas, G., D. Schwartz, K. Leelawong, and N. Vye, "Learning by Teaching: A New Agent Paradigm for Educational Software", Applied Artificial Intelligence, 19(3), 2005, pp. 363-392.
- [3]Bargh, J.A. and Y. Schul, "On the cognitive benefits of teaching", Journal of Educational Psychology, 72(5), 1980, pp. 593-604.
- [4]Roscoe, R.D., and M.T.H. Chi, "Understanding Tutor Learning: Knowledge-Building and Knowledge-Telling in Peer Tutors' Explanations and Questions", Review of Educational Research, 77(4), 2007, pp. 534-574.
- [5]Chi, M.T.H., S. A. Siler, et al., "Learning from Human Tutoring," Cognitive Science, 25(4), 2001, pp. 471-533.
- [6]Biswas, G., D. Schwartz, J. Bransford, and the Teachable Agents Group at Vanderbilt, "Technology Support for Complex Problem Solving: From SAD Environments to AI", In Forbus & Feltovich (eds.), *Smart Machines in Education*, Menlo Park, CA: AAAI Press, 2001, pp. 71-98.
- [7]Zimmerman, B. J., "A Social Cognitive View of Self-Regulated Academic Learning", Journal of Educational Psychology, 81(3), 1989, pp. 329-339.
- [8]Jeong, H., J. Wagster, R. Roscoe, G. Biswas, and D. Schwartz, "Using Hidden Markov Models to Characterize Student Behaviors Patterns in Computer-base Learning Environments", 9th Intl. Conference on Intelligent Tutoring Systems, Montreal, CA, 42(8), 2008.
- [9]Bryson, S., D. Kenwright, M. Cox, D. Elsworth, R. Haimes, "Visually exploring gigabyte data sets in real time", Communications of the ACM, 42(8), 1999.