

Investigating Self-Regulated Learning in Teachable Agent Environments

John S. Kinnebrew, Gautam Biswas, Brian Sulcer, and Roger S. Taylor
Department of EECSISIS
Box 351829, STA B
Vanderbilt University
Nashville, TN 37235

August 7, 2010

Abstract

We have developed a computer-based learning environment that helps students to learn about a variety of science topics by constructing causal concept maps. The system builds upon research in Learning-by-Teaching, and has each student take on the role and responsibilities of being the teacher of a virtual student named Betty. The environment is structured so that successfully instructing their teachable agent requires the students to learn and understand the science topic for themselves. This process is supported through the use of adaptive scaffolding provided by feedback from the system. This feedback is instantiated through the interactions with the teachable agent and a mentor agent, named Mr. Davis. For instance, if Betty performs poorly on a quiz, she might explicitly tell the student that she needs to learn more about that topic. Furthermore, Mr. Davis might suggest that the student ask Betty questions and then get her to explain her answers to find out where she may be making mistakes. Thus the system is designed to help students develop and refine their own knowledge construction and monitoring strategies as they teach their agent.

This chapter provides an overview of two studies that were conducted with 5th grade science students and a description of the analysis techniques that we have developed for interpreting students' activities in this learning environment. More specifically, we explain how we generated hidden Markov models (HMMs) of students' aggregated behavior patterns and how this exploratory technique helped us demonstrate how well students learned the metacognitive strategies presented in the Teachable Agents system. A second study uses verbal protocol analysis to compare students' acceptance of the strategies discussed by the two agents, and how this acceptance relates to their overall learning performance.

1 Introduction

Cognitive scientists have established that metacognition and self-regulation are important components for developing effective learning in the classroom and beyond [Bransford et al., 2000, Zimmerman, 2001]. The framework for self-regulated learning (SRL) originated from the social constructive theory of learning proposed by Bandura [Bandura, 1997], who postulated that learning is governed by three interacting factors: (i) *personal* (e.g., learners attitudes and beliefs); (ii) *behavioral* (e.g., the ability to invoke relevant prior knowledge, the ability to employ appropriate strategies to support learning); and (iii) *environmental* (e.g., type of instruction, quality of feedback, nature of interactions with parents and peers). A number of researchers (e.g., [Zimmerman et al., 1992, Pintrich, 2000, Zimmerman, 2001]) have demonstrated that students SRL capabilities can play a significant role in high school academic achievement. In addition, studies by [Brown and Palincsar, 1989] have demonstrated that through instruction younger students can acquire and apply metacognitive skills, such as planning and monitoring. However, students in typical classrooms are rarely provided opportunities to learn and exercise these strategies [Paris and Paris, 2001, Zimmerman, 1990].

For about eight years, our research team, the Teachable Agents Group, has been developing computer-based-learning environments that utilize the learning-by-teaching approach to instruction in order to foster students acquisition of knowledge and development of sophisticated metacognitive strategies. The system embodies the social constructive learning framework and provides students with opportunities for self-directed, open-ended learning in the domains of science and mathematics [Biswas et al., 2005, Blair et al., 2007, Leelawong and Biswas, 2008]. In the system, students are given a *knowledge construction* task in which they engage in the iterative process of reading and building causal concept maps for a range of instructional topics (e.g., ecology, thermoregulation). This process is enhanced through the *social interaction* component of the system in which students assume the role and responsibilities of being their agents teacher. The environment is structured so that successfully instructing their teachable agent (“Betty”) requires the students to learn and understand the topic for themselves. Our previous work has shown that

students find the task of teaching and interacting with Betty to be motivating, and it also helps them enhance their own learning [Chase et al., 2009, Schwartz et al., 2007, Schwartz et al., 2009]. The teachable agents performance is a function of how well it has been taught by the student, which provides the student with a less threatening way of assessing their own understanding and areas of confusion (e.g., “Ugh, Betty is so stupid, now Ive got to figure out another way to help her learn this stuff.”). Based upon the students level of progress and pattern of activities, the system triggers responses from Betty or Mr. Davis at appropriate times to provide the student with guidance on problem-solving and metacognitive strategies. As a result, the students are more likely to both increase their knowledge of the specific domain content and develop more sophisticated problem-solving and metacognitive strategies, which in turn helps their preparation for future learning [Biswas et al., 2005, Bransford and Schwartz, 1999, Schwartz and Martin, 2004, Schwartz et al., 2007].

This chapter presents analyses from several studies that were conducted in middle school science classrooms, in which students taught their agent about complex science topics, such as river ecosystems and global climate change. One of our goals was to determine the degree to which the agents metacognitive and SRL prompts could help improve students learning. Within this framework, we have developed analytical methods to identify and interpret students learning strategies based on their activity traces on the system. Such analyses can shed light on students underlying learning processes and the strategies they employ in achieving their learning tasks [Roscoe and Chi, 2007]. To date there has been very little work on deriving students self-regulated learning strategies from their activity sequences in computer-based learning environments (some exceptions are [Hadwin et al., 2007, Azevedo et al., 2009]). In this chapter, we present a novel methodology that derives hidden Markov models [Rabiner, 1989, Li and Biswas, 2002] from student activity sequences to quantify and assess student learning and metacognition. In addition, we report the results of a second study, where we performed verbal protocol analyses to determine students’ acceptance of the strategies discussed by the two agents, and how the feedback provided by the agents influenced their subsequent learning activities.

2 Measuring Self-Regulated Learning

To effectively design, test, and refine a system promoting SRL skills requires the ability to identify and measure metacognitive processes. The traditional approach to measuring students' self-regulated learning (SRL) has been through the use of self-report questionnaires (*e.g.*, [Pintrich et al., 1993b, Weinstein et al., 1987, Zimmerman and Martinez-Pons, 1986]). The underlying assumption in these questionnaires is that self-regulation is an aptitude that students possess. For example, the questionnaire items might attempt to assess students' inclination to elaborate as they read a passage, or to determine their approach to managing available time resources [Perry and Winne, 2006, Zimmerman, 2008]. This approach has been useful, as the self-report questionnaires have been shown to be good predictors of students' standard achievement test scores and they correlate well with achievement levels [Pintrich et al., 1993a, Zimmerman and Martinez-Pons, 1986]. However, Hadwin and others [Azevedo and Witherspoon, 2009, Hadwin et al., 2001, Hadwin et al., 2007, Perry and Winne, 2006] have argued that while the questionnaires provide valuable information about the learners' self-perceptions, they fail to capture the dynamic and adaptive nature of SRL as students are involved in learning, knowledge-building, and problem-solving tasks.

Increasingly, researchers have begun to utilize trace methodologies in order to examine the complex temporal patterns of self-regulated learning [Aleven et al., 2006, Azevedo and Witherspoon, 2009, Azevedo et al., 2009, Biswas et al., 2010, Hadwin et al., 2007, Jeong and Biswas, 2008, Zimmerman, 2008]. Perhaps the most common type of data collected, and the focus of this chapter, are computer logs, which can record every action that the student performs within a computer-based learning environment. An example of computer trace log analysis is presented in [Hadwin et al., 2007]. They performed a study that collected activity traces of 8 students using the gStudy system [Perry and Winne, 2006]. The activity traces were analyzed in four different ways: (i) frequency of studying events, (ii) patterns of studying activity, (iii) timing and sequencing of events, and (iv) content analyses of students' notes and summaries. The results of this analysis were compared against students' self-reports on their self-regulated learning. One

of the important findings was that many participants' self-reports of studying tactics, as determined by the MSLQ items, were not well calibrated with studying events traced in the gStudy system. The researchers found that the best matched item showed a 40% agreement, and the average agreement was 27%. The authors concluded from this study that trace data of student activity in e-learning environments are important for furthering our understanding of SRL.

More recently, trace data is being supplemented with other sources of data, such as concurrent verbal think-alouds (*e.g.*, [Azevedo and Witherspoon, 2009]) and measures of affect (*e.g.*, automatic recording of facial expression, posture, etc. [Burlison et al., 2004, D'Mello et al., 2007, DMello et al., 2008, Lester et al., 1997]. [Azevedo et al., 2009] have developed a hypermedia environment called MetaTutor to help students learn about complex and challenging science topics, such as the circulatory processes in human body systems. The system is also designed to train students in key SRL processes that relate to planning, metacognitive monitoring, learning strategies, and methods for handling task difficulties and demands. The authors used a combination of student trace data and think-aloud protocols to understand the nature of students' learning outcomes and their deployment of SRL processes. For example, one of their studies showed that students predominantly used strategies that pertained to the acquiring knowledge from the multimedia resources, and they only occasionally employ monitoring strategies to check on what they have learned [Azevedo and Witherspoon, 2009]. Combining trace and think-aloud protocols provides more insight into the students' thought processes that govern the use of strategies. Furthermore, they can be used to validate the results of the trace data analysis.

3 Betty's Brain and Self-Regulated Learning

The Betty's Brain system, illustrated in Figure 1, implements the learning-by-teaching paradigm to help middle school students develop cognitive and metacognitive skills in science and mathematics domains [Biswas et al., 2005, Blair et al., 2007, Schwartz et al., 2007]. The system supports five primary types of activities:

- *Read*: The system contains a set of indexed, hypermedia resources that students can access and use at any time while working on the system. These resources contain all of the science information needed for students to build their concept maps.
- *Edit*: Students explicitly teach Betty using a concept map representation [Novak, 1998], with the nodes being the relevant science concepts, and the links represent causal relations between the nodes. For example, fish eat (decrease) macroinvertebrates - therefore, an increase in fish causes a decrease in macroinvertebrates. Students teach Betty new concepts and links using a visual interface that includes menu selections and templates for adding and modifying information (e.g., the interface contains these four buttons: Teach Concept, Teach Link, Delete, and Edit).
- *Query*: Students use a template, illustrated in Figure 1, to check their teaching by asking Betty questions, which she answers using causal reasoning through chains of links [Forbus, 1984, Leelawong and Biswas, 2008].
- *Explain*: Students can probe Betty's reasoning, by asking her to explain her answer to a query. She demonstrates the use of causal reasoning processes to derive her answer, and verbalizes her reasoning process using speech and simultaneous animation on the concept map.
- *Quiz*: Students can assess how well they have taught Betty by having her take a quiz, which is a set of questions chosen by the Mentor agent. Betty's inability to answer some of the questions correctly typically motivates students to learn more so that they can make improvements to the concept map and help Betty do better on her quizzes.

Since our middle school students are novices in both the science topics and the teaching tasks, we provide them with a variety of scaffolds to help them overcome obstacles they may face in learning and teaching the domain material. In addition to answering queries and taking/administering quizzes, the agents also provide spontaneous feedback to the student on the relative effectiveness of

their teaching performance. This feedback is designed to help students develop and employ more metacognitive learning strategies [Schwartz et al., 2007, Tan et al., 2006, Wagster et al., 2007].

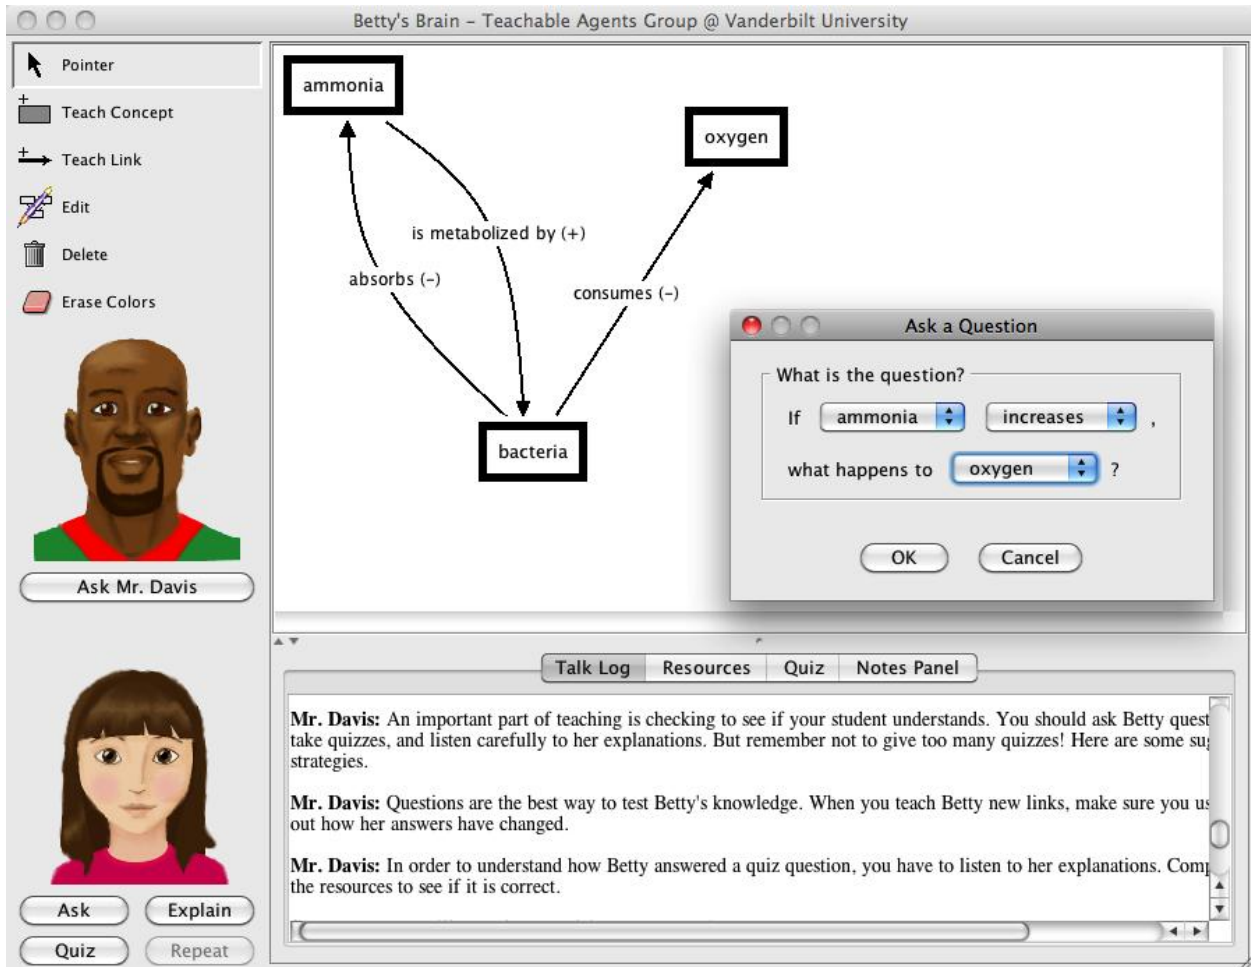


Figure 1: Betty's Brain system with query window

[Schunk and Zimmerman, 1997] point out that the self-regulation profiles of novice learners are quite distinct from those of experienced learners. Novices are often poor at forethought, and their self-judgment abilities are not well developed. These strategies can be taught, but students in typical classrooms are rarely provided opportunities needed to learn and master them. Our system addresses this problem by adopting a SRL framework that promotes a set of comprehensive skills: setting goals for learning new materials and applying them to map building tasks; deliberating about strategies to enable this learning; monitoring ones learning progress; and revising ones knowledge, beliefs, and strategies as new material and strategies are learned [Azevedo, 2005,

Schraw et al., 2002, Winne and Hadwin, 2008, Zimmerman, 2001].

Figure 2 illustrates our conceptual cognitive/metacognitive model that we employ in the design of the Betty’s Brain system. [Pintrich, 2002] differentiates between two major aspects of metacognition for learners: (i) *metacognitive knowledge* that includes knowledge of general strategies and when they apply, as well as awareness of one’s own abilities, and (ii) *metacognitive control* and self-regulatory processes that learners use to monitor and regulate their cognition and learning. In our model, metacognitive control is illustrated in the monitoring and knowledge construction strategies in Figure 2.

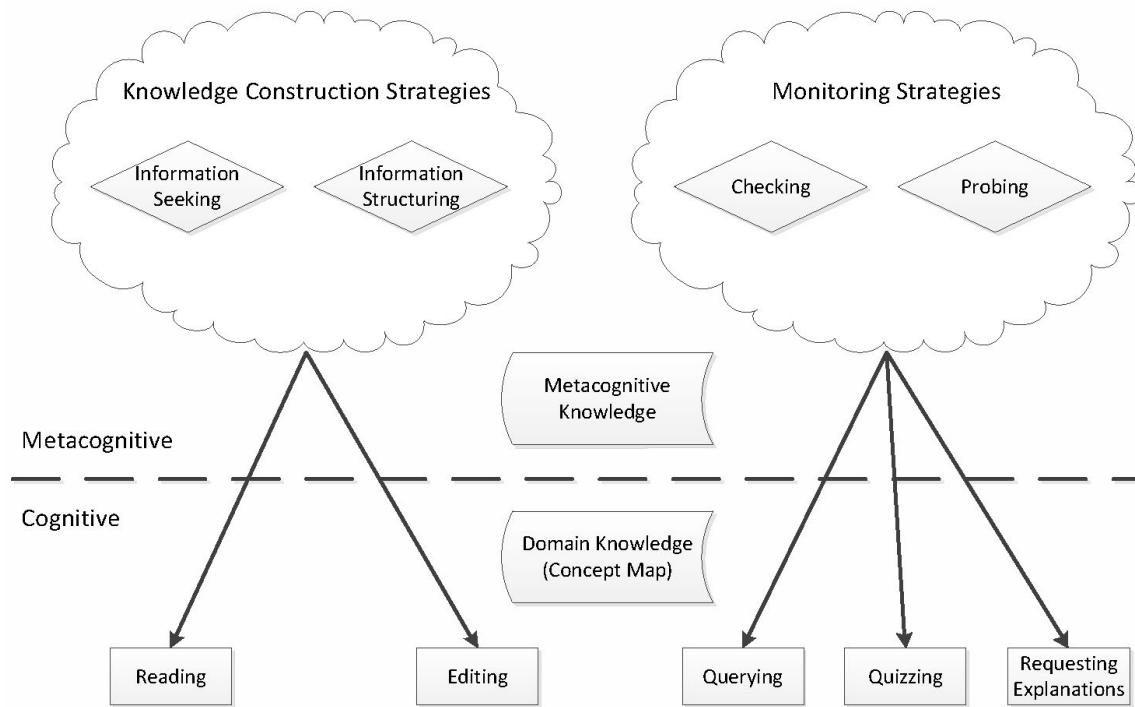


Figure 2: Our model of Self Regulated Learning strategies and Activities in the Betty’s Brain system linked to these strategies

For knowledge construction in the Betty’s Brain system (*i.e.*, building causal concept maps), we identify two key types of self-regulation strategies: (i) *information seeking*, in which students study and search available resources in order to gain missing domain information or remediate existing knowledge, and (ii) *information structuring*, in which students structure the information by causal and taxonomic relationships to build and revise their concept maps. Information seeking

involves strategies that focus on effective use of the resources in the system, whereas information structuring focuses on strategies for construction and revision of the concept map.

SRL Strategy	Triggering Activities	Betty Response	Mr. Davis Response
Knowledge Construction: Information Seeking	Student taught Betty several concepts, but not added many corresponding links.	<i>What do the resources say about that concept? Could we search the resources to learn more?</i>	<i>Reading is a very important part of learning. You should use the resources often to help you understand and double-check what you are teaching. Here are some strategies that you can use to improve your reading skills.</i>
Knowledge Construction: Information Structuring	Student has just added a number of links to the concept map.	<i>Can we double-check that these links are correct by reading the resources?</i>	<i>A good teacher explains how concepts affect each other. Use the links to teach Betty how the concepts cause other concepts to change.</i>
Monitoring: Checking	Betty has incorrect answers on the quiz she just took.	<i>Could you ask me some review questions to see how good my answers are?</i>	<i>Questions are the best way to test Betty's knowledge. When you teach Betty new links, make sure you use questions to find out how her answers have changed.</i>
Monitoring: Probing	Betty has got some of her quiz answers wrong.	<i>These quizzes can be tough. Can we go over my explanations to see how well I understand?</i>	<i>After each quiz, try to review a few of Betty's answers and explanations. Follow each step in her explanations carefully to be sure that they make sense.</i>
Monitoring: Pointing out a suboptimal strategy	Student is asking Betty to take too many quizzes without reading the resources or asking her questions.	<i>I want to do my best on the quizzes. It would help if I learned more between each quiz instead of taking so many.</i>	<i>You have been asking Betty to take a lot of quizzes recently. Try asking her questions and teach more information in between quizzes.</i>

Table 1: Examples of Agent Responses to Observed Student Behavior Patterns

The model also posits two types of monitoring strategies: (i) *checking*, where students use the query or the quiz features to test the correctness of their concept map, and (ii) *probing*, a stronger monitoring strategy, where students systematically analyze their map in greater detail, by asking for explanations, and following the causal reasoning steps generated by the agent to locate potential errors in the maps. Effective guidance (*i.e.*, relevant and timely feedback) based on this SRL model makes students aware and helps them develop good learning strategies, such as re-reading

the resources to check if there are errors in their concept maps (combining information seeking and checking strategies), and asking queries and checking explanations to find the source of an error (a probing strategy). Betty's persona, influenced by the social cognitive framework, incorporates metacognitive awareness that she conveys to the students at appropriate times to help them develop and apply monitoring and self-regulation strategies [Wagster et al., 2007, Schwartz et al., 2009]. Mr. Davis, the mentor, and, therefore, the more authoritative persona in the system, provides help in the form of suggested behavior patterns linked to effective SRL strategies (e.g., "if you are not sure, check the resources to see if Betty is answering her questions correctly."). In other words, the agent feedback is triggered by students' activity patterns and linked to strategies for knowledge construction and monitoring implied by our model. Examples of these embedded dialog prompts, guided by observed behavior patterns, appear in Table 1. Betty's responses combine motivational and self-regulation cues, whereas the Mentor, Mr. Davis', responses focus primarily on self-regulation strategies.

4 Experimental Studies

We have conducted several classroom studies where students use the teachable agents system to learn and gain a better understanding of a variety of science topics, such as river ecosystems, thermoregulation, and climate change. In these studies, the topics and specific science content provided by the system are closely linked to the 5th grade science curriculum. At the beginning of each study, the science teacher introduces students to the topic during regular classroom instruction. The intervention phase starts with an overview of causal relations and causal mapping during a 45-minute class period. This is followed by a hands-on training session with the system the next day. Over the next 4 or 5 days, the students teach Betty by building up a causal concept map for the science topic, which represents what Betty knows

To assess students acquisition of science domain knowledge and causal reasoning skills, we employ two measures. The first was pretest to posttest gain score. These tests contain a mix of

content-related multiple choice and free response items [?, ?] that were administered before the students were introduced to causal reasoning and at the end of the intervention, respectively. The second measure examined students' final maps, in terms of completeness and accuracy

In this chapter, we analyze the results from two classroom studies. The first study compared the students use of self regulated learning (SRL) strategies as they used the system. We had two questions: (1) Would students who taught an agent use more SRL strategies in their learning and teaching tasks? and (2) Would students who received SRL feedback from the agents use more sophisticated SRL strategies than students who did not? The second study used verbal protocol analysis to study the effectiveness of different kinds of SRL strategies, and also compared if the feedback provided by one agent was more effective than the feedback provided by the other agent. The results, and a discussion of these results are presented in the rest o this study.

4.1 Study 1: Modeling student's Self Regulated Learning (SRL) strategies

In this study, our goal was to determine if teaching the agent, Betty and providing metacognitive feedback would help students become better learners than those who did not teach or receive the feedback. Our participants were 56 students in two 5th grade science classrooms taught by the same teacher. Students were assigned to one of three conditions using stratified random assignment based on standardized test scores. All students created river ecosystem concept maps over five 45-minute sessions. Two of the conditions (i) the learning-by-teaching (LBT) group, and (ii) the self-regulated learning-by-teaching (SRL) group created their map to teach Betty so that she could pass a test on her own. The third group, the intelligent coaching system (ICS) group, provided the control condition and was told to create the map to learn for themselves. They could also query their map and ask for explanations, but in this case, it was Mr. Davis, and not Betty, who responded to them. Analyses that do not take into account the sequential nature of student interactions with the system, such as counting the frequency of student activities, can only capture limited or indirect information about learning strategies and students' cognitive state. We believe that employing a representation capable of modeling internal states and sequential characteristics of

student activities can aid analysis of student learning behaviors. One way of representing internal states and their interactions over time is a hidden Markov model (HMM) [Rabiner, 1989], which contains a set of states and probabilistic transitions between those states (more likely transitions are assigned higher probabilities). The states in a HMM are hidden, meaning that they cannot be directly observed in the environment/system. Instead, they produce output (*e.g.*, student activities in the Betty's Brain system) that can be observed. Deriving a hidden Markov model from activity traces requires simultaneous estimation of (1) the number of states; (2) the probabilities associated with transitions between states; (3) the probabilities associated with observing certain outputs (*i.e.*, particular student activities, such as reading or querying activities); and (4) the probability of a state being the initial state in an activity sequence.

By providing a concise representation of student learning strategies and behaviors, HMMs have the potential of providing a higher-level view of how students approach their learning tasks (*e.g.*, what strategies they use and how they switch between strategies) [Jeong and Biswas, 2008, Biswas et al., 2010]. Algorithms for learning an HMM from output sequences are well-known but require appropriate configuration/initialization parameters for effective use [Rabiner, 1989]. Specifically, HMM learning algorithms require an initial HMM whose parameters are then modified to maximize the likelihood of producing observed output sequences. In particular, the number of states in the HMM and their initial output probabilities can have a significant effect on the resulting, learned HMM.

We have developed an algorithm designed to generate HMMs from a set of student activity sequences [Li and Biswas, 2000, Li and Biswas, 2002, Jeong and Biswas, 2008]. The first step in the analysis is to extract each students' activity sequences over the period of the study from the log files. Each element of the activity sequence is labeled as one of the five primary activity types: READ (read the resources), QUER (query Betty on a portion of the map), EXPL (ask Betty to explain her answer to a query), QUIZ (ask Betty to take a quiz), and EDIT (add, edit, or delete concepts and links in the concept map).

Although all students had access to the full set of actions, not all of them used them effectively.

Using queries to check whether recent revisions to the map were correct, or to locate errors in the concept map are examples of effective use of queries. On the other hand, asking questions simply to make Betty speak, so that the student could make fun of her mechanical, computer-generated voice is clearly an ineffective use of queries for the learning task. When students generated questions that were not related to parts of the map they had worked on recently, it was unclear whether these queries could be associated with effective learning. We addressed this issue by developing a *relevance score* that took into account how much the current action could be linked to other recent actions.

Each student action was assigned a relevance score that depended on the number of relevant previous actions within a pre-specified window. This score provides a measure of *informedness* for knowledge construction activities and, similarly, a measure of *diagnosticity* for monitoring activities. Overall, the relevance score provides a rough measure of strategy consistency or coherence over a sequence of actions. For this analysis, a prior action was considered relevant to the current action if it was related to, or operated on, one of the same map concepts or links. For example, if a student edited a link that was used to generate an answer a recent query, the query action was counted in the edit's relevance score. The increased relevance score suggested a more informed edit action because it was related to a recent query.

The relevance score is employed in HMM generation by refining the labeling of student activities. Each of the actions in an activity sequence is assigned a label, H (high) or L (low), based on its relevance score, in order to maintain the context and relevance information of the actions in the sequence. For example, a QUER-H activity implies that the query the student asked is related to other activities recently performed, while a QUER-L implies that the query activity is unrelated to the students' recent activities.

The HMM models derived for the ICS, LBT, and SRL groups are shown in Figure 3. The possible transitions between states are shown as arrows, and the transition probabilities are expressed as percentages. For example, the ICS behavior model indicates that there is an 84% likelihood that a student who just performed an applied reading action will next perform another applied reading

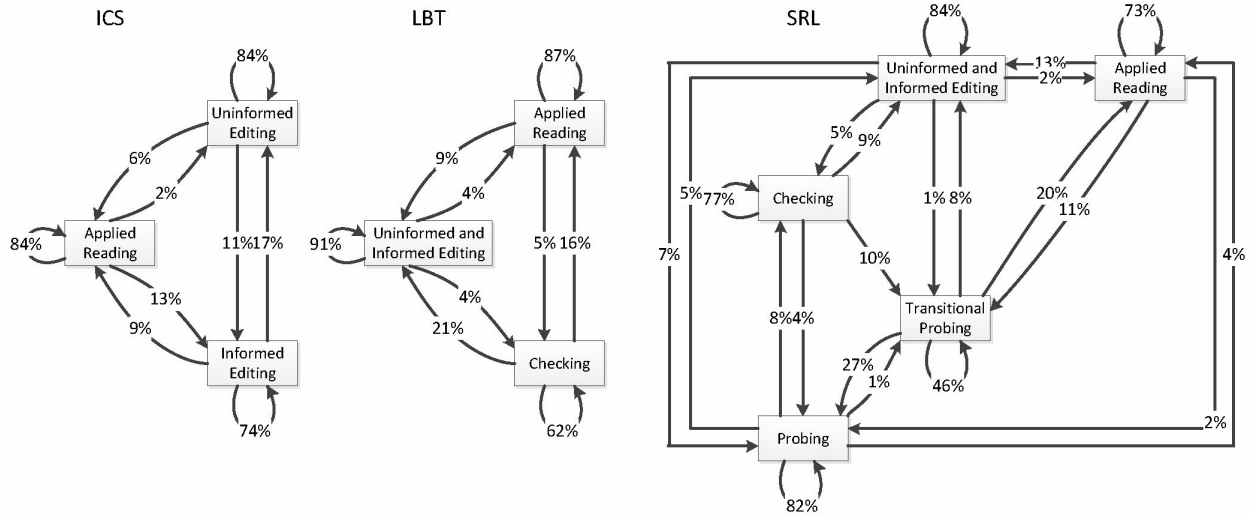


Figure 3: HMMs Derived from Student Activity Sequences

action, but there is a 13% chance that the student will perform an informed editing action next. States in the models are named based on an interpretation of their outputs (activities) illustrated in Figures 4 and 5. The models for the ICS and LBT groups each have three states, but the activities associated with some of those states differ significantly. Therefore, the states are interpreted, and named, differently for those groups. Further, the derived model for the SRL group has five states instead of three and shows some interesting differences in the set of actions associated with those states.

We used the activities associated with a state to categorize the states of the three derived HMM models. This produced seven different states that are described below.

1. *Applied reading* - students are primarily engaged in reading the resources and applying the knowledge gained from reading by editing their maps. This state combines information-seeking strategies with informed information structuring.
2. *Uninformed editing* - students are primarily making uninformed changes to their map, indicating the use of trial-and-error or guessing strategies for information structuring. Students may generate queries, but the queries generally do not relate directly to the editing activities. This represents a suboptimal information structuring strategy.

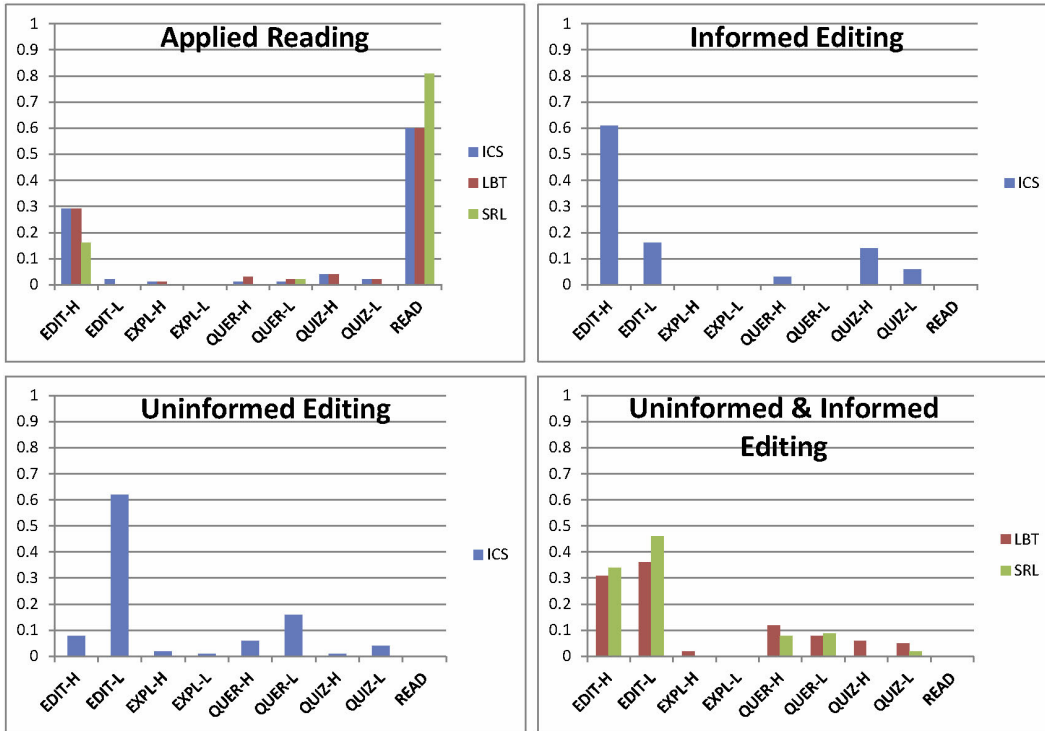


Figure 4: Activities in Knowledge Construction States

3. *Informed editing* - students are primarily making informed changes to their map (information structuring) based on relevant queries or quiz questions. As opposed to uninformed editing, the students are using queries and quizzes to guide their map editing actions.
4. *Uninformed and informed editing* - students are primarily making changes to their map, some of which are based on relevant queries or quizzes. This state combines the uninformed editing and informed editing states, including situations where students are making edits relevant to recent queries and quizzes, as well as situations in which students are making edits without focusing on a single area of the map.
5. *Checking* - students are querying and quizzing Betty to check the correctness of their concept maps. However, the use of queries and quizzes is unfocused. For example, queries are not related to recently edited areas of the map, and it is not clear that students are using the quiz results to focus on areas on the map where there are errors. Therefore, this state corresponds to a weak monitoring strategy.

6. *Probing* - students combine querying and quizzing with the explanation feature, which illustrates the chain of links that were followed to generate an answer to a question. Further, the queries, explanations, and quizzes are focused on a particular area of the map, and the results inform map editing. This combination implies a deeper, more focused monitoring strategy than the checking state and may be evidence of metacognitive reflection on the quality of the student's map/knowledge.

7. *Transitional probing* - students perform activities similar to the probing state, but generally with lower relevance scores, suggesting that they may be transitioning to probing a different area of the concept map.

As discussed above, each of the interpreted states can be mapped onto one or more knowledge construction and monitoring strategies outlined in our conceptual SRL model that was illustrated in Figure 2. The models provide evidence that the SRL condition uses more effective monitoring and knowledge construction strategies than the LBT and ICS conditions.

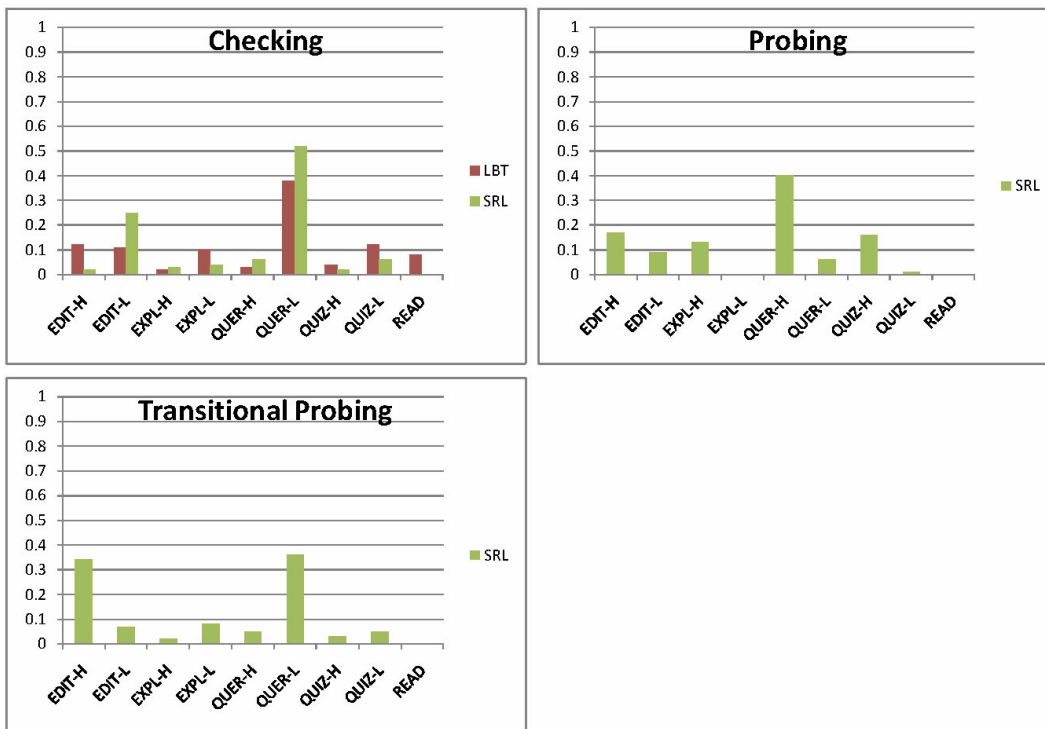


Figure 5: Activities in Monitoring States

We probed further to determine the prevalence of individual states suggested by a generated HMM. To do this, we calculate the proportion of *expected state occurrences* by condition in Table 2. This measure employs state transition probabilities and average activity sequence lengths to calculate an expected value for the proportion of individual state occurrences. Although states corresponding to knowledge construction behaviors account for a significant percentage of behaviors in all groups, the HMMs for LBT and SRL groups also show use of monitoring strategies (10% for LBT and 49% for SRL). The SRL HMM also includes more states suggesting a greater number (and possibly greater complexity) in the types of strategies employed. Further, the activities involved in these additional states suggest use of the more advanced monitoring behavior of probing that is absent from the ICS and LBT HMMs.

Behaviors	ICS	LBT	SRL
	Proportion (%)	Proportion (%)	Proportion (%)
Applied Reading	33	30	17
Uninformed Editing	36	-	-
Uninformed and Informed Editing	-	60	34
Informed Editing	31	-	-
Checking	-	10	13
Transitional Probing	-	-	7
Probing	-	-	29

Table 2: Proportion of expected state occurrences by condition

The results of the HMM analysis identify differences in strategies employed by the different groups of students, but it does not directly indicate how this affected student learning. Table 3 shows the learning gains measured by tests and map scores for each condition in the study. Results show that the two groups that taught Betty (LBT and SRL) outperformed the ICS group on gains in both test and map scores with statistically-significant performance differences between the SRL and ICS groups. The gain score differences between conditions for the multiple choice test questions were not statistically significant. However, for the free response questions, the SRL group showed greater gains than the ICS group ($p < 0.1$ and a relatively large effect size of $\hat{d} = 0.72$). For the gain in correct map concepts, the SRL group outperformed the ICS group ($p < 0.05$, $\hat{d} = 0.81$), and the LBT group ($p < 0.01$, $\hat{d} = 1.05$). Similarly, for the gain in correct map links, the

SRL group again outperformed the ICS and LBT groups ($p < 0.05$, $\hat{d} = 0.97$ and $p < 0.1$, $\hat{d} = 0.72$, respectively).

Gain Score	Conditions		
	ICS	LBT	SRL
Multiple Choice	0.4 (2.4)	1.1 (3.1)	0.4 (1.5)
Free Response	1.9 (3.0)	4.3 (3.2)	4.8 (4.7)
Map Concepts	8.1 (2.4)	7.3 (2.7)	10.4 (3.1)
Map Links	12.2 (3.8)	12.7 (5.3)	16.2 (4.4)

Table 3: Mean pre-to-post test and concept map score gains

It is clear from these results that the students who taught Betty outperformed the other students, both in learning gains and the use of more effective strategies. The ICS students use of ineffective strategies, as apparent in the hidden Markov models, may explain their smaller learning gains. Further, the fact that the SRL group had higher free response and map score gains than the LBT group (although not all of the differences were statistically significant for the number of students in this study), further supports the results of the HMM analysis that the SRL feedback led to use of better learning strategies, and as a result, students performed better in their learning tasks. Moreover, while 60% of SRL students completed their concept maps during the five sessions, only 44% of LBT students and 31% of ICS students were able to complete their concept maps. The results of the HMM analysis, combined with the results on learning gains, suggest that the metacognitive feedback helped students learn SRL strategies, which allowed them to more effectively learn the science content. Although the HMM analysis illustrates the effectiveness of providing metacognitive feedback in Betty’s Brain, it does not indicate which agent or types of feedback were most effective. This was the focus of the second study, which we describe next.

4.2 Study 2: Comparing the Mentor and Teachable Agent Feedback

In order to assess the effectiveness of different forms of feedback in our system, we conducted a study, which included a talk-aloud protocol to determine students’ reactions to the agent feedback. The study was conducted in three 5th grade science classrooms in the same school as study 1¹.

¹study 1 and study 2 were conducted in different years.

Two of the classrooms had the same science teacher as in study 1. The third classroom had a different teacher, but teacher 2 worked closely with teacher 1 for the unit taught in this study. All students worked on a newer version of the SRL system from study 1. In this version of the system, the feedback from the two agents was better organized into the categories described in Table 1.

Students worked in pairs chosen by the teachers to ensure that the paired students were at similar academic levels and had compatible personalities. Before the study began, the teachers instructed students on how to collaborate on the system. The students had to discuss with one another and come to a consensus before they performed an action on the system. Control of the keyboard and mouse alternated between the students (*e.g.*, if one student had control of the input devices on day 1, then the partner was given control on day 2). The science teachers ran a brief practice session on working with the system before the students started this phase of the study. All students had worked individually with the Teachable Agent system on another science unit (ecosystem), so they were familiar with the system.

Students worked on the topic of pollution in river ecosystems for three 45 minute periods. We recorded student conversations and interactions using webcams. After the study was concluded, two coders reviewed all of the video data and recorded students responses to the feedback. For every instance in which the TA or the mentor provided feedback, the coders noted whether the students' subsequent discussion *affirmed*, *dismissed*, or *deferred* the agent's feedback. Inter-rater reliability for the coding scheme was over 85%, and the results are summarized in Table 4.

Students explicitly referenced the feedback from the TA, Betty and the mentor, Mr. Davis about a third of the time (34% for Betty and 30% for Mr. Davis). In these discussions, they sometimes affirmed (*e.g.*, "We should do that" or "We need to read more" responding to feedback suggesting the students read the resources), sometimes dismissed (*e.g.*, "No, I don't want to read" or "Let's just keep giving her quizzes" responding to feedback suggesting students teach Betty more between giving her quizzes), and occasionally deferred (*e.g.*, "Hold on, we will get to that in a second") agent feedback².

²Some student discussions that referenced the feedback neither affirmed, dismissed, or deferred the feedback.

As illustrated in Table 4, there were differences between how frequently the students affirmed or dismissed feedback from the two agents. Students were more likely to affirm feedback from Mr. Davis, and were more likely to dismiss feedback from Betty. This suggests that students paid less attention to the self-reflective feedback from Betty, who was treated like a peer or even someone who did not seem to know as much as the students, than they did to the more explicit, strategy-oriented feedback from Mr. Davis, who had the status of an authority figure.

Agent →	N	Referenced	Affirmed	Dismissed	Deferred
Betty	649	34%	6%	16%	2%
Mr. Davis	275	30%	18%	10%	2%

Table 4: Student Verbal Response to Agent Feedback

To understand how students’ verbal responses related to learning, we analyzed the study results for the two metacognitive categories of feedback from each agent: (1) knowledge construction strategies, and (2) monitoring strategies. Table 5 summarizes the percentages of each type of verbal reaction to the different forms of feedback, as well as their correlation with the student pair’s final map score. Students who more frequently affirmed the knowledge construction strategy feedback from either the TA or the mentor had higher map scores, but the correlations were not statistically significant. Students who dismissed either the knowledge construction or the monitoring feedback from either agent had lower map scores (negative correlations). However, when the students affirmed the monitoring feedback, the results were surprising. Affirming Mr. Davis’s monitoring feedback showed a positive correlation with map score (not statistically significant), but affirming Betty’s monitoring feedback resulted in a negative correlation with map score ($p < 0.05$). We discuss this result in greater detail later, but overall the students seemed to affirm the knowledge construction feedback more, and affirming this feedback implied higher map scores.

The verbal responses to feedback listed in Table 4 suggest a difference in the way students react to the two agents, and these reactions affect their concept map building performance. For example, those who affirmed Mr. Davis’s knowledge construction and monitoring feedback seemed to do better in their map building task. To determine whether the students verbal responses to feedback matched their expected actions in the system, we analyzed student actions immediately following

Agent → Category ↓	Betty			Mentor		
	Percent Verbal Response (Correlation with Map Score)			Percent Verbal Response (Correlation with Map Score)		
	Affirm	Dismiss	Defer	Affirm	Dismiss	Defer
Knowledge Construction	9% (0.20)	27% (-0.45 ^b)	2% (-0.16)	20% (0.37)	22% (-0.31)	2% (0.24)
Monitoring	4% (-0.46 ^b)	11% (-0.35)	2% (0.04)	17% (0.28)	6% (-0.44 ^b)	2% (-0.16)

Table 5: Verbal Responses to feedback and Corresponding Map Score Correlations (^b p < .05)

each agent feedback statement. For example, if Betty said “Can we go over my explanations to see if I am missing anything,” we checked to see if the subsequent student actions included asking Betty to explain an answer. Table 6 reports for both Betty and Mr. Davis (1) the average number of feedback events by category per student, (2) the average proportion of subsequent activities that matched the actions advised by the feedback (using a window size of 3 actions³), and (3) the correlation between the percentage of matched actions and the students’ final map scores.

Feedback Category → Agent ↓	Measures	Knowledge Construction	Monitoring
Betty	Feedback Events	5.77	13.08
	Action Match	28.2%	0.5%
	Map score Correlation	0.52 ^a	-0.41
Mr. Davis	Feedback Events	2.46	5.54
	Action Match	24.02%	33.0%
	Map score Correlation	0.11	0.26

Table 6: Action Response to Feedback and Corresponding Map Score Correlations (^a p < .1)

Overall, the correlation between percentage of matching actions and the students’ final map scores was positive (0.34 for Betty and 0.36 for the Mentor), but the correlations were not statistically significant. More detailed analysis by category of metacognitive feedback, showed a positive correlation between students’ final map scores and their following Betty’s and Mr. Davis’s advice on knowledge construction feedback. Students were more likely to follow Betty’s feedback suggestions than that of the Mentor(but the differences were small (28% to 24%). The more the

³We tried a number of different window sizes, and all of them produced similar results.

students subsequent actions matched the feedback the better was their map scores, as measured by the correlations: 0.52 for following Betty's knowledge construction advice ($p < .1$) and 0.11 for matching the Mentor's knowledge construction advice. These results differ from the verbal responses to feedback, where students affirmed Mr. Davis's knowledge construction feedback more than they did Betty's, and the corresponding correlations with map scores was also higher for the mentor (0.37 versus 0.2).

On the other hand, for monitoring feedback Mr. Davis was more effective than Betty. Though the relative number of Betty monitoring feedback events was high compared to Mr. Davis's (13.1 to 5.5), student actions after Betty feedback showed a poor match to the feedback content (about 0.5%). For the mentor feedback the match was at 33%. Combining this information with the verbal response results indicates that the students were more dismissive of Betty's monitoring feedback, and at the same time they rarely followed up with activities that matched the feedback content. In addition, the correlation between the activity match percentage and students' map scores was negative, implying those who affirmed Betty's monitoring feedback or tried to apply it ended up with lower map scores. On the other hand, Mr. Davis's monitoring feedback had more affirmations and there were more attempts to follow his suggestion, and also resulted in higher map scores (though the correlations were not statistically significant.)

Together the verbal and action response results show clear differences in the way students responded to the two agents. Overall, the students affirmed the mentor's feedback more than they did the teachable agent's, and in general, higher affirmation levels implied better final map scores. These results also indicate that the monitoring feedback was less effective than knowledge construction feedback. With the exception of Betty's monitoring feedback, the results showed positive correlations with map scores for both verbal and action response measures. Although there are many potential explanations for the negative correlation between Betty's monitoring feedback and map scores, the results suggest that her monitoring feedback was generally ineffective in helping students improve their concept maps. This could imply that students could not understand Betty's feedback and, therefore, they did not apply her suggestions during their learning and teaching tasks.

The few who did, applied them inappropriately, and, therefore, used up time that could have been more productively spent in other activities. However, those who followed similar feedback from Mr. Davis did better on their map. Overall, the results indicate that the metacognitive feedback had a generally positive effect on students' learning, but the more explicit strategy feedback from the Mentor agent was more effective than Betty's self-evaluative statements and suggestions to use strategies.

5 Discussion and Conclusions

The Betty's Brain system is designed to leverage the benefits of learning by teaching and causal reasoning to help students learn science. The teaching interactions and agent feedback support students' engagement and promote the development and use of educationally-productive cognitive and metacognitive processes. In study 1, students who utilized learning by teaching versions of our system (*i.e.*, the LBT and SRL groups) constructed better concept maps than students who used the non-teaching ICS version of the system. Moreover, students' performance was strongest when the system explicitly supported their use of self-regulated learning strategies by having Betty model and prompt for such behaviors, and having the mentor provide additional strategy-oriented advice.

Our approach to analyzing students activity sequences using hidden Markov models produced good results. We were able to characterize students activity patterns into a number of (good and bad) knowledge construction and monitoring strategies. The interpretation of student behaviors with the HMMs also matched the SRL feedback model we implemented in the Betty's Brain system.

Although the HMM analysis illustrated the effectiveness of providing metacognitive feedback in the Betty's Brain system, it did not indicate which agent or types of feedback were most effective in promoting SRL behaviors. Our second study included a talk-aloud protocol to determine students' reactions to the agent feedback. We combined the talk-aloud protocols with analysis of student activity traces to develop a more complete picture of how well students employed the feed-

back to their map building tasks. Overall, students' verbal responses to agent feedback indicated that they were more receptive to the explicit, strategy-oriented advice from the mentor agent, as opposed to the self-reflective, but less explicit, feedback from the teachable agent. Further, students were more likely to affirm the knowledge construction feedback from each agent than the monitoring feedback. This analysis also showed a positive correlation between affirming feedback and students' map scores, except in the case of Betty's monitoring feedback.

Additional analysis of student responses to feedback, in terms of actions taken following feedback events, showed a similar differentiation between knowledge construction and monitoring feedback. Students taking more actions consistent with an attempt to apply the knowledge construction feedback tended to have better map scores. However, students taking more actions advised by Betty's monitoring feedback tended to have lower map scores, suggesting they were unable to apply the strategies suggested by that feedback. This brings up a number of issues. It suggests that students find it easier to understand and apply knowledge construction strategies (*e.g.*, read the resources to find the correct relation between two concepts or check the resources to see if all of the required concepts appear on the map) than monitoring strategies (*e.g.*, ask a query to check if the map is correct or ask for an explanation to check an answer step-by-step to identify errors). Other studies, such as [Azevedo et al., 2009] also suggest that students rarely employ monitoring strategies during learning and knowledge construction tasks, but frequently apply a variety of other metacognitive strategies. It may also be true that students may understand a monitoring strategy, but may not know when to apply it. For example, when constructing their concept map, students may not know when to switch from map building to map checking and back in an effective way. Moreover, they may have difficulty in formulating "good" queries that helps them check a relevant part of their map. Therefore, monitoring strategy feedback may need to be presented in more elaborate detail with justification of its importance in the learning task. For example, analysis of the context and details of advised actions (*e.g.*, Betty's feedback "Can we go over my explanation step by step and check it with the resources?") suggests the use of explain and read actions, but effective application of the feedback involves reading sections of the resources

related to the map concepts and links in the current query and explanation. Some of these details may need to be built into the feedback mechanisms, especially in the early stages, to help students learn *when* and *how* to apply strategies in an effective way.

Since students appeared to be more receptive to the explicitly strategy-oriented feedback from the more authoritative agent, i.e., the mentor, it may be especially fruitful to improve the mentor agent's feedback. We intend to continue analyzing the data from this and future studies in order to better understand how specific phrasing and different forms of metacognitive feedback affect student behavior. We have also been conducting studies to determine how to make the timing and content of strategy feedback more relevant to the student's current activities on the system.

In addition to analyzing and enhancing the agent feedback to promote metacognitive strategies and prepare students for future learning, we also plan to refine our HMM analysis technique. Enhanced HMM analysis could provide a better understanding of the different strategies employed by students when learning complex science topics and allow for more adaptive feedback suited to the current context of the student's activities. In particular, we intend to employ clustering of individual student HMMs to improve the accuracy of our HMM analysis and use sequence mining to pre-process the trace data in the HMM analysis to maintain more of the temporal information in the aggregated behaviors of HMM states.

Acknowledgments

This work has been supported by Dept. of ED IES grant #R305H060089, NSF REESE Award #0633856, and NSF IIS Award #0904387.

References

[Aleven et al., 2006] Aleven, V., McLaren, B., Roll, I., and Koedinger, K. (2006). Toward meta-cognitive tutoring: A model of help seeking with a Cognitive Tutor. *International Journal of Artificial Intelligence in Education*, 16(2):101–128.

- [Azevedo, 2005] Azevedo, R. (2005). Using hypermedia as a metacognitive tool for enhancing student learning? The role of self-regulated learning. *Educational Psychologist*, 40(4):199–209.
- [Azevedo and Witherspoon, 2009] Azevedo, R. and Witherspoon, A. (2009). Self-regulated use of hypermedia. In Graesser, A., Dunlosky, J., and Hacker, D., editors, *Handbook of metacognition in education*. Erlbaum, Mahwah, NJ.
- [Azevedo et al., 2009] Azevedo, R., Witherspoon, A., Chauncey, A., Burkett, C., and Fike, A. (2009). MetaTutor: A MetaCognitive Tool for Enhancing Self-Regulated Learning. In *Annual Meeting of the American Association for Artificial Intelligence, Symposium on Metacognitive and Cognitive Educational Systems*, pages 14–19.
- [Bandura, 1997] Bandura, A. (1997). *Self-efficacy: The exercise of control*. Freeman, New York, NY.
- [Biswas et al., 2010] Biswas, G., Jeong, H., Kinnebrew, J., Sulcer, B., and Roscoe, R. (2010). Measuring Self-regulated Learning Skills through Social Interactions in a Teachable Agent Environment. *Research and Practice in Technology-Enhanced Learning (RPTEL)*. In Review.
- [Biswas et al., 2005] Biswas, G., Leelawong, K., Schwartz, D., Vye, N., and Vanderbilt, T. (2005). Learning by teaching: A new agent paradigm for educational software. *Applied Artificial Intelligence*, 19(3):363–392.
- [Blair et al., 2007] Blair, K., Schwartz, D., Biswas, G., and Leelawong, K. (2007). Pedagogical agents for learning by teaching: teachable agents. *Educational Technology & Society: Special Issue on Pedagogical Agents*, 47(1).
- [Bransford et al., 2000] Bransford, J., Brown, A., and Cocking, R., editors (2000). *How people learn*. National Academy Press Washington, DC, Washington, D.C.
- [Bransford and Schwartz, 1999] Bransford, J. and Schwartz, D. (1999). Rethinking transfer: A simple proposal with multiple implications. *Review of research in education*, 24(1):61.
- [Brown and Palincsar, 1989] Brown, A. and Palincsar, A. (1989). Guided, cooperative learning and individual knowledge acquisition. In Resnick, L., editor, *Knowing, learning, and instruction: Essays in honor of Robert Glaser*, pages 393–451. Lawrence Erlbaum Associates, Hillsdale, NJ.
- [Burlison et al., 2004] Burlison, W., Picard, R., Perlin, K., and Lippincott, J. (2004). A platform for affective agent research. In *Workshop on Empathetic Agents, International Conference on Autonomous Agents and Multiagent Systems*, New York, NY.
- [Chase et al., 2009] Chase, C., Chin, D., Opezzo, M., and Schwartz, D. (2009). Teachable agents and the protégé effect: Increasing the effort towards learning. *Journal of Science Education and Technology*, 18(4):334–352.
- [DMello et al., 2008] DMello, S., Craig, S., Witherspoon, A., Mcdaniel, B., and Graesser, A. (2008). Automatic detection of learners affect from conversational cues. *User Modeling and User-Adapted Interaction*, 18(1):45–80.

- [D’Mello et al., 2007] D’Mello, S., Picard, R., and Graesser, A. (2007). Toward an affect-sensitive AutoTutor. *IEEE Intelligent Systems*, 22:53–61.
- [Forbus, 1984] Forbus, K. (1984). Qualitative process theory. *Artificial intelligence*, 24(1-3):85–168.
- [Hadwin et al., 2007] Hadwin, A., Nesbit, J., Jamieson-Noel, D., Code, J., and Winne, P. (2007). Examining trace data to explore self-regulated learning. *Metacognition and Learning*, 2(2):107–124.
- [Hadwin et al., 2001] Hadwin, A., Winne, P., Stockley, D., Nesbit, J., and Woszczyna, C. (2001). Context moderates students’ self-reports about how they study. *Journal of Educational Psychology*, 93(3):477–487.
- [Jeong and Biswas, 2008] Jeong, H. and Biswas, G. (2008). Mining student behavior models in learning-by-teaching environments. In *Proceedings of The First International Conference on Educational Data Mining*, pages 127–136, Montreal, Quebec, Canada.
- [Leelawong and Biswas, 2008] Leelawong, K. and Biswas, G. (2008). Designing learning by teaching agents: The Betty’s Brain system. *International Journal of Artificial Intelligence in Education*, 18(3):181–208.
- [Lester et al., 1997] Lester, J., Converse, S., Kahler, S., Barlow, S., Stone, B., and Bhogal, R. (1997). The persona effect: affective impact of animated pedagogical agents. In *Proceedings of the SIGCHI conference on Human factors in computing systems (CHI ’97)*, pages 359–366.
- [Li and Biswas, 2000] Li, C. and Biswas, G. (2000). A Bayesian approach to temporal data clustering using hidden Markov models. In *Proceedings of the Seventeenth International Conference on Machine Learning*, pages 543–550.
- [Li and Biswas, 2002] Li, C. and Biswas, G. (2002). A Bayesian Approach for Learning Hidden Markov Models from Data. *Scientific Programming: Special Issue on Markov Chain and Hidden Markov Models*, 10(3):201–219.
- [Novak, 1998] Novak, J. (1998). *Learning, Creating, and Using Knowledge: Concept Maps as Facilitative Tools in Schools and Corporations*. Lawrence Erlbaum Associates, Mahwah, NJ.
- [Paris and Paris, 2001] Paris, S. and Paris, A. (2001). Classroom applications of research on self-regulated learning. *Educational Psychologist*, 36(2):89–101.
- [Perry and Winne, 2006] Perry, N. and Winne, P. (2006). Learning from learning kits: gStudy traces of students self-regulated engagements with computerized content. *Educational Psychology Review*, 18(3):211–228.
- [Pintrich, 2000] Pintrich, P. (2000). An Achievement Goal Theory Perspective on Issues in Motivation Terminology, Theory, and Research* 1. *Contemporary Educational Psychology*, 25(1):92–104.

- [Pintrich, 2002] Pintrich, P. (2002). The role of metacognitive knowledge in learning, teaching, and assessing. *Theory into Practice*, 41(4):219–225.
- [Pintrich et al., 1993a] Pintrich, P., Marx, R., and Boyle, R. (1993a). Beyond cold conceptual change: The role of motivational beliefs and classroom contextual factors in the process of conceptual change. *Review of Educational Research*, 63(2):167–199.
- [Pintrich et al., 1993b] Pintrich, P., Smith, D., Garcia, T., and McKeachie, W. (1993b). Reliability and predictive validity of the Motivated Strategies for Learning Questionnaire (MSLQ). *Educational and psychological measurement*, 53(3):801–813.
- [Rabiner, 1989] Rabiner, L. (1989). A tutorial on hidden Markov models and selected applications in speech recognition. *Proceedings of the IEEE*, 77(2):257–286.
- [Roscoe and Chi, 2007] Roscoe, R. and Chi, M. (2007). Understanding tutor learning: knowledge-building and knowledge-telling in peer tutors’ explanations and questions. *Review of Educational Research*, 77(4):534–574.
- [Schraw et al., 2002] Schraw, G., Kauffman, D., and Lehman, S. (2002). Self-regulated learning theory. In Nadel, L., editor, *The encyclopedia of cognitive science*, pages 1063–1073. Nature Publishing Company, London.
- [Schunk and Zimmerman, 1997] Schunk, D. and Zimmerman, B. (1997). Social origins of self-regulatory competence. *Educational Psychologist*, 32(4):195–208.
- [Schwartz et al., 2007] Schwartz, D., Blair, K., Biswas, G., Leelawong, K., and Davis, J. (2007). Animations of thought: Interactivity in the teachable agent paradigm. In Lowe, R. and Schnotz, W., editors, *Learning with animation: Research and implications for design*, pages 114–140. Cambridge University Press, UK.
- [Schwartz et al., 2009] Schwartz, D., Chase, C., Chin, D., Oppezzo, M., Kwong, H., Okita, S., Biswas, G., Roscoe, R., Jeong, H., and Wagster, J. (2009). Interactive metacognition: Monitoring and regulating a teachable agent. In Hacker, D., Dunlosky, J., and Graesser, A., editors, *Handbook of metacognition in education*, pages 340–358. Routledge Press.
- [Schwartz and Martin, 2004] Schwartz, D. and Martin, T. (2004). Inventing to prepare for future learning: The hidden efficiency of encouraging original student production in statistics instruction. *Cognition and Instruction*, 22(2):129–184.
- [Tan et al., 2006] Tan, J., Biswas, G., and Schwartz, D. (2006). Feedback for metacognitive support in learning by teaching environments. In *Proceedings of the 28th Annual Meeting of the Cognitive Science Society*, pages 828–833, Vancouver, Canada.
- [Wagster et al., 2007] Wagster, J., Tan, J., Wu, Y., Biswas, G., and Schwartz, D. (2007). Do learning by teaching environments with metacognitive support help students develop better learning behaviors. In *Proceedings of the 29th Annual Meeting of the Cognitive Science Society*, pages 695–700, Nashville, TN.

- [Weinstein et al., 1987] Weinstein, C., Schulte, A., and Palmer, D. (1987). *The Learning and Study Strategies Inventory*. H & H Publishing, Clearwater, FL.
- [Winne and Hadwin, 2008] Winne, P. and Hadwin, A. (2008). The weave of motivation and self-regulated learning. In Schunk, D. and Zimmerman, B., editors, *Motivation and self-regulated learning: Theory, research, and applications*, pages 297–314. Taylor & Francis, NY.
- [Zimmerman, 1990] Zimmerman, B. (1990). Self-regulating academic learning and achievement: The emergence of a social cognitive perspective. *Educational Psychology Review*, 2(2):173–201.
- [Zimmerman, 2001] Zimmerman, B. (2001). Theories of self-regulated learning and academic achievement: An overview and analysis. In Zimmerman, B. and Schunk, D., editors, *Self-regulated learning and academic achievement: Theoretical perspectives*, pages 1–37. Erlbaum, Mahwah, NJ.
- [Zimmerman, 2008] Zimmerman, B. (2008). Investigating self-regulation and motivation: Historical background, methodological developments, and future prospects. *American Educational Research Journal*, 45(1):166–183.
- [Zimmerman et al., 1992] Zimmerman, B., Bandura, A., and Martinez-Pons, M. (1992). Self-motivation for academic attainment: The role of self-efficacy beliefs and personal goal setting. *American Educational Research Journal*, 29(3):663–676.
- [Zimmerman and Martinez-Pons, 1986] Zimmerman, B. and Martinez-Pons, M. (1986). Development of a structured interview for assessing student use of self-regulated learning strategies. *American Educational Research Journal*, 23(4):614–628.