Using Hidden Markov Models to Characterize Student Behaviors in Learning-by-Teaching Environments

Hogyeong Jeong[1], Amit Gupta[1], Rod Roscoe[1], John Wagster[1], Gautam Biswas[1], and Daniel Schwartz[2]

[1]Vanderbilt University & [2]Stanford University
Nashville, TN 37235. USA.

{hogyeong.jeong, rod.roscoe, john.wagster, gautam.biswas}@vanderbilt.edu; danls@stanford.edu

Abstract. Using hidden Markov models (HMMs) and traditional behavior analysis, we have examined the effect of metacognitive prompting on students' learning in the context of our computer-based learning-by-teaching environment. This paper discusses our analysis techniques, and presents evidence that HMMs can be used to effectively determine students' pattern of activities. The results indicate clear differences between different interventions, and links between students learning performance and their interactions with the system.

Keywords: Learning by Teaching environments, Metacognition, Behavior Analysis, hidden Markov modeling.

1 Introduction

We have developed exploratory learning environments called teachable agents that use a learning-by-teaching paradigm to promote learning and reasoning skills with middle school science students [1][2]. The students are typically not domain experts and have little knowledge of teaching practices. In these environments, students teach a computer agent called Betty using structured graphical representations called concept maps [3]. Since the concept maps are purported to be representations of Betty's knowledge, the students are teaching Betty and fixing her errors by revising the maps. Of course, the maps are generated by the students based on their own knowledge, thus they are actually representations of the students' own domain understanding (Fig. 1).

The teaching aspects of this environment build upon research showing that students can benefit academically by teaching other students [1][4]. Biswas, Schwartz, & Bransford have reported that students preparing to teach felt that the responsibility to teach encouraged them to gain deeper understanding of the materials [5]. Beyond preparing to teach, actual teaching taps into three critical aspects of learning interactions – *structuring, taking responsibility, and reflecting.* These interactions facilitate self-monitoring and reflective knowledge-building for the teacher [6]. Effective teaching requires the monitoring of how well students understand and use ideas. Tutors and teachers often reflect on their interactions with students during and after the teaching process in order to better prepare for future sessions [8][9].

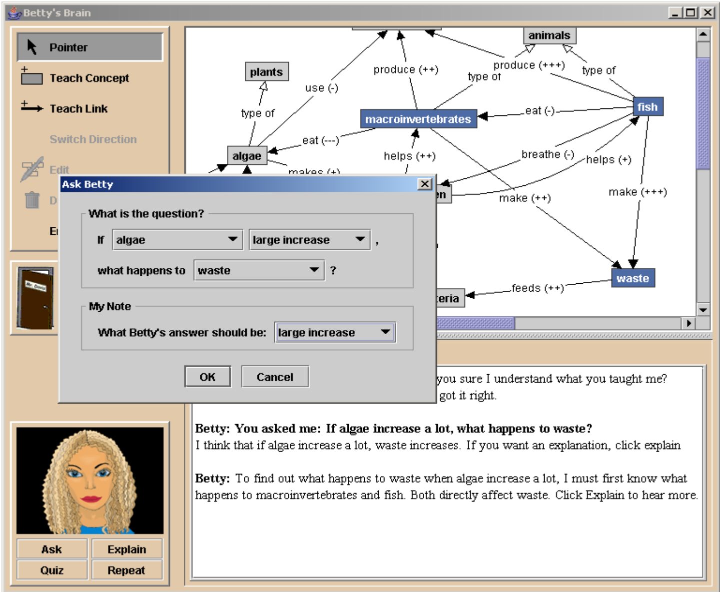

Fig. 1: Betty's Brain system with Query Window

The visual concept map structure also helps students make concepts and relationships explicit, which supports self-monitoring and knowledge organization [3]. The concept mapping also occurs in a context where the students can query Betty, ask her to explain her reasoning, and assess her knowledge by having her to take quizzes. For these reasons, we have hypothesized that working with Betty can help students to better understand science concepts, and engage in productive learning strategies that promote metacognition, organization, and reasoning with causal knowledge

Our previous work has focused on students' learning as measured by the quality of their concept maps. We found that learning-by-teaching with metacognitive support helped students learn about river ecosystems, and also better prepared them for future learning on related topics [1][11]. We compared several versions of 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, in which she exhibited self-regulated learning behaviors that the student teacher could appropriate to improve their own learning. 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.

We have recently turned our attention to analyses of students' behaviors as they teach Betty and create concept maps. Such analyses are important because they shed light on students' choices of interactive behaviors that influence learning, and the strategies they bring to the learning task [4]. Preliminary analyses of prior data showed that the quality of students' concept maps was paralleled by patterns in their behaviors [2]. These results suggest that self-regulated learning prompts and feedback from Betty helped student teachers engage in productive learning interactions.

In this paper, we discuss data from a new study testing the benefits of the Betty's Brain system. In particular, we present a refined methodology for exploring students' strategies using hidden Markov models (HMMs) to capture students' behaviors as they use the system [13]. Ours is a specific implementation of the generic process outlined by Fisher and Sanderson [14]. We discuss our methods for extracting the student interaction patterns from system log files, and describe our procedure for deriving and interpreting the HMMs of student behaviors. We then compare the HMMs across three experimental conditions in the main and transfer phases. We believe that this approach has merit for analyzing student behaviors for several reasons. First, HMMs allow us to go beyond frequency counts or proportions of individual behaviors, instead examining how these behaviors cohere in larger patterns or strategies. Similarly, this approach takes into account the entire sample of students' behaviors, rather than focusing only on specific behaviors or moments in time. The holistic nature of our analysis may provide a useful global view of how students approach the learning task.

2 Experimental Design and System Features

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. The conditions varied on the type of scaffolding provided by the mentor agent and/or the Betty agent. The students first created concept maps on river ecosystems during the main phase (seven 45-minute sessions). After an eight-week delay, students participated in the transfer phase (five 45-minute sessions) in which they learned about a new domain, the land-based nitrogen cycle. All students used an identical system during the transfer phase.

The three versions of the system were: (i) a learning by teaching (LBT) version in which students taught Betty, (ii) a self-regulated learning by teaching (SRL) version in which students taught Betty and received metacognitive prompts from Betty, and (iii) an intelligent coaching system (ICS) version in which students created a map for themselves with guidance from the mentor agent.

Students' interactions with Betty include three main activities: "teaching" by generating the concept map; "querying" by using a template to ask Betty questions; and "quizzing" Betty by asking set of predefined questions that have been "assigned" by the mentor agent. Betty answers questions using qualitative reasoning methods [1] to follow chains of links to determine how changes in one concept affect other concepts. After asking Betty a question, students can ask Betty to explain her reasoning steps.

The ICS version was our control condition. Students constructed a concept map to answer three sets of quiz questions. These students had access to the same teach, query, and quiz functions, but they were not presented in terms of teaching Betty. Students directly edited and queried their own maps. When students submitted their maps for the quizzes, Mr. Davis, the mentor agent, provided corrective feedback in the form of hints on how to correct errors [1]. The LBT group received the same corrective feedback from Mr. Davis after Betty took a quiz. The feedback to the SRL students focused on higher level concepts (e.g., read about the food chain or the waste cycle) and suggestions on how they could become better learners and teachers.

For the transfer phase, all students used a stripped down version of the LBT system with no feedback provided by Betty or the mentor. Students could still check their maps by asking questions or submitting them for a quiz.

2.1 Metacognitive Support in Betty's Brain

An important part of our system is the self-regulated learning support provided to the students. Self-regulated learning theory describes a set of comprehensive skills such as setting learning goals, selecting appropriate strategies, monitoring one's learning progress, and revising one's knowledge and strategies as necessary [14][18].

Table 1. Some Interactive Action Patterns and Betty's responses

Regulation Goal	Pattern Description	Betty Response
MONITORING BY ASKING QUERIES	**Successive quiz requests but no queries asked of Betty in between quizzes**	I'm still unsure of this material and I would like to do well. Mr. Davis said "take the quiz only if you think you will do well." *(Betty refuses to take quiz)*
MONITORING THROUGH EXPLANATIONS	**Multiple requests for Betty to give an answer but no request for explanation**	Let's see, you have asked me a lot of questions, but you have not asked for my explanations lately. Please make me explain my answers so you will know if I really understand.
TRACKING PROGRESS	**The most recent quiz score is significantly worse than the previous score**	I would really like to do better. Please check the resources, teach me, and make sure I understand by asking me questions that are on the quiz. My explanation will help you find out why I am making mistakes in my answers. Also, be sure to check out the new tips from Mr. Davis.

Betty's SRL persona incorporates aspects of this metacognitive knowledge that she conveys to the students to help them develop and apply monitoring and self regulation strategies [2]. For example, when the student is building the concept map, Betty occasionally responds by demonstrating reasoning through chains of events. She may 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 students to reflect on what they are teaching and perhaps 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.

We have identified several recurrent sequences where metacognitive feedback might be useful. When the system detects such patterns, Betty provides suggestions the students may employ to improve their own understanding. Some of the triggering patterns along with Betty's response are shown in Table 1. After Betty takes a quiz, the mentor agent also reminds Betty and her teacher about the importance of reading resources, and checking one's understanding after learning (teaching) new material.

3 Learning Results

In the main phase, the SRL condition generated maps with significantly more correct concepts and links than the LBT, $p < .05$, and ICS students, $p < .05$ [2]. These results suggest that the metacognitive prompting improved the students' learning. However, the LBT students also generated more correct maps than the ICS students, which suggests an overall benefit for learning by teaching.

Table 2. Concept map quality: main and transfer studies

Condition	Mean (SD) Map Scores	
	Main Phase	**Transfer Phase**
ICS	22.83 (5.3)	22.65 (13.7)
LBT	25.65 (6.5)^{c}	31.81 (12.0)
SRL	31.58 (6.6)^{a,b}	32.56 (9.9)^{a}

^{a} SRL > ICS, $p < .05$; ^{b} SRL > LBT, $p < .05$; ^{c} LBT > ICS, $p < .05$.

Students' transfer map scores provide indications whether a given version of the system better prepared students to learn in a new domain without scaffolds and prompts. Students in the SRL condition still had the highest map scores after the transfer phase, and scored significantly higher than the ICS students, $p < .05$. Interestingly, the LBT students' scores were now comparable to the SRL students. However, LBT students did not differ significantly from the ICS group, in part because of the high level of variability within that group.

Our interpretation is that working with Betty, especially with metacognitive prompts, helped students develop metacognitive strategies that supported their abilities to learn subsequently. However, another possible explanation is that ICS students were at a disadvantage because they switched from a non-teaching environment to a teaching environment in the transfer phase of the study, whereas the other students had not. By directly analyzing patterns of how students interact with the system, we explore the validity of these explanations.

4 Analysis of Behavior Patterns

We recorded log files of students' interactions with the system. From these log files, we identified the six main activities summarized in Table 3.

Sequences of these activities were then extracted from the log files, and mined using statistical learning methods to search for patterns that defined students' interactions with system. Our goal was to determine if there was evidence of different activity patterns between groups during the main phase, and whether these patterns persisted or changed when students worked in a new environment during the transfer phase. We thus derived two sets of HMMs, one set from students' accumulated activity sequences from all main phase sessions, and the second set from the transfer phase sessions [13].

Table 3. Student Activities and Related Actions

Activity	Student Actions
Edit Map (**EM**)	adding, modifying, or deleting concepts and links
Ask Query (**AQ**)	asking Betty queries
Request Quiz (**RQ**)	asking Betty to take the quiz
Resource Access (**RA**)	accessing the resources
Request Explanation (**RE**)	asking Betty for an explanation to her query answer
Continue Explanation (**CE**)	asking Betty to provide a more detailed explanation

HMMs are so named because their states are hidden. That is, they are not directly observed in the input sequences, but provide an aggregated description of the students' interactions with the system. Sequences of states may be interpreted as the students' learning behavior patterns. The set of parameters that define a HMM comprise (i) the transition probabilities between the states, (ii) observation probabilities for detecting a particular observation in a state, and (iii) initial probabilities for each state [13]. The particular learning method used, developed by Li and Biswas [19] utilizes the Bayesian information criterion (BIC) to find the optimal number of states that define the HMM.

The HMM models derived for the three conditions in the two phases of our study are summarized in Figure 3. For convenience, each hidden state is labeled by the predominant activity (or activities) comprising that state. (Only those activities whose likelihood of occurrence exceed 10% are listed). Some states are dominated by a single activity (e.g., editing the map), whereas others represent a composite of more than one activity (e.g., requesting quizzes and accessing the resources). Figure 3 also provides the likelihood, expressed as a percentage, of a student in a given state transitioning to a different state or remaining in the current state.

4.1 Interpreting the HMMs

Much like factor analysis, it is up to the researcher to give meaning to the derived interactive states, and hypothesize strategies for learning that may be associated with these states. Our analyses suggest several interpretable patterns that are relevant to interactive metacognition. These patterns combine several activity states and transitions to define higher level behavior patterns with links to metacognitive strategies.

One pattern is *basic map building*. This activity pattern is characterized by editing the map (EM), submitting the map for a quiz (RQ), and occasionally accessing the reading resources (RA). The pattern reflects a basic and important metacognitive strategy. Students work on their maps, check the map by taking a quiz to see if there are flaws, and occasionally refer to the readings.

A second pattern is *map probing*. Students edit the map (EM) and then ask a question (AQ) to check for specific relations between two concepts (e.g., if fish increase, what happens to algae?). This pattern exhibits a more proactive, conceptually driven strategy, because students are targeting specific relations rather than relying on the quiz to identify errors. Students also need to formulate their own questions to do so.

The third pattern is *map tracing*. This pattern reflects students asking Betty or the mentor (depending on the system) to explain the reasoning step by step (RE and CE). When Betty or the mentor initially answers a question, they state that a change in one entity causes a change in another entity and highlight the paths they followed to reach their answer. To follow the details of the inference chain, students had to ask Betty or Mr. Davis to explain their reasoning. The agents did so by hierarchically decomposing the chain of inference; for each explanation request, they showed how a particular path within the larger chain contributed to the final answer. Receiving more details about the reasoning process is particularly useful when maps become complex, and there are multiple paths between two concepts.

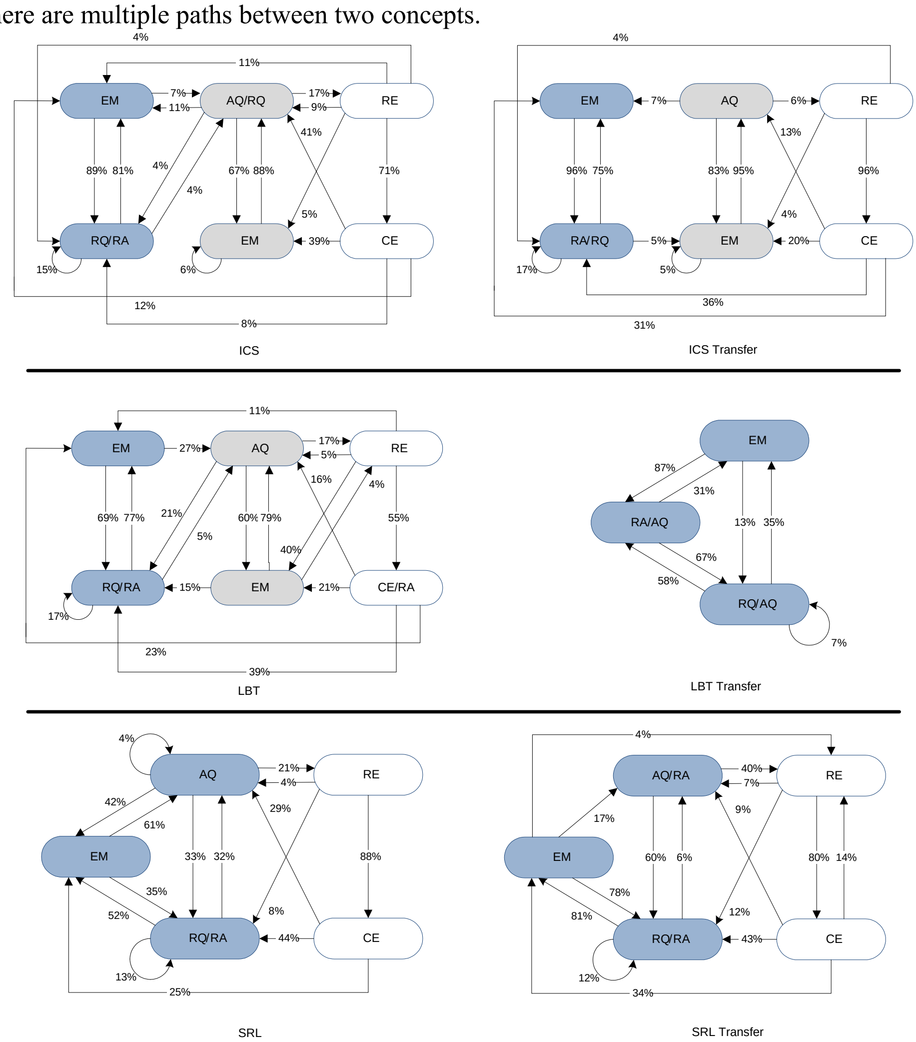

Fig. 3: HMMs for the three conditions in the main and transfer phases

The individual states portrayed in the original HMMs (Fig. 3) can be combined and re-represented in order to reflect these higher level aggregate states. These aggregate states are shown in Fig. 4, which are separated by condition and phase of the study. The percentages accompanying each arrow indicate the likelihood of transitioning from one aggregate to another or remaining in a given aggregate state. In addition, we exploited the stationary nature of these models to calculate the steady state probabilities of each aggregate state as the sum of the stationary probability of the individual states that make up the aggregate state (Table 4). These values indicate the probability that a student would be in a given state. For example, during the main phase, ICS students had a 62% chance of engaging in map building, but only a 7% chance of engaging in map tracing. The individual states' steady-state probabilities were also taken into account when calculating the transition probabilities between aggregate states.

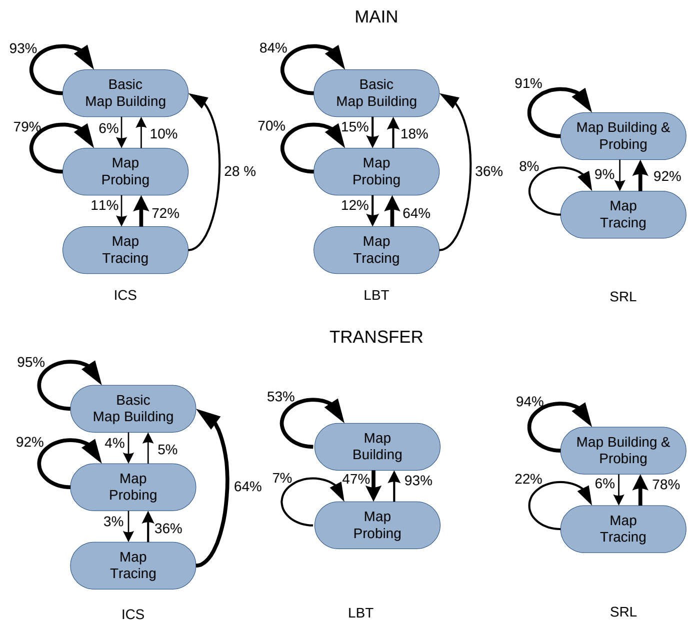

Figure. 4. Behavior Patterns for the three groups in the main and transfer study

Table 4. Aggregate state stationary probablities (i.e., probability of being in a given state).

State	ICS	LBT	SRL	Transfer ICS	Transfer LBT	Transfer SRL
Map Building	0.62	0.54	x	0.60	0.66	x
Map Probing	0.31	0.38	x	0.36	0.34	x
Map Tracing	0.07	0.08	0.15	0.04	x	0.11
Building & Probing	x	x	0.85	x	x	0.89

4.1.1 Main Phase Patterns. One way to approach each figure is to assume that students begin with basic map building. As an example, the ICS models show that there is a greater than 90% chance that the students remain in the map building condition, and less than a 10% chance that they transition to the map probing state. The LBT behavior model is very similar to the ICS model, except that these students in this group were more likely to transition to the map probing state from the map building state (15%). Both the ICS and LBT groups rarely use Map Tracing as a learning behavior.

The SRL behavior model is different in that the map building and map probing states are tightly coupled, and thus aggregated into one state. This is not surprising because Betty's prompts required the students to ask queries and check her answers between quizzes. The aggregated model indicates that the SRL students were more likely to engage in map tracing behaviors (15% as opposed to 8% for the LBT group and 7% for the ICS group) perhaps to understand how she reasoned with the concept map to derive her answers. Overall, the SRL condition exhibited a more versatile repertoire of interactive strategies for completing the cognitive task of teaching Betty. This offers one explanation for why the SRL students generated higher quality maps in the main phase, even though they were never explicitly told how to correct their maps. The support for interactive metacognition, primarily in terms of seeking information from the resources, and monitoring Betty's learning helped them learn the content better than the other two conditions.

4.1.2 Transfer Phase Patterns. In the transfer phase, all students taught Betty but all scaffolds were removed. The only feedback students received was how well Betty performed on the quizzes. The question was whether there was any continuation of the patterns developed during the main phase. The ICS condition continued to focus on the basic map building pattern (60%). Their map probing behavior occurrence increased marginally (31% to 36%), and their use of the tracing mechanisms was limited (4%), even though they were now teaching Betty just like the other two groups. We inferred that the teaching aspect alone did not override the students' desire to simply get quiz answers right. The ICS students did not seem inclined to probe Betty's understanding, and by extension their own understanding.

The transfer phase behavior patterns exhibited by the SRL group were also similar to their main phase behaviors. The map building and map probing states were still aggregated, and occurred with high frequency (89%). The transitions from the building/probing state to map tracing decreased (9% to 6%). It is possible that the SRL students had internalized the reasoning mechanism and did not need to probe Betty as often or ask her to explain. However, once these students transitioned to the map tracing state, there were more internal transitions in that state (transition likelihood was 22%). This may indicate that when the concept map and the answer generation process became complex, the students did spend more time in map tracing activities.

The LBT condition behavior model also remained similar with map building and map probing dominating their learning activities. However, a more careful study of the more detailed model in Figure 3 reveals that within the map building phase these students spent almost twice as much time reading the resources as they did in editing their maps (41% to 25%). The amount of map tracing by this group seemed to decrease and all map tracing activity was integrated with the map building and map probing states. This version of Betty used in the transfer phase was most similar to the original LBT condition, except there was no corrective feedback provided by Mr. Davis. Therefore, it is reasonable to expect that the LBT condition would show the same interactive patterns across study phases. Instead, the LBT students seemed to develop a different learning strategy that included more time spent in reading the resources to learn about the new domain during the map building phase.

Unlike the SRL group the LBT group did not receive feedback on monitoring strategies in the main phase of the study. As a result, the LBT group did not seem to use map tracing as a learning strategy. Instead, their strategy was to track down Betty's wrong answers by querying her and then reading the resources to find how to change their map to get the correct answer (given that they were no longer given corrective feedback). Teaching in the main phase of the study seemed to have a positive effect on the LBT groups learning behaviors because they performed better than the ICS group in the transfer phase where both groups worked with the same system. The differences in the feedback received during the main phase of the study also produced differences in learning behaviors between the LBT and SRL groups. Whereas the LBT group combined map probing and reading of resources to learn the new domain, the SRL group also used map tracing when faced with complex reasoning situations with their concept maps.

5 Discussion and Conclusions

The Betty's Brain system is designed to leverage the benefits of learning by teaching and concept mapping to facilitate students' science learning and causal reasoning. Our hypothesis that working with Betty helped students engage in educationally productive cognitive and metacognitive processes is supported by the results reported here. Students who utilized the LBT and SRL systems constructed better concept maps with more causal relationships between entities than students who used the non-teaching ICS version of the system. Moreover, students' performance was strongest when we explicitly supported their use of self-regulated learning strategies by having Betty model and prompt for such behaviors. Not only did these students do well in the main phase of our study when the prompts were present, they also continued to outperform other groups in the transfer phase when the prompts were absent.

Although assessments of learning outcomes were in agreement with our hypotheses, it was also critical to explore students' actual behaviors during the teaching and learning process. Using HMMs to characterize students' behaviors allowed us to identify several meaningful patterns that distinguished our three experimental conditions in the main and transfer phases.

Examining the map building, map probing, and map tracing patterns across the three conditions provided interesting results. First, these results support the claim that our SRL system with metacognitive prompting was beneficial because it altered students' behaviors in positive ways. Whereas LBT and ICS students relied heavily on basic map building, we were successful in encouraging SRL students to engage in more probing and tracing. In addition, these beneficial patterns tended to persist in the transfer phase of the study. An interesting result is the learning strategy that the LBT students developed as they progressed from the main to the transfer phase. Although these students used more map building during the main study, they spontaneously showed a shift toward probing and resource reading to correct errors, but they did not develop the tracing behavior during the transfer phase. Because the LBT students used fairly similar systems in both the main and transfer phases, these results suggest that use of a learning-by-teaching system over a period of time may help students gradually *develop* better learning strategies. But there is also added value to focusing on metacognitive and self-regulated learning strategies through social interactions between tutors, their students, and agents who play the role of mentors. More research is needed to look at the benefits of extended use of our system.

Acknowledgments. This work has been supported by a Dept. of Education IES grant #R305H060089 and NSF REESE Award #0633856.

References

1. Biswas, G., Leelawong, K., Schwartz, D., Vye, N., TAG-V: Learning By Teaching: A New Agent Paradigm for Educational Software. In: Applied Artificial Intelligence. AAI, vol. 19, pp. 363-392. (2005)
2. Wagster, J., Tan, J., Wu, Y., Biswas, G., Schwartz, D: Do Learning by Teaching Environments with Metacognitive Support Help Students Develop Better Learning Behaviors? The 29th Annual Meeting of the Cognitive Science Society, Nashville, TN, pp. 695-700. (2007)
3. Novak, J.D: Learning, Creating, and Using Knowledge: Concept Maps as Facilitative Tools in Schools and Corporations. LEA, Mahwah, NJ. (1998)
4. Roscoe, R. D., Chi, M: Understanding tutor learning: Knowledge-building and knowledge-telling in peer tutors' explanations and questions. In: Review of Educational Research. vol. 4, pp. 534-574. (2007)
5. Biswas, G., Schwartz, D., Bransford, J., TAG-V: Technology support for complex problem solving: From SAD environments to AI. In: K.D. Forbus and P.J. Feltovich (Eds.), Smart Machines in Education. AAAI Press, pp. 71-98. Menlo Park, CA (2001)
6. Roscoe, R. D., Chi, M: Tutor learning: The role of instructional explaining and responding to questions. In: Instructional Science. (in press)
8. Chi, M.T.H., Siler, S. A., Jeong, H., Yamauchi, T., Hausmann, R. G: Learning from Human Tutoring. In: Cognitive Science. vol. 25(4), pp. 471-533. (2001)
9. Lin, X., Schwartz D. L., Hatano G.: Toward Teachers' Adaptive Metacognition. In: Educational Psychologist. vol. 40(4), pp. 245-255. (2005)
11. Leelawong, K. and Biswas, G.: Designing Learning by Teaching Systems: The Betty's Brain System. In International Journal of Artificial Intelligence in Education. (2008).
13. Rabiner L. R.: A Tutorial on Hidden Markov Models and Selected Applications in Speech Recognition. In: Proc. IEEE. vol. 77. (1989)
14. Azevedo, R.: Using Hypermedia as a Metacognitive Tool for Enhancing Student Learning? The Role of Self-Regulated Learning. In: Educational Psychologist. vol. 40(4), pp. 199-209. (2005)
18. Zimmerman, B. J.: A Social Cognitive View of Self-Regulated Academic Learning. In: Journal of Educational Psychology. vol. 81(3), pp. 329-339. (1989)
19. Li, C. and Biswas, G.: A Bayesian Approach for Learning Hidden Markov Models from Data. Special issue on Markov Chain and Hidden Markov Models, Scientific Programming, vol. 10, pp. 201-219. (2002).
20. Fisher, C. and Sanderson, P.: Exploratory sequential data analysis: exploring continuous observational data. In: Interactions 3(2): pp. 25-34 (1996).