Modeling and Measuring Self-Regulated Learning in Teachable Agent Environments

John S. Kinnebrew, Gautam Biswas

Vanderbilt University (USA)
john.s.kinnebrew@vanderbilt.edu, gautam.biswas@vanderbilt.edu

Our learning-by-teaching environment has students take on the role and responsibilities of a teacher to a virtual student named Betty. The environment is designed to help students learn and understand science topics for themselves as they teach and monitor their agent. This process is supported by adaptive scaffolding and feedback through interactions with the teachable agent and a mentor agent. This paper discusses the results of a comparative study conducted in an 8th-grade science classroom, where students received two kinds of metacognitive and learning strategy feedback. We analyze student performance and learning gains as a result of the intervention. To gain further insight into student learning behaviors exhibited during the intervention, we employ a data mining methodology incorporating hidden Markov modeling and sequence mining techniques. The results illustrate both the effectiveness of the experimental agent feedback in encouraging metacognitive learning strategies and the utility of the data

for citations:
Kinnebrew J. S., Biswas G. (2011), *Modeling and Measuring Self-Regulated Learning in Teachable Agent Environments,* Journal of e-Learning and Knowledge Society, English Edition, v.7, n.2, 19-35. ISSN: 1826-6223, e-ISSN:1971-8829

mining methodology.

1 Introduction

Cognitive scientists have established that metacognition and self-regulation are important components for developing effective learning in the classroom and beyond (Pintrich, 2000, Zimmerman, 2001). Brown and Palincsar (1989) have demonstrated that younger students can acquire and apply metacognitive skills, such as planning and monitoring, through instruction. However, students in typical classrooms are rarely provided opportunities to learn and exercise these strategies (Paris & Paris, 2001).

Our research team has been developing computer-based-learning environments that utilize the learning-by-teaching approach in order to foster student 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 science and mathematics domains (Biswas *et al.*, 2005).

This paper discusses the results of a study conducted in an 8th-grade science classroom in which students taught their agent about global climate change. One of our goals was to determine the degree to which the agents' metacognitive and SRL prompts could help improve student learning. In particular, we augmented the existing agent feedback, which promoted metacognitive awareness, with more explicitly strategy-oriented feedback from the mentor agent. We report results on student learning gains from pre- to post-tests and the quality of the maps that they created during the intervention. Further, we apply a novel method to describe student learning behaviors and compare them between groups. This methodology employs a variety of techniques, including constructing hidden Markov models (HMMs) from student activity traces and applying sequence mining methods to develop more refined interpretations of the students' learning behaviors. The results of this analysis illustrate important differences in learning behaviors between two student groups receiving either the metacognitive-awareness or the strategy-augmented feedback.

2 Related Work in Measuring Self-Regulated Learning

Recently, researchers have begun to utilize trace methodologies in order to examine the complex temporal patterns of self-regulated learning (Aleven *et al.*, 2006; Azevedo & Witherspoon, 2009; Biswas *et al.*, 2010; Hadwin *et al.*, 2007; Jeong & Biswas, 2008; Zimmerman, 2008). Underlying these approaches is a move away from assessing self-regulation as an intrinsic aptitude and, instead, assessing it as dynamic and adaptive event occurrences. By identifying and analyzing temporal event sequences from trace data (e.g., student actions in a computer-based learning environment), we hope to better understand student's self-regulated learning (SRL) strategies and develop online measurement schemes to provide more adaptive feedback and scaffolding.

In recent years, a number of researchers have had success in applying HMM generation techniques to this learning interaction data. Pardos and Heffernan (2011) have successfully applied HMMs (in conjunction with a bagged decision tree classifier) to predict correctness of student answers in computer math learning environments. Boyer *et al.* (2009) have used HMMs for identifying strategies and high-level flow in student dialogues with human tutors. Some of the more closely related work with HMMs for identification of student learning strategies is the HMM clustering technique employed in (Shih *et al.*, 2010). Performance and learning metrics are used to direct generation of HMM collections that predict learning outcomes. In contrast, our technique generates HMMs from a set of student activity sequences to build a model that represents an aggregated view of the group's learning behaviors (Biswas *et al.*, 2010).

To identify additional details of students' learning behaviors, we employ sequence mining techniques. Sequential pattern mining (Agrawal & Srikant, 1995) is used to identify frequent patterns of actions within a group of students. However, this can also result in a very large number of frequent patterns, presenting difficulties for effective analysis. To find the most important patterns in a comparative analysis between student groups, our methodology combines sequential pattern mining with episode discovery/mining to identify differentially frequent patterns. In contrast to sequential pattern mining, episode mining (Mannila *et al*., 1997) defines frequency as the number of times the pattern can be matched within a single sequence of actions. Combining sequential pattern mining and episode mining, our methodology focuses on learning activity patterns that are used differentially between two groups of students.

3 Betty's Brain and Self-Regulated Learning Feedback

The Betty's Brain system 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; Schwartz *et al.*, 2009). Students interact with a Teachable Agent (TA), named Betty, and a Mentor Agent, named Mr. Davis, to learn and understand a science topic. The system includes a set of indexed, hypermedia resources that students can access and use at any time while working on the system. Using the visual interface shown in Figure 1, students explicitly teach the Betty agent by constructing a causal concept map representation (Leelawong & Biswas, 2008).

The system supports five primary types of activities:

1. Read: students access one of the pages in the resources;
2. Edit: students add, delete, or modify a concept or link in the map;
3. Query: students use a template, illustrated in Figure 1, to check their teaching by querying Betty;
4. Explain: students probe Betty's reasoning by asking her to explain her answer to a query;
5. Quiz: students assess how well they have taught Betty by having her take a quiz, which is a set of questions chosen and graded by the mentor agent.

Since our middle school students are novices in both the science topics and the teaching tasks, the agents provide them with a variety of scaffolds to help overcome obstacles they may face in learning and teaching the domain material. For example, Betty answers student queries using qualitative reasoning methods through chains of links (Leelawong & Biswas, 2008). If asked, she also explains her reasoning through text, speech, and animation schemes. Students reflect on Betty's answers and her explanations, potentially revising their own knowledge as they make changes to the concept maps to better teach Betty. When students feel they are not making progress, they can seek help from the mentor agent by clicking on an "Ask Mr. Davis" button.

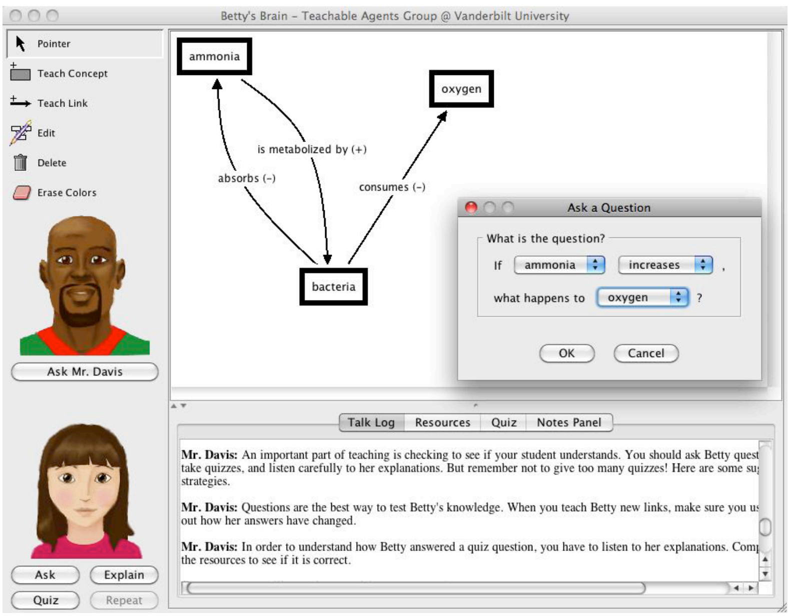

Fig. 1 - Betty's Brain system with query window.

In particular, the agent feedback and dialogues focus on metacognitive awareness and strategies related to knowledge construction and monitoring (Biswas *et al.*, 2010). For knowledge construction (i.e., using the resources to acquire knowledge and structure it in a causal map), the feedback addresses 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 causal concept maps. The agent feedback also addresses 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 following the causal reasoning steps generated by the agent to locate potential errors in the maps. Proper guidance (i.e., relevant and timely feedback) provides opportunities to help the students develop better learning strategies.

3.1 Classroom Study

We recently performed a study with 49 8th-grade students in a middle Tennessee science classroom. At the beginning of the study, students were introduced to the science topic (global climate change) by the classroom teacher, and then provided an overview of causal relations and concept maps with hands-on system training by the researchers. For the next five days, students taught their agent about climate change and received feedback from both agents with minimal intervention by the teachers and the researchers.

In both the control and experimental conditions, students received feedback on how much their agent had learned by asking her questions and getting her to take quizzes that were graded by the mentor agent. In both conditions, students also received a variety of feedback from the agents to promote metacognitive awareness and the use of metacognitive strategies during learning. For example, as the students taught Betty, occasionally she would remark that the relation implied by a link between two concepts did not make sense to her, and that the students should check to see if they were teaching her correctly (i.e., the student should monitor their own learning). Mr. Davis similarly provided spontaneous metacognitive awareness feedback, which was triggered by the students' recent activities or the state of their causal map (Biswas *et al.*, 2010).

The difference between the two conditions was in the additional feedback providing explicit advice on learning strategies that students received from the mentor in the experimental condition. In the control condition the mentor feedback was quite general, and provided as short statements, for example, "You should check if Betty's explanations are correct by comparing them to the information in the resources." The purpose of this feedback was to provide cues to help students think of strategies to improve their own learning processes. In the experimental condition, this feedback was supplemented with extended feedback, which suggested specific actions to the student in the context of explicit strategies. For example, occasionally when Betty got a quiz question wrong due to an incorrect link, Mr. Davis would say, "Betty got this question wrong because she does not have a good understanding of how these concepts are related. You might want to read the resources related to the links she used to see if any of the relationships are wrong," which is followed by a list of specific pages in the resources to read about the relevant concepts and relationships. Further the timing and presentation of the experimental feedback was designed to be more context-relevant. For example, Mr. Davis kept track of the part of the map that had a number of errors, and when the student seemed to be working on this part of the map, would point to resources the student should read to find out more about concepts and links from that part of the map.

4 Learning Behavior Analysis

We analyzed student learning in the Betty's Brain system by extending standard measures of learning gain and task performance with an exploratory data mining methodology that combines HMM generation and sequence mining techniques. In this section, we describe the major components of our data mining methodology and present the results of its application to learning traces from the classroom study described in Section 3.1.

4.1 Overall Learning and Performance

For this study, we employ two measures to assess learning gain and task performance: (i) normalized gain in scores from pre-test to post-test[1], and (ii) final map score[2], which provides a measure of overall performance in teaching the agent. The pre- and post-tests included two kinds of multiple-choice questions: (1) 11 science definition questions, which tested students' understanding of primary concepts and simple relations among concepts; and (2) 16 causal-reasoning questions, where students were given a simple causal map, and asked to answer questions about chains of cause-and-effect in the map. Because the

[1] Normalized learning gain was calculated as: (post - pre)/(max - pre)

[2] Map score was computed as the number of correct links (in comparison to the expert map) minus the number of incorrect links. The maximum possible map score was 18.

students were assigned to the control or experimental condition by section instead of academic performance, we control for these differences between groups using the Tennessee Comprehensive Assessment Program (TCAP) test scores as a measure of prior academic achievement. Table 1 presents the mean difference[3] between the experimental and control groups in the study.

TABLE 1
MEAN DIFFERENCE (EXPERIMENTAL - CONTROL)
WITH TCAP AS A COVARIATE

Metric	Mean Difference	Significance (p)	95% CI
Definition Learning Gain	-0,200	0,291	[-0,580,0,179]
Causal Learning Gain	1,015	0,002	[0,398,1,632]
Total Learning Gain	0,301	0,085	[-0,044,0,645]
Map Score	7,956	0,000	[3,958,11,954]

The results in Table 1 illustrate statistically significant differences in final map scores and casual reasoning learning gain between the control and experimental groups. For both of these metrics, the experimental group outperformed the control group, indicating that the experimental feedback helped students better understand causal reasoning, as well as build better causal maps. Since both groups had access to the same resources and were involved in the same map building task, the lack of a significant difference in learning gain for the science definition questions is not surprising. Overall, these results illustrate the effectiveness of the explicit strategy-oriented feedback provided in a context-relevant manner. To analyze how the differences in feedback may have affected student learning behaviors and strategies, we employ an exploratory data mining methodology on the students' learning activity sequences in the following sections.

4.2 Modeling Learning Behaviors with HMMs

The student activities logged by a learning environment result from a variety of internal cognitive/metacognitive states, strategies, and processes used by the student. Employing a direct representation of these internal states and strategies with a probabilistic automaton, such as a hidden Markov model (HMM) (Rabiner, 1989), has the potential for facilitating identification, interpretation, and comparison of student learning behaviors.

Like a student's mental processes, the states of an HMM are hidden, mea-

3 The difference in means is computed using estimated marginal means from a multivariate linear model with TCAP score as the covariate.

ning that they cannot be directly observed, but they produce observable emissions, such as actions in a learning environment. Together, three sets of probabilities form a complete model: (1) transition probabilities, which determine the likelihood of going from one state to another at each step; (2) state output probabilities, which define the likelihood of observing different outputs in each state; and (3) state initial probabilities, which define the likelihood that a state will be the starting state for an output sequence.

Before we can employ HMM and other data mining techniques to the student interaction traces, we must process them to convert raw log events into categorical actions. We categorize relevant events into the five available action types detailed in Section 3. To maintain a balance between minimizing the number of distinct actions in the sequences and keeping important context information, we employ a measure of the relevance of each action to recent, previous actions. Based on this relevance metric we split each categorized action into two distinct actions: (1) high relevance and (2) low relevance to recent actions (Biswas *et al.*, 2010).

The processed interaction traces can then be used as input to our HMM technique. Algorithms for learning an HMM from sequences are well-known but require appropriate configuration/initialization parameters for effective use (Rabiner, 1989). We have developed an algorithm that addresses these concerns in the construction of HMMs from a set of student activity sequences (Jeong & Biswas, 2008, Biswas *et al.*, 2010):

1. Model Initialization: To determine the appropriate number of states for the HMM, our algorithm employs the Bayesian information criterion (BIC) (Heckerman *et al.*, 1995), which balances a preference for concise models (i.e., fewer states) with a preference for better-fitting models (i.e., a greater likelihood of the model producing the observed activity sequences) (Li & Biswas, 2002). For each group of student interaction traces, we also cluster vectors of student activities (from each step in the student activity sequences), to find similar sets of activities at different points in the sequences (Biswas *et al.*, 2010). This provides an initial model that helps explore a targeted portion of the possible model space.
2. Model Generation: We employ the popular Baum-Welch algorithm (Baum *et al.*, 1970) to the initial model and interaction traces to derive the optimal group HMM.
3. State Occurrence Calculation: To determine the prevalence of individual states suggested by a generated HMM, we calculate the proportion of expected state occurrences (Biswas *et al.*, 2010). This metric employs the generated HMMs to calculate an expected value for the proportion of individual state occurrences. To maintain its relevance to the trace

data and prevent the marginalization of states only occurring near the beginning of learning activities (as can happen with a stationary probability calculation), the expected value is calculated with lengths of student interaction traces in the group (Biswas *et al.*, 2010).

4. Model Interpretation and Comparison: This is the most difficult and subjective step in the process of analyzing student activity data through HMM generation. We assign meaning to the derived states of the model and generate behavior descriptions in terms of the interpreted states, transitions, and the proportion of expected state occurrences.

We employed this HMM generation and analysis to identify learning behaviors in the control and experimental groups. Based on the resulting state action emission probabilities in Figure 3, and using the methods detailed in (Biswas *et al.*, 2010), we interpreted the states as representing:

- *Reading*: students are primarily engaged in reading the resources.
- *Informed editing*: students are primarily making informed edits related to recent activities, such as editing one area of the map at a time and linking their editing to recent reading, querying, and quizzing.
- *Uninformed editing*: students are primarily making unfocused or uninformed changes to their map, possibly indicating the use of trial-and-error and guessing strategies.
- *Checking*: students are querying and quizzing Betty to check the correctness of their causal maps and are making some changes to their map. However, their attempts at checking and editing their map are relatively unfocused and uninformed. This state likely corresponds to less effective attempts at employing monitoring strategies prompted by some of the Mentor agent feedback.
- *Uninformed editing and checking*: students are performing checking behaviors like querying and quizzing and are also making a significant number of uninformed edits to their map. This state (in the control group HMM) is similar to a combination of the uninformed editing and checking states (in the experimental group HMM).
- *Probing*: students are using focused queries and quizzes to check the correctness of their causal maps and are making informed changes to their map. This state likely corresponds to more effective attempts at employing monitoring strategies prompted by some of the Mentor agent feedback.

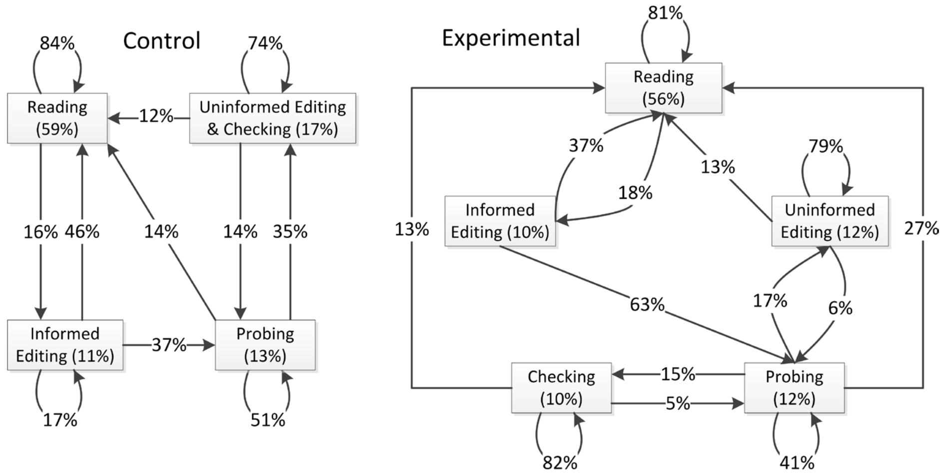

Fig. 2 - Control and experimental group HMM structures.

The HMM results illustrate a similar set of behaviors employed by both the control and experimental groups, although the uninformed editing and checking behaviors were combined in a single state for the control group. The proportions of expected state occurrences are also relatively similar between the two groups. However, there are distinct differences in the patterns in which these behaviors were employed, illustrated by the transition probabilities between states in Figure 2. The control HMM has a much smaller likelihood (37%) of transitioning from informed editing to probing activities compared to the experimental HMM (63% transition probability). Further, the experimental HMM exhibits a higher likelihood of following probing activities with more reading (27% versus 14% for the control HMM). Although the groups performed similar total amounts of probing (13% and 12% expected state occurrences for the control and experimental HMMs, respectively), as well as similar amounts of reading and informed editing, the transition probabilities suggest that the experimental group performed these activities in a more systematic fashion. To provide a complementary, finer-grained level of analysis, these HMM results are augmented with a comparison of frequent learning activity patterns in the following section.

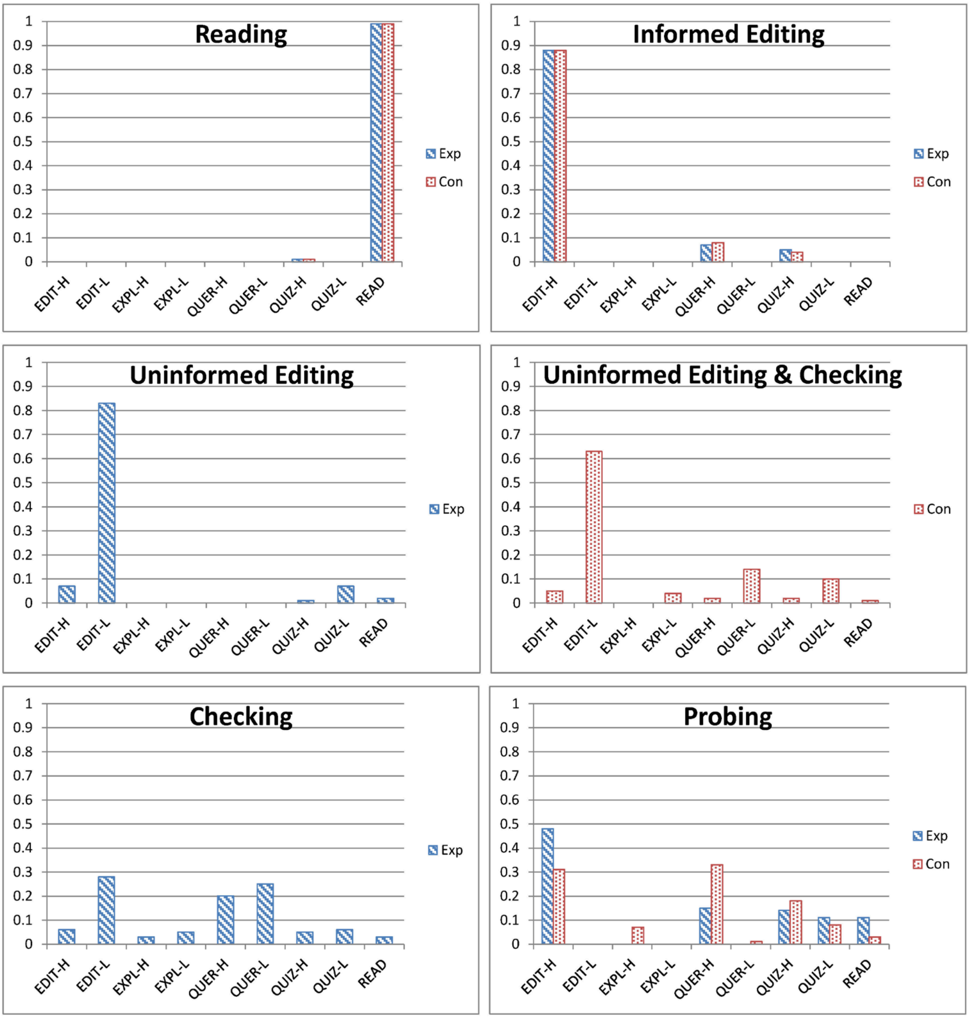

Fig. 3 - Learning activity emissions in HMM states.

4.3 Comparative Sequence Mining

The control and experimental group HMMs in the previous section provide an overview of student learning behaviors and illustrate some important differences between the two groups. However, these models of learning behavior lack detail for some analyses. This issue becomes even more pronounced when trying to compare across groups of students. When the HMMs generated for two groups contain similar states, it is still possible that those states represent different, although often related, learning strategies. To provide a finer-grained, complementary level of analysis, our methodology employs a novel combination of sequence mining techniques (Kinnebrew & Biswas, In Review).

For this analysis, the sequential pattern mining frequency measure (i.e., how many students exhibit the given pattern, called "sequence frequency/support" or s-frequency) is employed to identify patterns common to a group of students. We employ an s-frequency threshold of 50% to analyze patterns that were evident in the majority of either group of students. To quantitatively compare patterns across groups, we employ the episode frequency (i.e., the frequency with which the pattern is repeated within an interaction trace). For a given trace, we refer to this as the "instance frequency/support" or i-frequency. To calculate the i-support of a pattern in a group of traces, we use the mean of the pattern's i-support values across all traces in the group. This methodology combines the group s-frequency and i-frequency measures to identify differentially frequent patterns across two groups of interaction traces (Kinnebrew & Biswas, in review).

Applying this comparison to the control and experimental groups in the study, we immediately noticed that many of the frequent patterns differed only by the number of consecutive reads or edits in the pattern. To improve this exploratory analysis, we revised the log pre-processing to distinguish a single action from repeated actions, which were condensed to a single "action" with the "-MULT" identifier. Using the re-transformed sequences, our comparative sequence mining technique allowed us to identify a trend that was not apparent from separately considering s-frequent or i-frequent patterns, which were largely the same between the two groups. Table 2 presents the top three differentially frequent patterns in four comparison categories (i.e., categorized by whether the patterns were s-frequent in one or both groups and in which group they were more i-frequent).

TABLE 2
TOP DIFFERENTIALLY FREQUENT PATTERNS

Learning Activity Pattern	i-Support (Hi - Lo)	s-Frequent Group	Group Trend
EDIT-H-MULT -> READ	6.16	Experimental	Experimental group has greater tendency to use single reads & edits in read-edit sequences
READ-MULT -> EDIT-H-MULT -> READ	4.47	Experimental	
EDIT-L -> READ	4.00	Experimental	
EDIT-H -> READ	2.47	Both	
QUIZ-H -> EDIT-H	1.72	Both	
READ-MULT -> EDIT-H	1.58	Both	

Learning Activity Pattern	i-Support (Hi - Lo)	s-Frequent Group	Group Trend
READ-MULT -> QUER-H	-0.80	Both	Control group has greater tendency to use short patterns including queries
QUER-H -> QUIZ-H	-0.95	Both	
EDIT-L -> QUER-H	-1.51	Both	
QUER-H -> QUER-H	-2.16	Control	
EDIT-L -> EDIT-L	-2.16	Control	
QUER-L -> EDIT-L	-2.21	Control	

This analysis illustrates that many of the patterns that were differentially frequent in the experimental group were repeated read-edit patterns. Many similar sequences were frequent in the control group, but the differentially frequent sequences for the experimental group often included single reads and edits, while the control group relied more on multiple reads and edits. There are a number of possible explanations for the observed behavior differences between the groups. For example, the experimental group students may have employed a knowledge construction strategy of reading a small portion of the resources and adding a link to their map each time they identified a causal relationship. Conversely, the control group students may have engaged in less systematic knowledge construction strategies. In particular, they may have been less focused and encountered more difficulties during extended series of reads and edits.

Some differences in knowledge construction strategies between the two groups may also be related to differences in how they combined monitoring activities with their reading and editing. Therefore, we extended our comparative sequence mining analysis by using a regular expression constraint to focus on monitoring patterns. In particular, we identified interesting results for patterns including Query actions, which are important for effective monitoring of the knowledge construction and map building process. Table 3 presents the top three actions in each comparative category. This analysis illustrates that the experimental group tended to use queries before and after read-edit sequences. In contrast, the control group had a differential preference for using queries before and after individual (uninformed) edits. This provides further confirmation of the HMM results that illustrated a difference in transitions to and from the probing states in these groups. Overall, these results suggest that the experimental group employed monitoring activities, such as queries and explanations, more systematically in combination with short sequences of reading and editing. Further, the query-read patterns exhibited by the experimental group suggests that they may have used monitoring to identify weak points for further exploration with the resources, as well as to confirm their current understanding.

TABLE 3
TOP DIFFERENTIALLY FREQUENT PATTERNS INCLUDING QUERIES

Learning Activity Pattern	i-Support (Hi - Lo)	s-Frequent Group	Group Trend
QUER-L -> READ-MULT -> EDIT-H	1.19	Experimental	Experimental group has greater tendency to use queries in between read-edit sequences
QUIZ-H -> EDIT-H -> QUER-H	0.97	Experimental	
QUER-H -> READ-MULT -> EDIT-H-MULT	0.97	Experimental	
READ-MULT -> EDIT-H -> QUER-H -> READ-MULT	0.50	Both	
EDIT-H -> EDIT-H -> QUER-H	0.50	Both	
QUER-H -> READ-MULT	0.44	Both	
READ-MULT -> QUER-H	-0.80	Both	Control group has greater tendency to use queries before and after uninformed edits
QUER-H -> QUIZ-H	-0.95	Both	
EDIT-L -> QUER-H	-1.51	Both	
QUER-H -> EXPL-H	-2.11	Control	
QUER-H -> QUER-H	-2.16	Control	
QUER-L -> EDIT-L	-2.21	Control	

Conclusion

In this paper, we presented a comparative study conducted in an 8th-grade science classroom, where students received two kinds of metacognitive and learning strategy feedback. In order to better understand how the agents' metacognitive and SRL prompts could help improve student learning, we employed two types of metacognitive feedback/dialogues: (1) spontaneous feedback from both agents using general statements and questions to promote metacognitive awareness and the use of metacognitive strategies, and (2) context-relevant advice from the mentor agent suggesting specific actions in terms of explicit strategies. The control group received only the first, metacognitive awareness, feedback, while the experimental group received both kinds of feedback.

The analysis of student learning gains and performance showed that the experimental feedback helped students better understand causal reasoning, as well as build better causal maps. To gain further insight into how the explicitly strategy-oriented feedback affected learning behavior, we presented and applied an exploratory data mining methodology incorporating hidden Markov modeling and sequence mining techniques. The comparative learning behavior results suggested that students in the experimental group combined knowledge construction activities with metacognitive monitoring activities in a more systematic and effective fashion than the control group. In particular, the experimental group more often used single reads and edits in conjunctions with queries to monitor their understanding, which were followed up with further reading (e.g., possibly exploring weak points in their understanding by re-reading portions of the resources). Overall, these results illustrate both the effectiveness of the explicitly strategy-oriented agent feedback in encouraging systematic use of metacognitive learning strategies and the utility of the data mining methodology for analysis of learning behaviors.

In future work, we intend to expand upon this research through a variety of enhancements to both the agent feedback and the data mining methodology. We will increase the length of possible dialogues with the agents to allow more targeted and detailed metacognitive feedback. Further, we are working on creating a library of interaction trace segments that are representative of identified learning strategies for use in online HMM analysis of student activities. Online analysis to categorize student activities by comparison with this HMM library can facilitate more targeted scaffolding and feedback in the Betty's Brain system.

Acknowledgements

This work has been supported by U.S. NSF IIS Award #0904387.

REFERENCES

Agrawal R., Srikant, R. (1995), *Mining sequential patterns*. In Proceedings of the Eleventh IEEE International Conference on Data Engineering (ICDE), pages 3–14.

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

Azevedo R., 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.

Baum L., Petrie T., Soules G., Weiss N. (1970), *A maximization technique occurring in the statistical analysis of probabilistic functions of Markov chains*. Annals of Mathematical Statistics, 41(1):164–171.

Biswas G., Jeong H., Kinnebrew J., Sulcer B., Roscoe R. (2010), M*easuring Self-regulated Learning Skills through Social Interactions in a Teachable Agent Environment*. Research and Practice in Technology-Enhanced Learning (RPTEL), 5(2):123–152.

Biswas G., Leelawong K., Schwartz D., Vye N., Vanderbilt T. (2005), *Learning by teaching: A new agent paradigm for educational software*. Applied Artificial Intelligence, 19(3):363–392.

Boyer K., Ha E., Wallis M., Phillips R., Vouk M., Lester J. (2009), *Discovering Tutorial Dialogue Strategies with Hidden Markov Models*. In Proceedings of the 14th International Conference on Artificial Intelligence in Education, Brighton, UK, pages 141–148. IOS Press.

Brown A., 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, Hills-dale, NJ.

Hadwin A., Nesbit J., Jamieson-Noel D., Code J., Winne P. (2007), *Examining trace data to explore self-regulated learning. Metacognition and Learning*, 2(2):107–124.

Heckerman D., Geiger D., Chickering D. M. (1995), *A tutorial on learning with Bayesian networks*. Machine Learning, 20:197–243.

Jeong H., 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.

Kinnebrew J., Biswas G. (2011), *Comparative Action Sequence Analysis with Hidden Markov Models and Sequence Mining*. In Under Review for The 4th International Conference on Educational Data Mining (EDM).

Leelawong K., Biswas G. (2008), *Designing learning by teaching agents: The Betty's Brain system*. International Journal of Artificial Intelligence in Education, 18(3):181–208.

Li C., 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.

Mannila H., Toivonen H., Inkeri Verkamo A. (1997), *Discovery of frequent episodes in event sequences*. Data Mining and Knowledge Discovery, 1(3):259–289.

Pardos Z., Heffernan N. (2011), *Using HMMs and bagged decision trees to leverage rich features of user and skill from an intelligent tutoring system dataset*. To Appear in the Journal of Machine Learning Research (JMLR).

Paris S., Paris A. (2001), *Classroom applications of research on self-regulated learning*. Educational Psychologist, 36(2):89–101.

Pintrich P. (2000), *An Achievement Goal Theory Perspective on Issues in Motivation Terminology*, Theory, and Research. Contemporary Educational Psychology, 25(1):92–104.

Rabiner L. (1989), *A tutorial on hidden Markov models and selected applications in speech recognition*. Proceedings of the IEEE, 77(2):257–286.

Schwartz D., Chase C., Chin D., Oppezzo M., Kwong H., Okita S., Biswas G., Roscoe R., Jeong H., Wagster J. (2009). *Interactive metacognition: Moni-toring and* *regulating a teachable agent*. In Hacker, D., Dunlosky, J., and Graesser, A., editors, Handbook of metacognition in education, pages 340–358. Routledge Press.

Shih B., Koedinger K., Scheines R. (2010), *Unsupervised Discovery of Student Learning Tactics.* In Proceedings of the Third International Conference on Educational Data Mining.

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 B. (2008), *Investigating self-regulation and motivation: Historical background, methodological developments, and future prospects*. American Educational Research Journal, 45(1):166–183.