

Modeling Learner's Cognitive and Metacognitive Strategies in an Open-Ended Learning Environment

James R. Segedy, John S. Kinnebrew, Gautam Biswas

Vanderbilt University

{james.segedy, john.kinnebrew, gautam.biswas}@vanderbilt.edu

Abstract

The Betty's Brain computer-based learning system provides an open-ended and choice-rich environment for science learning. Using the learning-by-teaching paradigm paired with feedback and support provided by two pedagogical agents, the system also promotes the development of self-regulated learning strategies to support preparation for future learning. We apply metacognitive learning theories and experiential analysis to interpret the results from previous classroom studies. We propose an integrated cognitive and metacognitive model for effective, self-regulated student learning in the Betty's Brain environment, and then apply this model to interpret and analyze common suboptimal learning strategies students apply during their learning. This comparison is used to derive feedback for helping learners overcome these difficulties and adopt more effective strategies for regulating their learning. Preliminary results demonstrate that students who were responsive to the feedback had better learning performance.

Introduction

In recent years, the Betty's Brain Learning Environment (Biswas *et al.*, 2005) has been shown to add value to course instruction (Chin *et al.*, 2010) and aid learners in developing cognitive and metacognitive strategies to support science learning (Biswas *et al.*, 2010). In Betty's Brain, learners are given the task of teaching a computer agent (also referred to as a Teachable Agent, or TA for short) by acquiring and organizing science knowledge from hypermedia resources using a causal concept map representation (Leelawong & Biswas, 2008). As they teach their TA, they may ask her questions or have her take a quiz to assess her learning.

Betty's Brain is open-ended and choice-rich: learners must decide when and how to acquire information, build or modify their concept map, check their TA's progress, reflect on their own understanding of both the science knowledge and the evolving causal map structure, and seek help or feedback from a Mentor agent (also referred to as

the mentor). Equally important is how the students sequence these activities to make their learning effective and efficient.

One important goal of working with learners in such an open-ended environment is to develop their abilities to independently regulate their learning process, a notion that supports *preparation for future learning* (Bransford and Schwartz, 1999). Cognitive science researchers have established that this preparation through development of metacognitive and self-regulatory skills is important for developing effective learners in the classroom and beyond (*e.g.*, Zimmerman, 2001). However, previous studies have shown that middle school students do not have well-developed independent learning strategies; novices often adopt suboptimal trial and error methods when they encounter difficult, open-ended, and exploratory learning tasks (Azevedo, 2005; Bandura, 1986; Schunk & Zimmerman 1997). Empirical data demonstrates that adaptive scaffolding and feedback, which address both the domain content and metacognitive strategies, can enhance learning and prepare students for future learning (*e.g.*, Azevedo & Witherspoon, 2009; Biswas *et al.*, 2010; Schwartz *et al.*, 2009).

When students reach a decision point in teaching their agent, they must determine their next step using the knowledge and strategies they possess (note that seeking help is also an important strategy (Roll, *et al.*, 2007)). Although both the TA and the mentor provide scaffolding and feedback to promote metacognitive strategies and science understanding, Betty's Brain is distinguished from many other intelligent learning environments because *learners are never required to make a specific, "correct" next step in order to succeed*. Instead, the system provides scaffolding and support to encourage students to adopt "good" self-regulated cognitive and metacognitive learning strategies and thereby become more independent learners. Our hypothesis is that if students learn to adopt good self-regulation strategies, they are more likely to develop the

ability to make beneficial choices on their own, not just for the present, but also for future learning.

Properly-designed support requires understanding how students make learning choices as they perform the various knowledge-building and monitoring activities required in the system. In our recent work, we have explored learner choices by analyzing traces of their learning activities on the system using sequence mining techniques to discover frequent behavior patterns (expected and unexpected), and hidden Markov models to represent aggregate learning behaviors (Biswas *et al.*, 2010; Kinnebrew & Biswas, 2011b). These approaches have led to some important insights about how students go about their learning tasks in the Betty's Brain system. For example, one study illustrated that students who taught their agents most effectively were more likely to interleave small amounts of knowledge-acquisition activities (*e.g.* reading) with small amounts of knowledge-construction activities (*e.g.* creating causal links). Conversely, lower-performing students were more likely to perform large amounts of both knowledge-

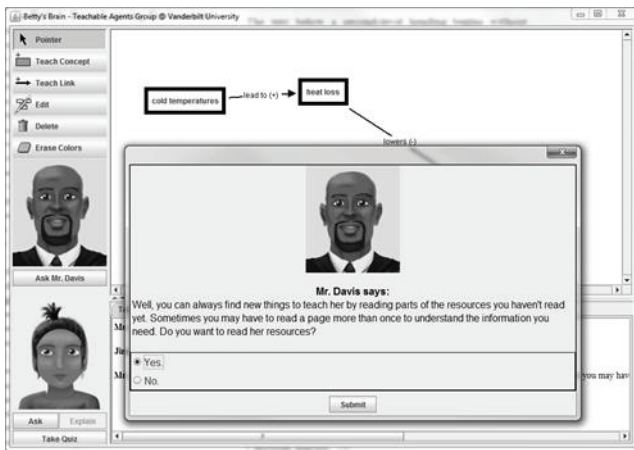


Figure 1. Betty's Brain during a conversation with the Mentor

acquisition and knowledge-structuring activities without interleaving them (Kinnebrew & Biswas, 2011b).

In this paper, we apply analytical methods and experiential analyses to interpret the results from previous classroom studies (Biswas *et al.*, 2005, 2010; Kinnebrew & Biswas, 2011a; Segedy, Kinnebrew, & Biswas 2011a) and propose an integrated cognitive and metacognitive model for effective, self-regulated student learning in the Betty's Brain system. This model helps us interpret the deficiencies in several observed learner strategies that produce suboptimal learning. An understanding of students' suboptimal learning choices in comparison to the idealized model forms the basis for systematically defining scaffolds and feedback in Betty's Brain. In recent studies, we are gaining evidence that such scaffolds, delivered through a mixed-initiative dialog between the student and a Mentor agent, result in positive shifts in students' learning

| Basic Action | Cognitive Activity |
|-------------------------|---|
| (1) Access Resources | Read domain information (<i>e.g.</i> , about entities, their role in the process(es) under study, and their relationships) |
| | Identify specific causal relation(s) |
| | Search for information (about specific concepts or causal relations) |
| (2) Edit Map | Add causal relation(s) between domain entities (concepts) |
| | Organize map (<i>e.g.</i> , group by sub-process based on concept-process associations) |
| | Revise causal relation(s) |
| (3) Query | Check indirect (<i>i.e.</i> , chain of links) relations between two concepts |
| (4) Request Explanation | Identify relevant links involved in answering a query or quiz question |
| (5) Quiz | Assess correctness of specific areas of map (by acquiring a set of correct and incorrect answers to questions posed by the Mentor) |
| | Gain progress feedback (from correct/incorrect quiz questions) |
| (6) Workbook | Note status of specific causal relation(s) (as to whether they have been checked for correctness or require checking) |
| | Write text notes as a memory aid (<i>e.g.</i> , for goals, plans, and monitoring strategies) |
| (7) Ask Mentor | Ask for strategy advice (<i>i.e.</i> , in the current situation what strategy could the student employ to achieve their current goal?) |
| | Ask about the correctness of, or any problems in, the TA's answer to a question. |
| | Review how to organize and build causal concept maps |

Table 1. Basic actions in Betty's Brain and related cognitive activities

behaviors and yield significant learning gains (Kinnebrew & Biswas, 2011a; Segedy, Kinnebrew, & Biswas 2011).

Betty's Brain

The Betty's Brain learning environment (*e.g.*, Leelawong and Biswas, 2008; Biswas *et al.*, 2010), employs the learning-by-teaching paradigm to help middle-school students learn science content while developing strategies for regulating their learning in open-ended, choice-rich learning tasks. It features both a Teachable Agent (TA), which learners teach about a science topic, such as climate change or body temperature regulation, and a Mentor agent that monitors student learning, offers advice when students initiate a conversation to seek help, and provides feedback when the learner demonstrates suboptimal learning

strategies. Learners teach their TA by constructing a causal concept map, where concepts are domain entities, and links represent causal relations between entities (e.g., heat reflected back to earth **increases** global temperature). Once taught, the TA can use the map to answer causal questions (e.g., if electricity generation is decreased, what will happen to global temperature?) and explain those answers by reasoning through chains of links (Leelawong & Biswas, 2008). The TA can also take quizzes (a set of questions created and graded by the mentor) at the learner's request. Learners know that they have completed their learning and teaching task when their TA can correctly answer all of the questions on the quiz.

As learners teach their TA, they may use several actions that are supported by the system. These *basic actions* include the following: (1) reading the hypertext resources; (2) editing the causal concept map; (3) querying the TA; (4) asking the TA to explain an answer; (5) quizzing the TA; (6) using the workbook to take notes and record important information about causal relations; and (7) requesting help from the Mentor. Using these actions, learners can accomplish several cognitive activities, such as identifying causal relationships, requesting performance feedback, and checking whether a link they created implies the right causal relation. Examples of important cognitive activities directly related to actions that can be performed on the system are listed in Table 1. These activities provide learners with a complete set of tools that they can use to accomplish their learning task.

The Integrated Cognitive and Metacognitive Model

In order to scaffold students' learning in Betty's Brain, we need an understanding of how students can employ the relevant cognitive activities (described in Table 1) to effectively and efficiently regulate their learning. Our cognitive and metacognitive model of effective, self-regulated learner behavior in Betty's Brain is illustrated in Figure 3. The model depicts an idealized set of metacognitive activities and strategies that students could employ in Betty's Brain for science learning with understanding.

The two inner layers of the model consist of the basic actions supported by the system and the relevant cognitive activities that are linked to the basic actions, as detailed in Table 1. The outermost metacognitive layer represents four related areas of metacognitive activity: (1) Goal-setting and planning to order and regulate the learners other cognitive and metacognitive activities; (2) knowledge construction, which involves processes related to acquiring information from the resources, integrating the information

with existing knowledge, and teaching the new information to the TA by structuring it in the causal concept map; (3) monitoring the learner's evolving understanding of the science topic and progress in teaching the TA; and (4) help-seeking, which involves identifying the need for help and interacting with the Mentor to get advice on relevant strategies and actions.

Within each modeled area of metacognition, Figure 3 illustrates useful metacognitive strategies that are further elaborated in Table 2, which provides a general description of the strategy and its relation to cognitive activities in the Betty's Brain environment. For example, knowledge construction involves two primary strategies: knowledge acquisition or verification, (acquiring new knowledge or verifying that the current knowledge is correct) and then knowledge structuring and integration (converting knowledge to a structured form to allow for integration with existing knowledge). In the Betty's Brain environment, these strategies are often implemented by determining that specific knowledge needs to be added to the map, then finding and reading the appropriate section of the resources, recognizing information from the text as a relationship between concepts, and finally adding the corresponding link to the map.

Learner Difficulties in Betty's Brain

We hope that learners using the Betty's Brain system will employ effective metacognitive strategies as they go about their open-ended learning tasks. However, the learners in our experiments, middle school students in grades 5-8, are novices in the science domains; they also usually lack the experience of independent learning and teaching. Consequently, they encounter a variety of difficulties in their learning tasks as they attempt to teach their student by building the causal map. In particular, many of these problems manifest because students lack awareness of their own learning strategies and rarely employ effective goal-setting and planning strategies.

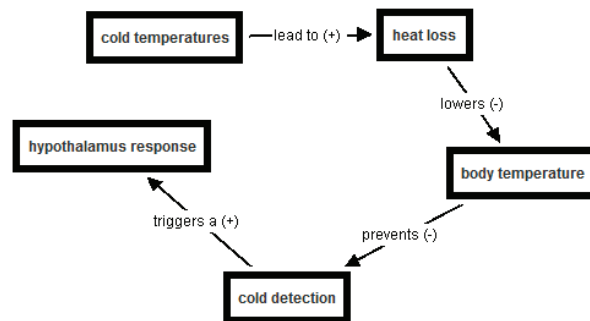


Figure 2. Example concept map for the body's response to cold temperatures

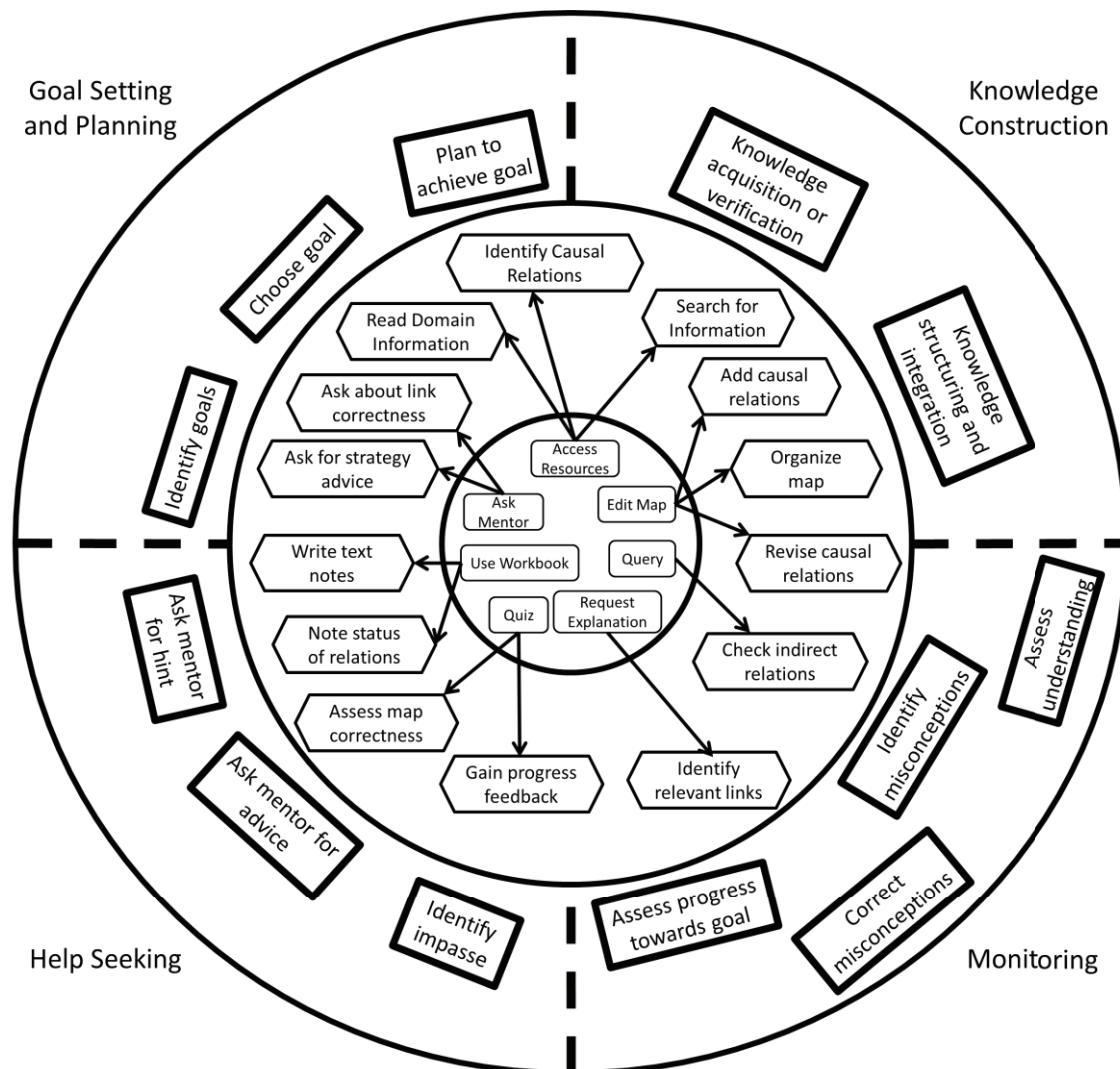


Figure 3. Integrated cognitive and metacognitive model for learning in Betty's Brain

Consider a learner trying to teach his TA how cold temperatures trigger a hypothalamus response based on a set of resources describing the simple chain of links illustrated in Figure 2. The first sub-optimal strategy learners often employ involves *teaching information to their TA based only on their prior knowledge*. Learners often bring misconceptions and incomplete understanding to a new learning situation (Basu, Biswas, & Sengupta, 2011). If they teach these misconceptions or this partial information to their TA through the causal map, it negatively affects the TA's ability to correctly answer quiz questions. So, in the scenario above, a learner might choose to teach a link based on incomplete prior knowledge: cold temperatures increase cold detection. In this case, the learner has taught his TA correct information because cold temperatures do increase cold detection. However, adding a link that directly explains this fails to elucidate *how* or *why* it takes place. Rather, to provide his

TA with complete information, the learner must explain the *process* by which cold temperatures affect cold detection, as illustrated in the map from Figure 2.

Another suboptimal strategy that learners often employ involves *utilizing quiz results for ad hoc map construction*. If our hypothetical learner, after teaching the link from cold temperatures to cold detection, were to have their TA take a quiz, the quiz would contain a question designed to point out the learner's incomplete information: "If cold temperatures increase, what happens to body temperature?" The learner in our scenario would notice that the TA was unable to answer this question, and so he may reason about what the correct answer should be. Without consulting the hypertext resources, he might reason that cold temperatures should reduce body temperature. Thus, he would teach this to his TA with a direct link from cold temperatures to body temperature.

The quiz might also contain another question the TA answered incorrectly: “If cold temperatures increase, what happens to the hypothalamus response?” The answer to

this question is less obvious, so our learner decides to turn to the resources to research it. After a quick skimming of a section of the text, he determines (incorrectly) that cold

| Metacognitive | | Integration | Cognitive |
|------------------------------------|--|--|---|
| <i>Metacognitive Area</i> | <i>Metacognitive Strategy/Activity</i> | <i>Applying Cognition and Metacognition</i> | <i>Supporting Cognitive Activities</i> |
| Goal-setting & Planning | Identify goals | Read introductory, overview, and organization material to identify different sub-processes or integration of sub-processes | Read resources |
| | Choose goal | Consider incorrect questions to identify sub-process with few or no answers correct on quiz | Gain progress feedback by taking quiz |
| | Plan to achieve goal (e.g., understanding and teaching TA about sub-process) | Skim introduction to sections and organization material to identify relevant resources for goal | Read resources |
| | | Take notes about which resources to read and relevant strategies to apply in achieving goal | Write text notes as a memory aid |
| Knowledge Construction | Knowledge acquisition/verification | Read identified resources relevant to goal | Read resources |
| | | Determine causal relations from reading | Identify causal relation(s) |
| | Knowledge structuring and integration | Edit map to add links for identified causal relations | Add causal relation(s) to map |
| | | Reorganize map by concept-process association | Organize map by sub-process |
| Monitoring | Assess understanding | Check direct and indirect relations that involve recently added links against current understanding | Check direct links by re-reading resources, and indirect relations by query, quiz, and explain activities |
| | Identify misconceptions and incomplete representations | Identify overlap in links between correct and incorrect quiz questions | Identify relevant links |
| | | Note links from correct answers as checked and unchecked links from incorrect answers as needing to be checked | Note status of causal relation(s) |
| | | Find resource material to check identified links | Search for information |
| | Correct misconceptions and incomplete representations | Determine corrected/ replacement relations | Identify specific causal relation(s) |
| | | Correct/replace links for identified causal relations | Revise causal relation(s) |
| | Assess progress toward goal | Assess change/progress on questions related to current goal and recent edits | Gain progress feedback |
| Help-Seeking | Identify impasse or difficulty | Identify inability to progress toward goal or correct misconceptions based on quiz results | Gain progress feedback |
| | Ask Mentor for specific advice | Initiate conversation with Mentor and direct conversation based on current goal and difficulties | Ask for relevant strategy advice |
| | Ask Mentor for hint | After the TA has explained her answer to a causal question, ask the Mentor if the explanation contained any problems. | Ask about correctness |

Table 2. Metacognitive and cognitive activities in Betty’s Brain

temperatures decrease the hypothalamus response. After all of this work, our learner's map looks like the one shown in Figure 4.

At this point, the learner's map has a number of errors, and this further confuses him because he is unable to pinpoint a source of error that he can correct. When he has his TA take a quiz again, the TA gets the same questions

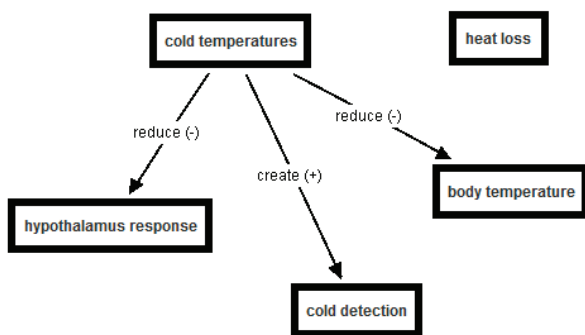


Figure 4. Example student concept map with errors

wrong, even after he seems to have taught the TA how to answer them correctly. As a result, the learner reaches an impasse. He has lost confidence in both his approach and his plans to continue building the map.

At this point, our hypothetical learner may resort to another suboptimal strategy: “guess and check,” in which he removes a link, adds a link, or changes the effect of a link and then checks whether his TA’s performance improves. For example, every time he has his TA take a quiz, he finds a link related to one of the wrong answers, makes a simple uninformed edit (e.g., switching the link from increase to decrease) and then administers another quiz. However, this confuses him more as his TA’s quiz results have the same number of errors. Moreover, some questions that his TA answered correctly on previous quizzes are now being answered incorrectly (because of his uninformed edits), and this naturally leads to discouragement, frustration, and further loss of confidence. At the end of this exercise, our learner has spent very little time learning about thermoregulation, instead spending a large portion of his time using suboptimal learning strategies.

Such a scenario represents real problems that many learners have faced while using Betty’s Brain. Difficulties such as these have led other researchers to argue that open-ended and constructivist learning environments may hinder learning (e.g., Klahr & Nigam, 2004). However, we believe that our approach, if properly scaffolded, will equip learners with an important understanding of how to assimilate and represent difficult scientific information, better preparing them for future learning. Therefore, the next section compares the idealized metacognitive

strategies from our model to the common suboptimal strategies employed by the learner in our scenario. Based on this comparison, we derive conversational feedback for guiding learners toward more productive learning strategies.

Using Feedback to Guide Learners Away From Ineffective Strategies

If we were to look back to our example learner, how might we characterize his strategies in light of our model (see Figure 3)? His first suboptimal strategy was a knowledge construction strategy that involved *teaching prior knowledge without reading the hypertext resources to check if he was building an accurate model*, which resulted in his teaching incomplete information to his TA. It is important to note here that the learner’s decision to teach without reading is not necessarily an ineffective choice; learners may prefer to start by teaching their current understanding before they begin reading. However, when the learner makes choices that, over time, continue to negatively impact his performance (e.g., the quality of his causal map), feedback can encourage him to modify his strategy, and help him improve his learning.

Once we decide to give feedback, we can analyze the difference between his learning strategy and those contained in our model. In this case, the learner’s strategy is negatively impacting his teaching (and learning) by adding incorrect or incomplete information to the map. One strategy suggested by the model involves first *acquiring information or verifying prior understanding*, and then *structuring it* by teaching the TA. Therefore, the Mentor agent might recommend that the learner adopt this strategy. He might say: “Hey Jake! I’ve noticed that you’re teaching information you haven’t read about. How do you know that you are teaching your student the right links for her to understand thermoregulation?” Feedback such as this encourages the learner to reflect on his approach to learning and teaching, which, with additional support from the Mentor agent in understanding effective metacognitive strategies, can lead to the adoption of a better strategy.

Our learner’s second sub-optimal strategy was a monitoring strategy that consisted of *utilizing quiz results for ad hoc map construction*. The learner monitored his progress by having his TA take a quiz, and then he used the result of that quiz to systematically teach several direct relationships between concepts that should not have been directly connected in a complete representation of the process. In our experience, learners who employ this strategy have a misconception about how to use the TA’s quiz results to identify problems and incomplete representations. Further, this may be related to a lack of understanding on how to model a scientific process. Thus,

when such an approach is detected (e.g., by comparing link adds to recent quiz results), the Mentor agent can again interact with the learner, saying *“It looks like you’re trying to help your TA do better on her next quiz. Would you like some help doing this?”* Should the learner agree, the Mentor might give the student choices on the feedback content, such as (a) *reviewing how to construct causal maps*, and (b) *learning how to interpret quiz results*, among others. If the student chooses (b), i.e., how to interpret quiz results, the Mentor would continue by explaining a key insight (e.g., when the TA correctly answers a quiz question, all of the links in her explanation are correct) in the context of how it can be applied to the current situation. The Mentor might also highlight the explain button while reminding students that when the TA explains an answer, she indicates exactly which links she used to generate her answer. Such feedback is directed to teaching learners a valuable monitoring strategy while also helping them achieve better results in their teaching and learning tasks.

Finally, our learner’s last ineffective strategy involved several iterations of first taking a quiz to receive performance feedback, and following that with an uninformed edit. This *“guess and check”* approach to map building generally indicates that the learner has reached (or will soon reach) an impasse: he has lost confidence in his approach, his self-efficacy is suffering, and as a result he is gradually disengaging from the learning task. In terms of the model, our learner is (1) using a highly-ineffective monitoring strategy, (2) no longer participating in any significant planning or goal-setting activities, and (3) choosing not to seek the help that he needs to move forward.

Feedback delivered by the Mentor at this time needs to accomplish three things. First, it should help the learner understand and *re-establish* his current goal as a context for moving forward. It should also *provide positive feedback* on work done so far in order to restore some of the learner’s self-efficacy. Finally, it should *provide a relevant suggestion* for moving forward, giving the learner the help and push he needs to re-engage in the learning task. As an example, the Mentor might say *“Hey Jake. Can I help you? I think that if we work together we can make sure that your TA learns everything she needs to know.”* Upon affirmation from the learner, the Mentor might continue: *“Remember, you need to read in order to learn what to teach your TA. You’ve already taught her some good information, but there are also some links on her map that are confusing her. If you’d like, I can help you by picking out an incorrect link and explaining why it’s wrong. Would you like me to do that?”* Feedback such as this has the potential to help learners who have reached an impasse and are no longer confident in their ability to accomplish the task. More examples of feedback based on

| Metacognitive Area | Example Feedback |
|-------------------------|---|
| Goal-setting & Planning | <i>Try keeping a list of which pages you have read and which pages you still need to work on.</i> |
| Knowledge Construction | <i>As you teach your agent, try to put concepts close to each other when they are connected by a link. This makes it easier to read the map later.</i> |
| Monitoring | <i>The results of your student’s quiz can tell you a lot about what your student understands correctly and what she needs help with.</i> |
| Help-Seeking | <i>If you’d like, I can tell you whether your student is answering your questions correctly. Next time you ask her a question, feel free to ask me if it’s right.</i> |

Table 2: Example Feedback Related to Each Metacognitive Area

our model, organized by metacognitive area, appear in Table 3.

Overall, our cognitive/metacognitive model provides a framework for analyzing learner behaviors and understanding learner difficulties in relation to effective, self-regulated learning strategies. Further, this comparison guides the development of feedback designed to help students overcome the obstacles in their open-ended learning task. We believe that feedback developed in this manner will enhance both science learning and preparation for future learning with the Betty’s Brain learning environment.

Identifying suboptimal strategies and designing agent feedback based on this model have shown promise in initial classroom studies (Segedy, Kinnebrew, and Biswas, 2011). In particular, analysis of a recent study with 7th-grade students in middle Tennessee science classrooms indicated a statistically significant correlation between normalized learning gain and students’ decisions to continue a dialogue with the Mentor agent when offering strategy feedback. Similarly, analysis of students’ actions following this strategy feedback indicated a statistically significant correlation between normalized learning gain and actions consistent with the Mentor’s advice. Moreover, these correlations were both stronger than the correlation between normalized learning gain and students’ prior performance on standardized academic tests. These preliminary results illustrate the potential of the feedback designed through comparison of students’ suboptimal learning strategies with the presented cognitive and metacognitive model.

Conclusion

In this paper, we have presented a cognitive and metacognitive model for effective learning strategies in the Betty's Brain learning environment. We further applied this model to interpret three commonly-observed difficulties that learners face while using the system. This comparison motivated the development of feedback to help learners overcome obstacles and engage in effective learning in an open-ended learning environment. We believe that such feedback, when delivered through a conversation with the Mentor agent and contextualized by the learner's current activities, encourages learners to practice and develop effective strategies for navigating open-ended learning activities.

In the continuing development of Betty's Brain and future classroom studies, we will apply our model to additional learner difficulties in order to develop a more comprehensive collection of feedback for both scaffolding learning in open-ended learning environments and promoting metacognitive, self-regulatory skills. This scaffolding and feedback is crucial to our goal of empowering students to take control of their learning as they practice using effective learning strategies.

Another important avenue for future work involves employing the cognitive and metacognitive model to systematically interpret activity traces of learners while using Betty's Brain. In particular, interpreting the cognitive and metacognitive states represented in hidden Markov models derived from student activity traces (Biswas *et al.*, 2010) with respect to this model may yield a deeper understanding of student learning strategies and how to detect suboptimal strategies in Betty's Brain.

References

- Azevedo, R. 2005. "Using Hypermedia as a Metacognitive Tool for Enhancing Student Learning: The Role of Self-regulated Learning." *Educational Psychologist*, 40, 199-209.
- Azevedo, R., and Witherspoon, A. M. 2009. "Self-regulated Learning with Hypermedia." In D. J. Hacker, J. Dunlosky, & A. C. Graesser (Eds.), *Handbook of Metacognition in Education* (pp. 319-339). Mahwah, NJ: Erlbaum.
- Bandura, A. (1986). "Social foundations of thought and action: A social cognitive theory." Englewood Cliffs, NJ: Prentice Hall.
- Basu, S., Biswas, G., & Sengupta, P. 2011. "Scaffolding to Support Learning of Ecology in Simulation Environments," *Proceedings of The 15th International Conference on Artificial Intelligence in Education*, Auckland, New Zealand, pp.
- Biswas, G., Leelawong, K., Schwartz, D., & Vye, N. 2005. "Learning by Teaching: A New Agent Paradigm for Educational Software." *Applied Artificial Intelligence*, 19: 363-392.
- Biswas, G., Jeong, H., Kinnebrew, J. S., 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)*, 5(2): 123-152.
- Bransford, J. D. & Schwartz, D. L. 1999. "Rethinking Transfer: A Simple Proposal with Multiple Implications. In A. Iran-Nejad and P. D. Pearson (Eds.), *Review of Research in Education*, 24: 61-100. Washington, D.C.: American Educational Research Association.
- Chin, D. B., Dohmen, I. M., Cheng, B. H., Opezzo, M. A., Chase, C. C., & Schwartz, D. L. (2010). "Preparing students for future learning with Teachable Agents." *Educational Technology Research and Development*.
- Klahr, D. and Nigam, M. 2004. "The Equivalence of Learning Paths in Early Science Instruction: Effects of Direct Instruction and Discovery Learning," *Psychological Science*, 15(10): 661-667.
- Kinnebrew, J.S., & Biswas, G. 2011a. "Modeling and Measuring Self-Regulated Learning in Teachable Agent Environments," *Journal of e-Learning and Knowledge Society*, 7(2).
- Kinnebrew, J.S., & Biswas, G. 2011b. "Comparative Action Sequence Analysis with Hidden Markov Models and Sequence Mining." In *The Workshop on Knowledge Discovery in Educational Data at The ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*, San Diego, CA, August 21-24, 2011.
- 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.
- Roll, I., Aleven, V., McLaren, B.M., & Koedinger, K.R. 2007. "Designing for metacognition—applying cognitive tutor principles to the tutoring of help seeking," *Metacognition and Learning*, 2(2): 125-140.
- Schunk, D. H. & Zimmerman, B. J. (1997). "Social origins of self-regulatory competence." *Educational Psychologist*, 32, 195–208.
- Schwartz, D. L., Chase, C., Wagster, J., Okita, S., Roscoe, R., Chin, D., & Biswas, G. 2009. "Interactive Metacognition: Monitoring and Regulating a Teachable Agent." In D. J. Hacker, J. Dunlosky, and A. C. Graesser (Eds.), *Handbook of Metacognition in Education* (pp. 340-359). Mahwah, NJ: Erlbaum.
- Segedy, J., Kinnebrew, J. S., and Biswas, G. 2011. "Investigating the Relationship Between Dialogue Responsiveness and Learning in a Teachable Agent Environment", *The 15th International Conference on Artificial Intelligence in Education*, Auckland, New Zealand.
- Zimmerman, B.J. 2001. "Theories of Self-regulated Learning and Academic Achievement: An Overview and Analysis." In B. Zimmerman & D. Schunk (Eds.), *Self-regulated Learning and Academic Achievement: Theoretical Perspectives* (pp. 1–37). Mahwah, NJ: Erlbaum.