Interface Design Issues for Teachable Agent Systems

Karun Viswanath Ramya Balachandran Mattie Ruth Kramer Bobby Bodenheimer
Department of Electrical Engineering and Computer Science
Vanderbilt University, Nashville, TN - 37235
{karun.viswanath, ramya.balachandran, mattie.r.kramer}@vanderbilt.edu, bobbyb@vuse.vanderbilt.edu

Abstract: This paper discusses improvements to the user interface of a system designed to promote learning through teaching. This system, called Betty's Brain, is an intelligent agent developed at Vanderbilt University for studying the learning by teaching paradigm. In the Betty's Brain system, students teach a computer agent by creating a concept map using a visual interface. Students themselves learn through the process of instructing the agent. Students aged 9-11 (U.S. fifth-grade) in Nashville public schools participated in studies using this system. In this work, we improve and evaluate the interface components of Betty's Brain that enable the user to organize and input problem solving knowledge about a domain for the purpose of instructing the intelligent agent. We then conducted a comparative user study to evaluate our changes to the interface. Both qualitative and quantitative improvements in the user's performance are reported. These results should provide useful guidance for designers of learning by teaching systems.

Introduction

The system described in this paper is founded on the idea that the process of teaching others can be a powerful method of learning. Studies have shown that people who prepared to teach others to take a quiz on a passage learned the passage better than those who prepared to take the quiz themselves [1]. Similarly, literature on tutoring suggests a similar conclusion in that tutors have been shown to benefit as much from tutoring as their tutees [2, 3]. Biswas et al. [4] report that students preparing to teach stated that the responsibility of teaching forced them to a deeper understanding of the material; other students focused on the importance of having a clear conceptual organization of the materials.

The Teachable Agents group at Vanderbilt University [5, 12, 15] has built an environment where students can explicitly teach and directly receive feedback about their teaching through interactions with a computer agent. In this environment, called Betty's Brain, students learn by teaching computer agents using cause and effect structures to solve problems and answer questions in a variety of scientific domains. Teachable agents have no a priori knowledge, only an ability to reason. The agent is explicitly taught knowledge by the user. Thus, teachable agents differ from pedagogical agents [16, 17] since pedagogical agents are primarily demonstrative. The effectiveness of the teachable agent environment in producing learning-related outcomes has been reported in [5, 14].

This paper presents an analysis evaluating the user interface of the teachable agent system. It analyzes the various trade-offs involved in determining which features should be present in the user interface and how the various components should be designed. Some of these refinements were suggested by Viswanath et al. [15], a preliminary study on the efficacy of the user interface. All refinements and changes were evaluated through studies on public school students age 9-11 (U.S. fifth grade), a target audience for this system.

The domain of knowledge for the system described in this paper is a river eco-system. In Tennessee, students in public schools age 9-11 learn about eco-systems as part of their science curriculum, including the major levels of biological classification, photosynthesis, etc. They understand direct causal relationships clearly, but they are not familiar with the concept of a balanced system, or how small changes can result in geometric propagation over time.

Background

An automated learning by teaching system requires a representation scheme for students to create their knowledge structure as a part of the teaching process. Since the primary users are students aged 9-14 (U.S. "middle-school") who are solving complex problems, this representation has to be intuitive but sufficiently expressive to help these students create, organize, and analyze their problem solving ideas. A widely accepted technique for constructing knowledge is a concept map [6, 7], and this structure is the one we employ.

Concept maps provide a mechanism for structuring and organizing knowledge into hierarchies, and allow the analysis of phenomena such as cause-effect relations [6, 8, 9, 10]. Moreover, an intelligent software agent based on concept maps can employ reasoning and explanation mechanisms that students can easily relate to. Thus the concept map provides an excellent representation that serves as the interface between the student and the teachable agent.

The Concept Map

A concept map is a collection of concepts and the relationships between these concepts. It is also a mechanism for representing domain knowledge [6]. A partial concept map in the domain of river ecosystems is shown in Figure 1. The labeled boxes correspond to entities (the labels are entity names), and the labeled links correspond to relations. The arrow indicates the direction of the relation, and its name appears by the arrow. The parenthesized phrase indicates the relation type.

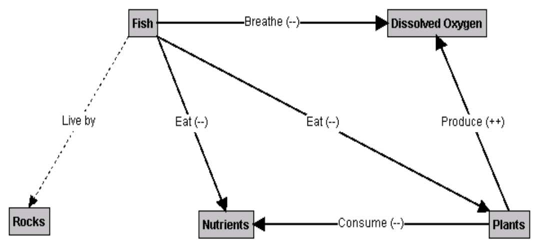

Figure 1: A partial concept map.

In our environment, concepts are **entities** that are of interest in the domain of study. Common entities in a river ecosystem are fish, plants, bacteria, dissolved oxygen, carbon dioxide, algae, and waste. Relations are unidirectional, binary links connecting two entities. They help to categorize groups of objects or express interactions among them.

In the current implementation of domain knowledge, such as for a river ecosystem, students can use three kinds of relations to build a concept map: (i) **cause-effect**, (ii) **descriptive**, and (iii) **type-of**. The primary relation students use to describe relations between entities is the cause-effect relation, such as “Fish eat Plants” and “Plants produce Dissolved oxygen”. The causal relations are further qualified by **increase** (‘++’) and **decrease** (‘**--**’) labels. For example, “eat” implies a decrease relation, and “produce” implies an increase. Therefore, an introduction of more fish into the ecosystem causes a decrease in plants, while an increase in plants causes an increase in oxygen.

The descriptive relation is similar to the cause-effect relation. It also expresses a dependency, but the change in one entity does not cause a change in the other entity. An example is the relation “Fish live by Rocks”. In this case, the “live by” relation is categorized as a descriptive relation. Fish use rocks, but an increase or decrease in fish does not directly change the amount of rocks. More complex forms of the “descriptive” relation, e.g., “Plants need Sunlight to produce Dissolved Oxygen” have not yet been implemented in Betty’s Brain.

Type-of relations let students establish class structures to organize the domain knowledge. Consider an example where students deal with a variety of fish, such as trout, bass, blue gill, and catfish. All of these fish types breathe dissolved oxygen and eat plants. To simplify the knowledge construction process, students can first create the entity “Fish”, and express the “Fish eat Plants” and “Fish breathe Dissolved Oxygen” relations. Then, they can create individual fish entities, such as “trout” and “bass”, and link them to the “Fish” entity using “type-of” links. All relations associated with the entity “Fish” are inherited by these individual types unless they are over-ridden by more specific links [11].

Description of Betty’s Brain

This section presents the user interface as described by Davis et al. [5]. Figure 2 illustrates the Betty’s Brain interface. The system possesses multimedia capabilities. Students use a graphical drag and drop interface to create and modify their concept maps. When queried, Betty can provide explanations for how she derives her answers, and simultaneously depict the derivation process on the concept map by animation. The interface of Betty’s Brain is implemented in Java with Java Swing components. In the sections below, we describe the software’s three modes: TEACH, QUERY and QUIZ.

TEACH Betty. As mentioned previously, students teach Betty by means of a concept map interface. Figure 2 displays an example of a concept map used to teach Betty about the river ecosystem. The map shown is not totally accurate or complete.

QUERY Betty. Students are able to query Betty about what they have taught her. Currently, Betty has templates for two types of queries:

1. What will happen to concept A when we increase/decrease concept B?
2. Tell me what you know about concept A.

To answer questions, Betty uses a simple chaining procedure to deduce the relationship between a set of connected concepts [13]. After Betty answers a question, the student can then ask her to explain how she derived her answer. Betty verbally explains her answer while simultaneously animating the concept map.

QUIZ Betty. During the quiz phase, students observe Betty's responses to a set of pre-scripted questions. A mentor agent informs Betty (and the student) if Betty's answers are right or wrong. The teaching expert also gives hints for each question to help the student debug errors in the concept map they have created. Currently, the system implements three levels of hinting, from suggesting that the student read a section of the resource materials to directly telling the student how to correct the concept or link in the concept map.

An Enhanced Interface For The Teachable Agents System

A conclusion of a study conducted in the Nashville Metropolitan school system and reported in Davis et al. [5] was that the user interface needed improvement. This study suggested that students should be more focused on domain knowledge rather than on specifics of the interface. The cognitive load involved in operating the interface should be minimized. Also, our goal was to make the interface self-explanatory so that the users found it easy to operate the interface with minimal training.

Figure 3 displays our enhanced interface for Betty's Brain. The enhancements focused on making the system easier to use for novices. These changes affect the visual and the cognitive aspects of the interface. Visual aspects of the interface deal with its aesthetic appearance and its look-and-feel. Cognitive aspects deal with functionality and how information is represented in the interface. These two aspects are clearly interrelated.

Causal links form an important part of the concept map structure. The existing interface differentiated between the increasing and decreasing effects by using the symbols (++) and (--), respectively. With hindsight, this design is confusing and a burden to young users. The new interface uses colors to depict the meaning of the causal links: red for a decreasing effect and green for an increasing effect. However, we decided to keep the symbols (++) and (--) along with the colors. This additional information has the positive impact of providing reinforcing information about the causal links. It should also be helpful to any users of the system who are color-blind.

Data from the prior study indicated that common errors students made were the incorrect representation of the effect of a causal link (creating a causal link with increasing effect instead of a decreasing effect) and specifying the direction of the arrow incorrectly. During the query or quiz phase, when students discovered these errors, they found it frustrating to correct them using the existing interface. For example, changing the direction of a "type-of" link from "Animal" to "Frog" involved first selecting the link, deleting it, and then recreating a "type-of" link from "Frog" to "Animal".

In the new interface, two buttons – "Switch Trend" and "Switch Direction" – were provided to help students easily rectify these common mistakes with a single click. For example, to change a trend in the relation "Animals eat plants causing plants to increase", the student would select the link and click on the "Switch Trend" button to rectify the mistake. The "Switch Direction" button changes the direction of the link selected. Figure 4 shows the use of the button "Switch Trend".

An issue in the introduction of these buttons was the concern that students might use them to employ a "test – modify – retest" strategy. This strategy would enable them to "game" the system by constructing a concept map simply designed to pass the quiz phase. This strategy is clearly undesirable and to guard against it, a pattern tracking agent was built to detect this pattern. If detected, then Betty refuses to take a quiz, both complaining that she is confused and suggesting that the student help her with the resources before making any further changes.

The existing interface provided three different buttons for adding a link in the concept map, corresponding to the three different types of links. Though currently screen space is not an issue, in the future when new features are added there may be space constraints. Therefore these three buttons were replaced with just one, labeled "Teach

Link”, in the new interface. On clicking this button, a window pops up as shown in Figure 5. The users can select from this window which type of link they want to create.

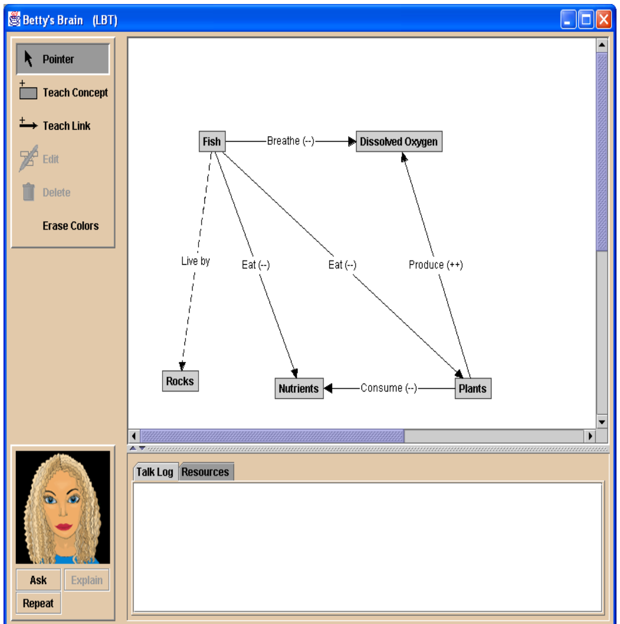

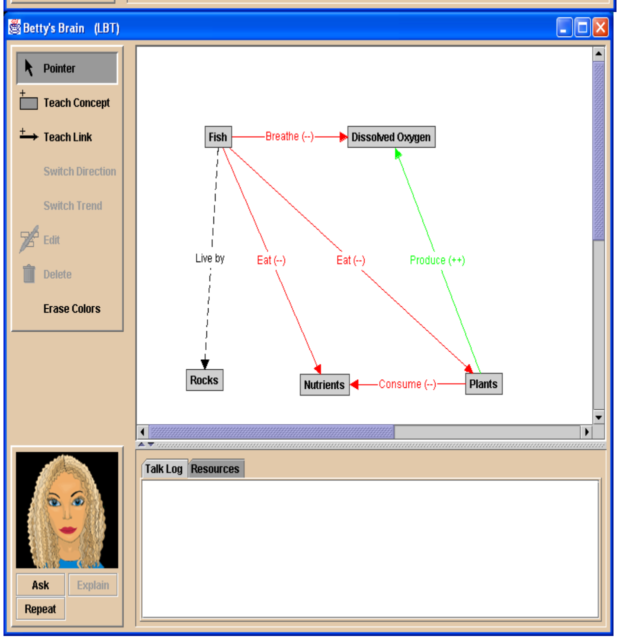

Figure 2: Existing Betty’s Brain Interface.

Figure 3: Enhanced Interface for Betty’s Brain.

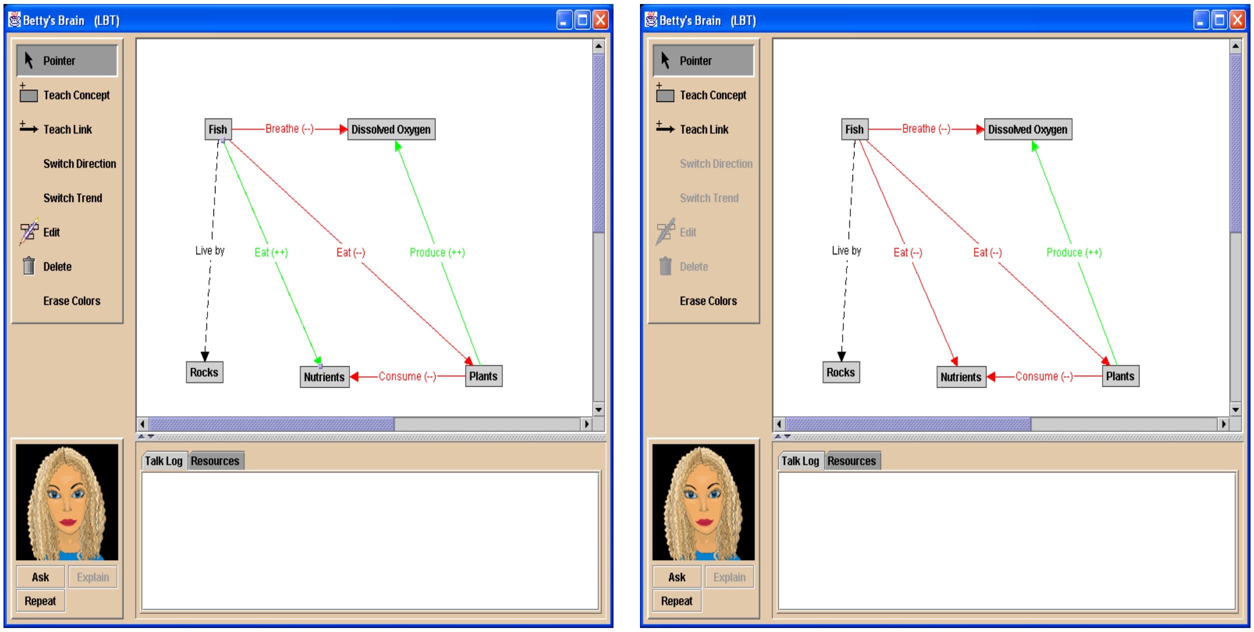

(a) (b)

Figure 4: The interface before (a) and after (b) use of the “Switch Trend” button.

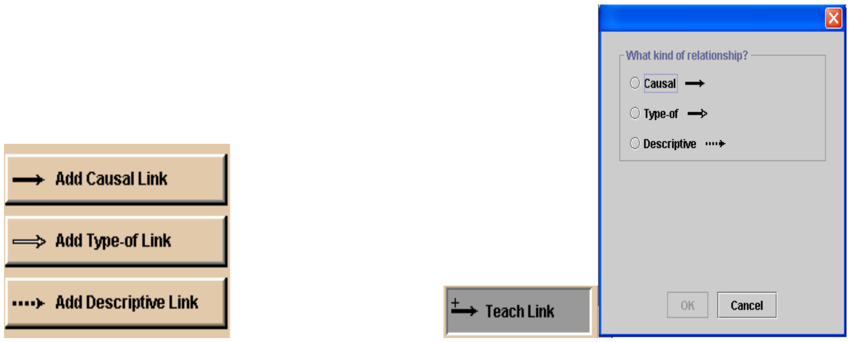

(a) (b)

Figure 5: (a) Original individual buttons
(b) Pop-up window displayed in the new interface when "Teach Link" button is clicked

User Study

To evaluate our changes, we conducted a usability study on the interface. Study participants were 48 high achieving students aged 9-11 (U.S. fifth-grade) from an urban public school located in Nashville. High achieving students were selected to avoid confounds associated with poor reading ability. The students were part of a larger study that investigated the effects of self-regulated learning using Betty's Brain, and were familiar with working on Betty's Brain by the time this study was conducted. Hence no training in working with the interfaces was needed, though the new features of the interfaces were shown and explained to them. The user study had three phases: an introduction, testing and questionnaire sections.

In the testing phase, the students were asked to debug a concept map that had two errors. Half the students worked on the new interface first while the other half worked on the old interface first. Then they switched the interfaces they were working on, now carrying out the tasks on a different concept map. The order of the concept maps was also switched to remove any bias that could possibly be created by the different maps used.

Finally users answered a paper questionnaire. In the questionnaire, they were asked to compare features in the two interfaces. Also, to find out how color affects the understandability, concept maps outside the river eco-system domain were given and they were asked to answer questions on them.

Discussion

Log files were generated for each section of the test tracking every action of the user along with a time stamp. These timings were used for comparing the two interfaces. They were analyzed using a two-way ANOVA to control the effect of the ordering of the interfaces.

When debugging a concept map, users were significantly faster using the new interface as opposed to the old ($F=14.92$, $p=0.0002203$, $MS=72680$). The effect of the ordering of the interfaces on their time was not significant. This result is reasonable since, as mentioned before, the users were experienced with the software and thus the order in which the interfaces were presented to them did not have any significant effect. Interaction effects were not significant.

The mean values of the time taken by the users are presented in Table 1. These values indicate that designing a user interface that demanded less cognitive load improves the efficiency of the system for students. Thus, creating a simple, easy-to-understand interface with special purpose tools enables the users to be more effective.

Table 1: Time taken by users in the tests.

	Avg. Time taken to debug (sec)
Old Interface	124.41
New Interface	66.93

In the questionnaire, the users were asked to mark their preference regarding adding the links – whether they liked to have separate buttons for each link or they liked the pop-up style. 33 out of the 48 users preferred the pop-up style, whereas the other 15 preferred to have separate buttons.

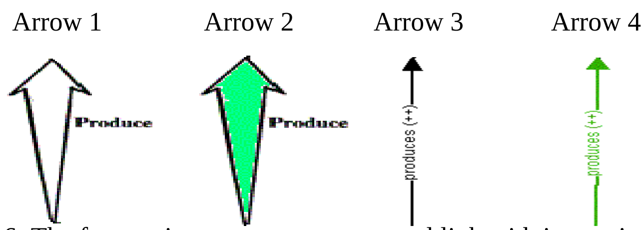

Figure 6: The four options to represent a causal link with increasing effect.

Table 2: User's feedback in the questionnaire section regarding the causal links.

	Highest Preference → Lowest Preference			
Arrow 1	2	10	15	21
Arrow 2	20	16	18	4
Arrow 3	2	8	16	22
Arrow 4	24	14	9	1

Table 2 shows other feedback obtained from the users. The users were given four different options, shown in Figure 6, for representing a causal link with an increasing effect and the users were asked to rate them according to the order of preference. As can be seen from Table 2 the users felt that the arrows used in the new interface were an improvement.

However, the analysis of the questionnaire presented a different picture. The questionnaire had two concept maps - one using normal block arrows (like in the old interface) and another, which used colored arrows (like in the new interface). The students were asked to answer two simple questions of the type "If A increases what happens to B and C ?" as shown in Figure 7. Nearly half of the students (23 out of 48) answered this question incorrectly for the colored concept map, compared to only one-fourth (12 out of 48) for the uncolored one. Thus, the number of incorrect answers to these questions was twice as high for the colored concept map as for the uncolored one. This result reveals that the students had incorrectly understood the color changes. This flaw was not exposed when the students worked on the full system as they were working on a concept map in a domain they were familiar with. Asking them questions outside this domain helped us to identify this flaw. Thus, though the changes to the user interface made it simpler to use and increased the efficiency, it suffered from the primary drawback of being incorrectly interpreted by the students.

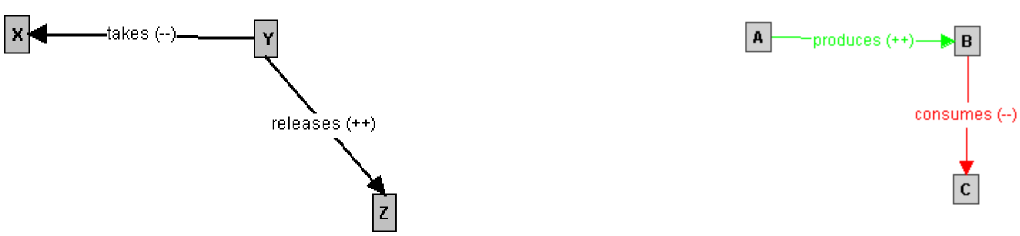

If **Y** decreases, what happens to **X** and **Z**?

If **A** decreases, what happens to **B** and **C**?

Figure 7: Questions asked to the students in the questionnaire

Conclusion and Future Work

This paper presented a usability study of a teachable agent system on its target audience, school children in U.S. "middle"-schools, aged 9-11. (U.S. middle schools typically encompass ages 9 through 14, but the domain of knowledge for our particular study is taught to students aged 9-11). The design of a teachable agent system differs from that of a tutoring system. A tutoring system emphasizes problem solving with its primary focus on modeling the student and tailoring feedback to address the student's immediate needs. As a result the student does not have opportunities to learn by exploration. The user interfaces of such systems reflect this reality both in their design and implementation. A teachable agent system, on the other hand, provides both focused feedback and opportunities for open-ended exploratory activities. The user interface must reflect this usage and provide features that enable students to explore productively and learn a domain of knowledge. Thus the interface needs to find a balance between encouraging the students to be exploratory and being cognitively burdening to children. The present study evaluated this trade-off.

First, students indicated that they preferred the pop-up mechanism for adding the links. This feature saves valuable screen space, which can be used for providing other tools. Second, the introduction of specialized buttons for “Switch Trend” and “Switch Direction” was found to be very effective in allowing students to debug concept maps. The intent of these buttons was to reduce the cognitive load of correcting mistakes, allowing students to focus on issues of domain knowledge. To guard against a “test-retest-modify” gaming strategy for the system, a pattern tracking agent was inserted that prevented students from engaging in this behavior, made easier by the introduction of these buttons.

With respect to the use of color in the concept map, earlier experiments conducted with college-age students [15] reported improved efficiency in using the system with this addition. However, the changes had not confused these more sophisticated users. In the present study, though the efficiency in using the system had similarly increased and the target students preferred the color changes, they misinterpreted them. Thus, changes made to a system must specifically keep in mind the cognitive ability of the users and a “one size fits all” policy for certain issues can only lead to problems. Careful analysis must be carried out before making any such changes.

Future experiments will evaluate how interface modifications affect learning-related outcomes. For example, many simple changes may create an interface that is sufficiently easy to use but will allow users to exploit shortcomings in the reasoning portions of the system with the result that they are able to “game” the system, i.e., have Betty achieve good results without themselves learning much. Also, future work in this field will explore the judicious use of colors and shapes to represent hierarchies and relationships between entities. The use of these in animations to explain the complex relations in a system will also be investigated.

Acknowledgments

This research was supported by National Science Foundation Grant NSF-9873520 and CAREER IIS-0237621. The opinions expressed here are those of the authors and should not be attributed to the funding agency. We thank Gautam Biswas, Nancy Vye, John Bransford, Kadira Belynne, Joan Davis, and Bethany Rittle-Johnson for their invaluable help on this project. We also acknowledge feedback from Krittaya Leelawong, and Thomas Katzlberger who are other members of the Teachable Agents project.

References

1. Bargh, J. A. & Schul, Y. (1980). On the cognitive benefits of teaching. *Journal of Educational Psychology, 72(5)*, 593-604.

2. Graesser, A. C., Person, N., & Magliano, J. (1995). Collaborative dialog patterns in naturalistic one-on-one tutoring. *Applied Cognitive Psychologist*, *9*, 359-387.

3. Chi, M. T.H., Siler, S. A., Jeong, H., Yamauchi, T., & Hausmann, R. G. (2001). Learning from human tutoring. *Cognitive Science*, *25(4*), 471-533.

4. Biswas, G., Schwartz, D., Bransford, J. & TAG-V. (2001). Technology support for complex problem solving: From SAD environments to AI. In Forbus and Feltovich, (Eds.), *Smart Machines in Education* (pp. 71-98). Menlo Park, CA: AAAI Press.

5. J. Davis, K. Leelawong, K. Belynne, B. Bodenheimer, G. Biswas, N. Vye, and J. Bransford. Intelligent User Interface Design for Teachable Agent Systems. *2003 International Conference on Intelligent User Interfaces, January2003, pp. 26-33. ISBN 1-58113-586-6*.

6. Novak, J. D. (1996). Concept mapping as a tool for improving science teaching and learning. In: D. F. Treagust; R. Duit; and B. J. Fraser (Eds.), *Improving Teaching and Learning in Science and Mathematics* (pp. 32–43). London: Teachers College Press.

7. Spiro, R. J., and Jehng, J. C. 1990. Cognitive flexibility and hypertext: Theory and technology for the nonlinear and multidimensional traversal of complex subject matter. In D. Nix and R. J. Spiro eds., *Cognition, education, and multimedia: Exploring ideas in high technology*. Hillsdale, NJ: Lawrence Erlbaum.

8. Kinchin, I. M. and Hay, D. B. 2000. How a qualitative approach to concept map analysis can be used to aid learning by illustrating patterns of conceptual development. *Educational Research* 42 (1): 43–57.

9. Stoyanov, S., & Kommers, P. (1999). Agent-support for problem solving through concept-mapping. *Journal of Interactive Learning Research*, *10(3/4*), 401-42.

10. Novak, J. D. 1998. Learning, *Creating, and Using Knowledge: Concept Maps as Facilitative Tools in Schools and Cooperations*. Hillsdale, NJ: Lawrence Erlbaum.

11. Russell, S. J., and Norvig, P. 1995. *Artificial Intelligence: A Modern Approach,* 319-320. Upper Saddle River, NJ: Prentice Hall.

12. Teachable Agents project web page http://teachableagents.org

13. Leelawong K., Yingbin, W., Biswas, G., Vye, N & Bransford, J.(2001). Qualitative reasoning techniques to support learning by teaching: The Teachable Agents project. In *Proceedings of the Fifteenth International Workshop on Qualitative Reasoning* (pp. 109–116). San Antonio, Texas: AAAI Press.

14. Leelawong, K., Davis, J., Vye, N., Biswas, G., Schwartz, D., Belynne, K., Katzlberger, T. and Bransford, J., The Effects of Feedback in Supporting Learning by Teaching in a Teachable Agent Environment. *The Fifth International Conference of the Learning Sciences.* 2002. Seattle, Washington.

15. Viswanath, K., Balachandran, R., Davis, J., Bodenheimer, B., "Effective User Interface design for Teachable Agent Systems", *ACM Southeast Regional Conference,* Savannah, Georgia, 2003.

16. Baylor, A. and Ryu, J. Effects of image and animation on agent persona. Journal of Educational Computing Research.

17. Graesser, A.C., Person, N. and Magliano, J. Collaborative dialog patterns in naturalistic one-on-one tutoring. Applied Cognitive Psychologist, 9. 359-387.