Developing Learning by Teaching Environments that support Self-Regulated Learning

Gautam Biswas[1], Krittaya Leelawong[1], Kadira Belynne[1], Karun Viswanath[1], Daniel Schwartz[2], and Joan Davis[2]

[1] Dept. of EECS & ISIS, Box 1824 Sta B, Vanderbilt University
Nashville, TN 37235. USA.
{gautam.biswas, krittaya.leelawong, kadira.belynne, karun.viswanath}@vanderbilt.edu
http://www.vuse.vanderbilt.edu/~biswas
[2] School of Education, Stanford University
Stanford, CA 94305. USA.
{daniel.schwartz, joan.davis}@stanford.edu

Abstract. Betty's Brain is a teachable agent system in the domain of river ecosystems that combines learning by teaching and self-regulation strategies to promote deep learning and understanding. Scaffolds in the form of hypertext resources, a Mentor agent, and a set of quiz questions help novice students learn and self-assess their own knowledge. The computational architecture is implemented as a multi-agent system to allow flexible and incremental design, and to provide a more realistic social context for interactions between students and the teachable agent. An extensive study that compared three versions of this system: a tutor only version, learning by teaching, and learning by teaching with self-regulation strategies demonstrates the effectiveness of learning by teaching environments, and the impact of self-regulation strategies in improving preparation for learning among novice learners.

1 Introduction

Advances in computer technology have facilitated the development of sophisticated computer-based Intelligent Tutoring Systems (ITS) [1]. The ITS paradigm is problem-based, and has been very successful in developing three core technologies: curriculum sequencing, intelligent analysis of student's solutions, and interactive problem solving support [2]. At the same time, these systems provide localized feedback, and do not emphasize practicing higher-order cognitive skills in complex domains, where problem solving requires active decision-making to set learning goals and to apply strategies for achieving these goals. Our goal has been to introduce effective learning paradigms that advance the state of the art in computer-based learning systems and support students' abilities to learn, even after they leave the computer environment. To achieve this, we have adopted a *learning by teaching* paradigm where students teach computer agents. This paper discusses the design and implementation of a teachable agent system, Betty's Brain, and reports the results of an experiment that manipulated

Biswas, G., Leelawong, K., Belynne, K., et al., Developing Learning by Teaching Environments that support Self-Regulated Learning. in *The seventh International Conference on Intelligent Tutoring Systems*, (Maceió, Brazil, 2004), forthcoming.

the metacognitive support students received when teaching the agent to determine its effects on the students' abilities to subsequently learn new content several weeks later.

Studies of expertise have shown that knowledge needs to be connected and organized around important concepts, and these structures should support transfer to other contexts. Other studies have established that improved learning happens when the students take control of their own learning, develop metacognitive strategies to assess what they know, and acquire more knowledge when needed. Thus the learning process must help students build new knowledge from existing knowledge (constructivist learning), guide students to discover learning opportunities while problem solving (exploratory learning), and help them to define learning goals and monitor their progress in achieving them (metacognitive strategies).

The cognitive science and education research literature supports the idea that teaching others is a powerful way to learn. Research in reciprocal teaching, peer-assisted tutoring, small-group interaction, and self-explanation hint at the potential of learning by teaching [3,4]. The literature on tutoring has shown that tutors benefit as much from tutoring as their tutees [5]. Biswas et al. [6] report that students preparing to teach made statements about how the responsibility to teach forced them to gain deeper understanding of the materials. Other students focused on the importance of having a clear conceptual organization of the materials.

Teaching is a problem solving activity [7]. Learning-by-teaching is an open-ended and self-directed activity, which shares a number of characteristics with exploratory and constructivist learning. A natural goal for effective teaching is to gain a good understanding of domain knowledge before teaching it to others. Teaching also includes a process for structuring knowledge in communicable form, and reflecting on interactions with students during and after the teaching task [5]. Good learners bring structure to a domain by asking the right questions to develop a systematic flow for their reasoning. Good teachers build on the learners' knowledge to organize information, and in the process, they find new knowledge organizations, and better ways for interpreting and using these organizations in problem solving tasks. From a system design and implementation viewpoint, this brings up an interesting question: "How do we design learning environments based on the learning by teaching paradigm?" This has led us to look more closely at the work on pedagogical and intelligent agents as a mechanism for modeling and analyzing student-teacher interaction.

2 Learning by Teaching: Previous Work

Intelligent agents have been introduced into learning environments to create better and more human-like support for exploratory learning and social interactions between tutor and tutee [8,9]. Pedagogical agents are defined as "animated characters designed to operate in an educational setting for supporting and facilitating learning" [8]. The agent adapts to the dynamic state of the learning environment, and it makes the user aware of learning opportunities as they arise, much as human mentor can. Agents use speech, animation, and gestures to extend the traditional textual mode of interaction, and this may increase students' motivation and engagement. They can gracefully combine individualized and collaborative learning, by allowing multiple students and their agents to interact in a shared environment [10]. However, the locus of control stays with the intelligent agent, which plays the role of the teacher or tutor.

Recently, there have been efforts to implement the learning by teaching paradigm using agents that learn from examples, advice, and explanations provided by the student-teacher [11]. A primary limitation of these systems is that the knowledge structures and reasoning mechanisms used by the agent are not made visible to the student, therefore, they find it difficult to uncover, analyze, and learn from their interactions with the agent. Moreover, some of the systems provide outcome feedback or no feedback at all. It is well known that outcome feedback is less effective in supporting learning and problem solving than *cognitive feedback* [12].

On the positive side, students like interacting with these agents. Some studies showed increased motivation but it was not clear that this approach helped achieve deep understanding of complex domain material. We have adopted a new approach to designing learning by teaching environments that supports constructivist and exploratory activities, and at the same time suggests the use of metacognitive strategies to promote learning that involves deep understanding and transfer.

3 A New Approach: Betty's Brain

Betty's Brain provides important visual structures that are tailored to a specific form of knowledge organization and inference to help shape the thinking of the *learner-as-teacher*. In general, our agents try to embody four principles of design: (i) they teach through visual representations that organize the reasoning structures of the domain; (ii) they build on well-known teaching interactions to organize student activity; (iii) they ensure the agents have independent performances that provide feedback on how well they have been taught, and (iv) they keep the start-up costs of teaching the agent very low (as compared to programming). This occurs by only implementing one modeling structure with its associated reasoning mechanisms.

Betty's Brain makes her qualitative reasoning visible through a dynamic, directed graph called a *concept map* [13]. The fact that the TA environments represent knowledge structures rather than the referent domain is a departure from many simulation-based learning environments. Simulations often show the behavior of a process, for example, how an algal bloom increases the death of fish. On the other hand, TAs simulate the behavior of a person's thoughts about a system. Learning empirical facts is important, but learning to use the expert structure that organizes those facts is equally important. Therefore, we have structured the agents to simulate particular forms of thought that help teacher-students structure their thinking about a domain.

Betty's Brain is designed to teach middle school students about the concepts of interdependence and balance in river ecosystems [6,14]. Fig. 1 illustrates the interface of Betty's Brain. Students explicitly teach Betty, using the *Teach Concept*, *Teach Link* and *Edit* buttons to create and modify their concept maps in the top pane of the window. Once taught, Betty can reason with her knowledge and answer questions. Users can formulate queries using the *Ask* button, and observe the effects of their teaching by analyzing Betty's answers. Betty provides explanations for her answers by depicting the derivation process using multiple modalities: text, animation, and speech. Betty uses qualitative reasoning to derive her answers to questions through a chain of causal inferences. Details of the reasoning and explanation mechanisms in Betty's Brain are presented elsewhere [15].

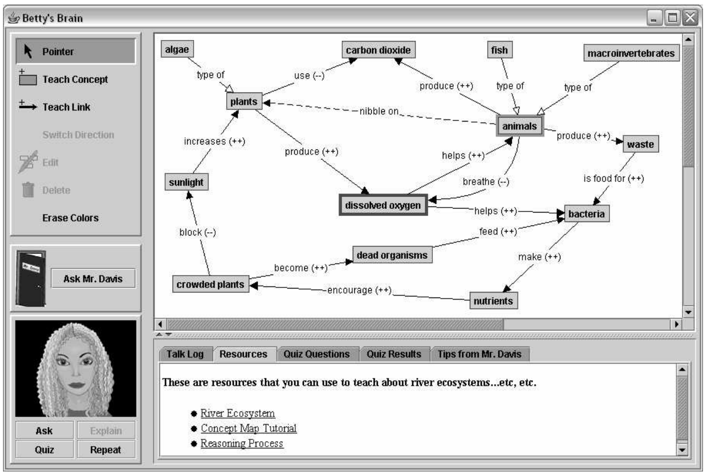

Figure 1: Betty's Brain

The visual display of the face with animation in the lower left is one way in which the user interface attempts to provide engagement and motivation to users by increasing social interactions with the system. We should clarify that Betty does not use machine learning algorithms to achieve automated learning. Our focus is on the well-defined schemas associated with teaching that support a process of instruction, assessment, and remediation. These schemas help organize student interaction with the computer.

To accommodate students who are novices in the domain knowledge and in teaching, the learning environment provides a number of scaffolds and feedback mechanisms. The scaffolds are in the form of well-organized online resources, structured quiz questions that support users in systematically building their knowledge, and Mentor feedback that is designed to provide hints on domain concepts along with strategies on how to learn and how to teach. We adopted the framework of *self-regulated learning*, described by Zimmerman [16] as situations where students are "*metacognitively, motivationally, and behaviorally participants in their own learning process.*" Self-regulated learning strategies involve actions and processes that can help one to acquire knowledge and develop problem-solving skills [17]. Zimmerman describes a number of self-regulated learning skills that include goal setting and planning, seeking information, organizing and transforming, self-consequating, keeping records and monitoring, and self-evaluation. We redesigned the characteristics of both Betty and the Mentor agent to help users develop these skills as they teach and learn.

This has produced a number of unique characteristics in the learning environment. For example, when a student begins the teach phase by constructing the initial concept map, both the Mentor and Betty make suggestions that the student *set goals* on what to teach, and make efforts to gain the relevant knowledge by studying the river ecosystem resources. The Mentor continues to emphasize the reading and understanding of resources, whenever students have questions on *how to improve their learning*. The user is given the opportunity to *evaluate her knowledge* while studying. If she is not satisfied with her understanding, she may *seek further information* by asking the Mentor for additional help. While teaching, the student as teacher can interact with Betty in many ways, such as asking her questions (*querying*), and getting her to take *quizzes* to evaluate her performance. Users are given a chance to predict how Betty will answer a question so they can check what Betty learned against what they were trying to teach.

Some of the self-regulation strategies manifest through Betty's persona. These strategies make Betty more involved during the teach phase, and drive her interactions and dialog with the student. For example, during concept map creation, Betty spontaneously tries to demonstrate *chains of reasoning*, and the conclusions she draws from this reasoning process. She may query the user, and sometimes remark (right or wrong) that an answer she has derived does not seem to make sense. This is likely to make users reflect on what they are teaching, and perhaps, like good teachers they will assess Betty's learning progress more often. At other times, Betty will prompt the user to *formulate queries* to check if her reasoning with the concept map produces correct results. There are situations when Betty emphatically refuses to take a quiz because she feels that she has *not been taught enough*, or that the student has not given her *sufficient practice by asking queries* before making her take a quiz.

After Betty takes a quiz offered by the Mentor agent, she discusses the results with the user. Betty reports: (i) her view of her performance on the particular quiz, and if her performance has improved or deteriorated from the last time she took the quiz, and (ii) the Mentor's comments on Betty's performance in the quiz, such as: "*Hi, I'm back. I'm feeling bad because I could not answer some questions in the quiz. Mr. Davis said that you could ask him more about river eco-systems.*" The Mentor agent's initial comments are general, but they become more specific (*"You may want to study the role of bacteria in the river"*) if errors persist, or if the student seeks further help. Specific mentor feedback explains *chains of events* to help students better understand Betty's reasoning processes. The online resources are structured to make explicit the concepts of interdependence and balance. A hypertext implementation and an advanced keyword search technique provide easy access to information.

Overall, we believe that the introduction of self-regulation strategies provides the right scaffolds to help students learn about a complex domain, while also developing metacognitive strategies that promote deep understanding, transfer, and life-long learning. All this is achieved in an exploratory environment, with the student primarily retaining the locus of control. Only when the student seems to be hopelessly stuck, does the Mentor intervene with specific help.

4 A Computational Architecture for Betty's Brain

With time, as we refined the system, it became clear that an incremental, modularized design strategy was required to keep to a minimum the changes to be made to the code as and when we felt the need to further refine the system. We turned to multi-agent architectures to achieve this goal. The current multi-agent architecture in Betty's Brain is organized into four agents: the teachable agent, Betty, the mentor agent, Mr. Davis, and two auxiliary agents, the student agent and the environment agent. The last two agents help achieve greater flexibility by making it easier to update the scenarios in which the agents operate without having to recode the communication protocols. The student agent represents the interface of the student teacher into the system. It provides facilities that allow the user to manipulate environmental functions and to teach the teachable agent.

All agents interact through the Environment Agent, which acts as a "Facilitator." This agent maintains information about the other agents and the services they provide. When an agent sends a request to the Environment Agent, it decomposes the request if different parts are to be handled by different agents and sends them to the respective agents, and translates the communicated information to match an agent's vocabulary.

A variation of the FIPA ACL agent communication language [18] is used for agent communication. Each message sent by an agent contains a description of the message, message sender, recipient, recipient class, and the actual content of the message. Communication is implemented using a Listener interface, where each agent listens only for messages from the Environment Agent and the Environment Agent listens for messages from all other agents.

The system is implemented using a generic agent architecture shown in Fig. 2. Each agent has a Monitor, Decision Maker, Memory, and an Executive. The Monitor listens for events from the environment, and using a pattern tracker, converts them to the appropriate format needed by the agent. The decision maker, the agent's brain, contains two components: the *reasoner* and the *emotion generator*. It performs reasoning tasks (e.g., answering questions) and updates on the state of the agent. The Executive posts

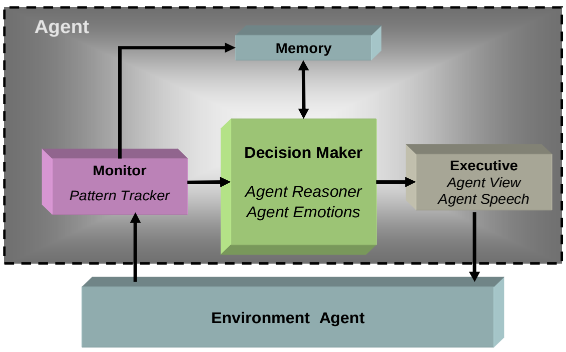

Figure 2: Agent Components

multimedia (speech, text, graphics, animation) information from an agent to the environment. This includes the agent's answer to a question, explanation of an answer and other dialog with the user. The Executive is made up of Agent Speech and Agent View, which handle speech and visual communication, respectively.

5 Experiments

An experiment was designed for fifth graders in a Nashville school to compare three different versions of the system. The version 1 baseline system (ITS) did not involve any teaching. Students interacted with the mentor, Mr. Davis, who asked them to construct concept maps to answer three sets of quiz questions. The quiz questions were ordered to meet curricular guidelines. When students submitted their maps for a quiz, Mr. Davis, the pedagogical agent, provided feedback based on errors in the quiz answers, and suggested how the students may correct their concept maps to improve their performance. The students taught Betty in the version 2 and 3 systems. In the version 2 (LBT) system, students could ask Betty to take a quiz after they taught her, and the mentor provided the same feedback as in the ITS system. Here the feedback was given to Betty because she took the quiz. The version 3 (SRL) system had the new, more responsive Betty with self-regulation behavior (section 3), and a more extensive mentor agent, who provided help on how to teach and how to learn in addition to domain knowledge. But this group had to explicitly query Mr. Davis to receive any feedback. Therefore, the SRL condition was set up to develop more active learners by promoting the use of self-regulation strategies. The ITS condition was created to contrast learning by teaching environments from tutoring environments. The two other groups, LBT and SRL, were told to teach Betty and help her pass a test so she could become a member of the school Science club. Both groups had access to the query and quiz features. All three groups had access to identical resources on river ecosystems, the same quiz questions, and the same access to the Mentor agent, Mr. Davis.

The two primary research questions we set out to answer were:

1. Are learning by teaching environments more effective in helping students to learn independently and gain deeper understanding of domain knowledge than pedagogical agents? More specifically, would LBT and SRL students gain a better understanding of interdependence and balance among the entities in river ecosystems than ITS students? Further, would SRL students demonstrate deeper understanding and better ability in transfer, both of which are hallmarks of effective learning?

2. Does self-regulated learning enhance learning in learning by teaching environments? Self-regulated learning should be an effective framework for providing feedback because it promotes the development of higher-order cognitive skills [17] and it is critical to the development of problem solving ability [13]. In addition, cognitive feedback is more effective than outcome feedback for decision-making tasks [10]. Cognitive feedback helps users monitor their learning needs (achievement relative to goals) and guides them in achieving their learning objectives (cognitive engagement by applying tactics and strategies).

Experimental Procedure

The fifth grade classroom in a Nashville Metro school was divided into three equal groups of 15 students each using a stratified sampling method based on standard achievement scores in mathematics and language. The students worked on a pretest with twelve questions before they were separately introduced to their particular versions of the system. The three groups worked for six 45-minute sessions over a period of three weeks to create their concept maps. All groups had access to the online resources while they worked on the system.

At the end of the six sessions, every student took a post-test that was identical to the pretest. Two other delayed posttests were conducted about seven weeks after the initial experiment: (i) a *memory test*, where students were asked to recreate their ecosystem concept maps from memory (there was no help or intervention when performing this task), and (ii) a *preparation for future learning transfer test*, where they were asked to construct a concept map and answer questions about the land-based nitrogen cycle. Students had not been taught about the nitrogen cycle, so they would have to learn from resources during the transfer phase.

In this study, we focus on the results of the two delayed tests, and the conclusions we can draw from these tests on the students' learning processes. As a quick review of the initial learning, students in all conditions improved from pre- to posttest on their knowledge of interdependence (p*'s*<.01, paired T-tests), but not in their understanding of ecosystem balance. There were few differences between conditions in terms of the quality of their maps (the LBT and SRL groups had a better grasp of the role of bacteria in processing waste at posttest). However, there were notable differences in their use of the system during the initial learning phase.

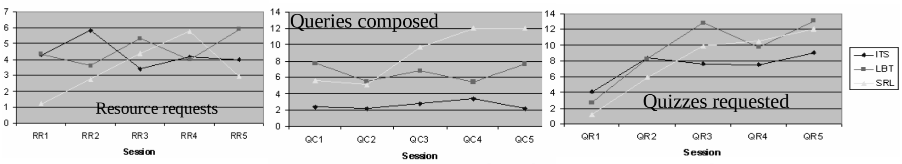

Figure 3: Resource Requests, Queries Composed, and Quizzes Requested per session

Fig. 3 shows the average number of resource, query, and quiz requests per session by the three groups. It is clear from the plots that the SRL group made a slow start as compared to the other two groups. This can primarily be attributed to the nature of the feedback; i.e., the ITS and LBT groups received specific content feedback after a quiz, whereas the SRL group tended to receive more generic feedback that focused on self-regulation strategies. Moreover, in the SRL condition, Betty would refuse to take a quiz unless she felt the user had taught her enough, and prepared her for the quiz by asking questions. After a couple of sessions the SRL group showed a surge in map creation and map analysis activities, and their final concept maps and quiz performance were comparable to the other groups. It seems the SRL group spent their first few sessions in learning self-regulation strategies, but once they learned them their performance improved significantly. Table 1 presents the mean number of expert concepts and expert causal links in the student maps for the delayed memory test. Results of an ANOVA test on the data, with Tukey's LSD to make pairwise comparisons showed that the SRL group recalled significantly more links that were also in the expert map (which nobody actually saw).

Table 1: Results of the Memory Test

Student Map Included:	SRL Mean (se)	LBT Mean (se)	ITS Mean (se)
Expert Concepts	6.7 (.6)	6.4 (.5)	5.8 (.6)
Expert Causal Links	3.3[a] (.6)	1.7 (.6)	2.0 (.6)

[a] Significantly greater than LBT, $p < .05$;

We thought that the effect of SRL would not be to improve memory, but rather to provide students with more skills for learning subsequently. When one looks at the results of the transfer task in the test on preparation for future learning, the differences between the SRL group and the other two groups are significant. Table 2 summarizes the results of the transfer test, where students read resources and created a concept map for the land-based nitrogen cycle with very little help from the Mentor agent (and which they had not studied previously). The Mentor agent's only feedback was on the correctness of the answers to the quiz questions. All three groups received the same treatment. There are significant differences in the number of expert concepts in the SRL versus ITS group maps, and the SRL group had significantly more expert causal links than the LBT and ITS groups. The effects of teaching self-regulation strategies had an impact on the students' abilities to learn a new domain.

Table 2: Results of the Transfer Study

Student Map Included:	SRL Mean (sd)	LBT Mean (sd)	ITS Mean (sd)
Expert Concepts	6.1[a] (.6)	5.2 (.5)	4.1 (.6)
Expert Causal Links	1.1[ab] (.3)	0.1 (.3)	0.2 (.3)

[a] Significantly greater than ITS, $p < .05$; [b] Significantly greater than LBT, $p < .05$

6 Conclusions

The results demonstrate the significant positive effects of SRL strategies in understanding and transfer in a learning by teaching environment. We believe that the differences between the SRL and the other two groups would have been even more pronounced if the transfer test study had been conducted over a longer period of time. Last, we believe that the concept map and reasoning schemes have to be extended to include temporal reasoning and cycles of behavior to facilitate students' learning about the concept of balance in ecosystems.

Acknowledgements
This work has been supported by a NSF ROLE Award #0231771. The assistance provided by the Teachable Agents group, especially John Bransford and Nancy Vye are gratefully acknowledged.

References

[1] Wenger, E. (1987). Artificial Intelligence and Tutoring Systems. Los Altos, California: Morgan Kaufmann Publishers.
[2] Brusilovsky, P. (1999). Adaptive and Intelligent Technologies for Web-based Education, Special Issue on Intelligent Systems and Teleteaching, C. Rollinger and C. Peylo (eds.), 4: 19-25.
[3] Palinscar, A. S. & Brown, A. L. (1984). Reciprocal teaching of comprehension-fostering and comprehension -monitoring activities. Cognition and instruction, 1: 117-175.
[4] Chi, M. T. H., De Leeuw, N., Mei-Hung, C., & Levancher, C. (1994). Eliciting self explanations. Cognitive Science, 18: 439-477.
[5] Chi, M. T. H., et al. (2001). "Learning from Human Tutoring." Cognitive Science 25(4): 471-533.
[6] Biswas, G., Schwartz, D., Bransford, J., & The Teachable Agents Group at Vanderbilt University. (2001). Technology Support for Complex Problem Solving: From SAD Environments to AI. In Forbus & Feltovich (eds.), Smart Machines in Education, 71-98. Menlo Park, CA: AAAI Press.
[7] Artzt, A. F. and E. Armour-Thomas (1999). "Cognitive Model for Examining Teachers' Instructional Practice in Mathematics: A Guide for Facilitating Teacher Reflection." Educational Studies in Mathematics 40(3): 211-335.
[8] G. Clarebout, J. Elen, W. L. Johnson, and E. Shaw. (2002). "Animated Pedagogical Agents: An Opportunity to be Grasped?" Journal of Educational Multimedia and Hypermedia, 11: 267-286.
[9] Johnson W., Rickel, J.W., and Lester J.C. (2001). "Animated Pedagogical Agents: Face-to-Face Interaction in Interactive Learning Environments", International Journal of Artificial Intelligence in Education11: 47-78
[10] Moreno, R. & Mayer, R. E. (2002). Learning science in virtual reality multimedia environments: Role of methods and media. Journal of Educational Psychology, 94: 598-610.
[11] Nichols, D. M. (1994). Intelligent Student Systems: an application of viewpoints to intelligent learning environments, Ph.D. thesis, Lancaster University, Lancaster, UK.
[12] Butler, D. L. and P. H. Winne (1995). "Feedback and Self-Regulated Learning: A Theoretical Synthesis." Review of Educational Research 65(3): 245-281.
[13] Novak, J.D. (1996). Concept Mapping as a tool for improving science teaching and learning, in Improving Teaching and Learning in Science and Mathematics, D.F. Treagust, R. Duit, and B.J. Fraser, eds. Teachers College Press: London. 32-43.
[14] K. Leelawong, K., et al. (2003), "Teachable Agents: Learning by Teaching Environments for Science Domains," Proc. 15^{th} Innovative Applications of Artificial Intelligence Conf., Acapulco, Mexico, 109-116.
[15] Leelawong, K., Y. Wang, et al. (2001). Qualitative reasoning techniques to support learning by teaching: The Teachable Agents project. 15^{th} International Workshop on Qualitative Reasoning, San Antonio, Texas, AAAI Press, 73-80.
[16] Zimmerman, B. J. (1989). "A Social Cognitive View of Self-Regulated Academic Learning." Journal of Educational Psychology 81(3): 329-339.
[17] Pintrich, P. R. and E. V. DeGroot (1990). "Motivational and self-regulated learning components of classroom academic performance." Journal of Educational Psychology 82: 33-40.
[18] Labrou. Y, T. Finin and Peng, Y. (1999) "Agent Communication Languages: The Current Landscape", IEEE Intelligent Systems, 14(2): 45-52.