Socially Intelligent Agents to Improve the Effectiveness of Educational Games

Cristina Conati, Maria Klawe

Department of Computer Science
University of British Columbia
2366 Main Mall
Vancouver, BC, V6T1Z4, CANADA,
conati@cs.ubc.ca, klawe@interchange.ubc.ca

Abstract

We describe preliminary research on devising intelligent agents that can improve the educational effectiveness of collaborative, educational computer games. We illustrate how these agents can overcome some of the shortcomings of educational games by explicitly monitoring how students interact with the games, by modeling both the students' cognitive and emotional states, and by generating calibrated interventions to trigger constructive reasoning and reflection when needed.

Introduction

In this paper, we explore the potential of enriching educational computer games with socially intelligent agents that can help students learn effectively from the games while maintaining the high level of engagement and motivation that constitutes the strong appeal of electronic games in non-educational settings. Our research is developed in the context of EGEMS, the Electronic Games for Education in Math and Science project at the University of British Columbia (UBC). EGEMS is an interdisciplinary project that aims to explore the potential of specially designed computer and video games in mathematics and science education for students aged 9 – 13. Several authors have suggested the potential of video and computer games as educational tools (e.g. [5, 14, 21]), but little concrete research has been devoted to how to turn this potential into reality. As a matter of fact, studies that have been performed within the EGEMS project to test the effectiveness of educational games have shown that, while they are almost always highly engaging, they often do not trigger the constructive reasoning and reflection necessary for learning [12].

For example, the empirical evaluation of Counting on Frank, a commercial mathematical computer game designed to help students learn how to solve algebra word problems, showed that when learning with Counting on Frank was compared with learning through more traditional spreadsheet-based exercises, the second modality was more effective. This despite the fact that, although the game included especially entertaining animations and click-ons that could have distracted the students, most of them spent a lot of time on the mathematical activities in the game. The empirical evaluation of Super Tangram, a computer game based on tangram puzzles to increase students' understanding of two-dimensional geometric transformations, showed that game effectiveness varied greatly with the conditions of testing, such as the presence of an active teacher during the evaluation sessions, the frequency of these sessions and the degree of connection with other geometry activities carried on in the classroom. Similar results were achieved with the evaluation of Phoenix Quest, a game that uses a story, mathematical puzzle activities embedded in the story and interactive communication between the player and the story characters to give students practice on fractions, ratios, negative numbers, coordinates, graph algorithms and number sequences. The game greatly enhanced learning when it was coupled with supporting classroom activities such as related pencil and paper worksheets, group discussions and journal writing. However, without these supporting activities, despite enthusiastic game playing, the learning increase was much more modest.

These results indicate that, although educational computer games can highly engage students in activities involving the targeted educational skills, such engagement, by itself, is often not enough to fulfill the learning and instructional needs of students. This could be due to several reasons.

One reason could be that even the most carefully designed game fails to make students reflect on the underlying domain knowledge and constructively react to the learning stimuli provided by the game. Insightful learning requires meta-cognitive skills that foster conscious reflection upon one's problem solving and performance [2, 4, 24], but reflective cognition is hard work. Possibly, the high level of engagement triggered by the game activities act as a distraction from reflective cognition, when the game is not integrated with external activities that help ground the game experience into the learning one.

A second reason might be that, while learning through educational games and free exploration suit students that already possess the meta-cognitive skills necessary to learn from these activities (such as self-monitoring, self-questioning and self-explanation [2]), they are not sufficient to trigger learning for those students that do not possess these skills.

A third reason could be that the exploratory nature of games requires more time to achieve the same amount of learning generated by more guided and focused educational activities.

Socially Intelligent Agents as Mediators in Educational Games

Although there is insufficient evidence to understand to what extent each of the above reasons limits the effectiveness of educational games, we believe that all three have an effect and we argue that this effect can be greatly reduced by enabling the games to

- explicitly monitor how students interact with and learn from the games; and
- generate calibrated interventions to trigger constructive reasoning and reflection when needed.

However, this must be done without interfering with the factors that make games fun and enjoyable, such as a feeling of control, curiosity, triggering of both intrinsic and extrinsic fantasies, and challenge [13]. Thus, it is not sufficient to provide educational games with the knowledge that makes more traditional Intelligent Tutoring Systems effective for learning: an explicit representation of the target cognitive skills, of pedagogical knowledge and of the student's cognitive state [20]. It is fundamental that the educational interventions be delivered within the spirit of the game, by characters that

- are an integral part of the game plot,
- are capable of detecting the student's frustration, boredom and lack of learning, and
- know how to effectively intervene to correct these negative emotional and cognitive states.

Basically, these characters must play, in the context of the game, the mediating role that teachers and external instructional activities have played during the most successful evaluations of the EGEMS prototypes. The requirement that these agents be socially intelligent is further enforced by the fact that we are currently interested in investigating the educational potential of multi-player networked computer games to support collaborative learning.

Socially Intelligent Agents to Support Game-Based Collaborative Learning

Effective collaborative interaction with peers has proven a successful and uniquely powerful learning method [1, 6]. Students learning effectively in groups encourage each other to ask questions, articulate, explain and justify their opinions, and elaborate and reflect upon their knowledge. However, effective group interaction does not just magically happen. Extensive research on collaborative learning has shown that successful collaboration depends upon a number of factors: the composition of the group, the features of the task, the medium available for communication, the roles that the group members play during the interaction, the level of constructive conflict that the interaction triggers, and the availability of more knowledgeable group members to help the less knowledgeable ones [6, 23]. Some of these factors (such as the composition of the group), need to be taken into account when creating the groups. Others can be enforced during the interaction if there is a human or artificial agent that oversees the collaboration process and detects when the conditions for effective collaboration are not met. We are working on creating artificial agents that can provide this mediating role within multi-player, multi-activity educational games designed to foster learning through collaboration. As a test-bed for our research we are using Avalanche, one of the EGEMS prototype games, in which four players assume the roles of the leading citizens in a mountain ski town, and work together to deal with the problems caused by a series of avalanches. The current set of Avalanche activities include:

- Meet with all the team members in the Town Hall and collaboratively answer questions during a press conference. The questions target central points of the game goals and rules. This activity aims at forcing the students to understand the game before they start playing.
- Prime Climb: teams of two players climb ice-faces by selecting numbers relatively prime to those occupied by other team members (see Figure 1), in order to collect the data needed to evaluate the critical mountain zones that are possible sites for the next avalanche. Before climbing a mountain, the players need to get certified for that mountain climbing level, by doing practice climbs with an instructor known as "Cool Guy" (the character to the right in Figure 1).
- Zone Size: players estimate the area and volume of snow in the critical zones using the data they have obtained from climbing the mountains and from maps.
- Snow Release: based on the information from the Zone Size activity, players decide where to place their limited set of explosives and barriers to release snow from the

most dangerous zones without releasing so much snow that another avalanche is triggered.

Preliminary pilot studies, which focused mainly on the first two activities above, have shown that the collaborative nature of the game triggers a tremendous level of engagement in the students. However, they also uncovered the following problems.

Figure 1: Prime Climb activity in Avalanche

First, given the complexity of the game and the large amount of available activities, students often lose track of the game goals and of the means available to achieve them. Although an on-line hypertext help is available, students almost never look at it. They also tend not to read the fixed set of instructions that are given by game characters upon entering some of the game activities.

Second, students eventually learn how to play the game, but they do not necessarily learn the target instructional knowledge. For example, we observed that, during Prime Climb, students who did not know about common factors generally did not learn from the fixed sequence of instructions they received from Cool Guy before they started climbing (see Figure 1). Hence, as they were climbing they often fell. When they tried to find out why, the information was not easily available. Thus, these students ended up climbing by trial and error, exhaustively trying hexes until they made it to the top. The falling did not discourage this behavior. On the contrary, students seemed to enjoy seeing their character fall, because the animation is quite cute. The problem here is twofold: the game fails to provide easy access to the information that may help students achieve the educational objectives and allows students to progress toward the game objectives while sidestepping the educational ones.

Third, the game at times fails to trigger effective collaboration. For instance, some students that seemed to be shyer and less familiar with the other group members, ended up being more and more isolated during the interaction. When they had problems and questions, none of the other players seemed to pay attention or tried to help them. Other students focused on competition rather than collaboration. For instance, during Prime Climb they were focusing on getting to the top of a mountain before their hiking partner or on getting more climbing certificates, even when another player was waiting for them to climb a mountain they were already certified for.

A Comprehensive Computational Model of Effective Collaborative Learning

The above examples clearly show that Avalanche can greatly benefit from the addition of socially intelligent agents that help students find their way through the game, trigger constructive learning and reflection, and help mediate and structure the collaborative interaction. The game already uses a few characters to deliver canned instructions, like Cool Guy in Figure 1, but as we said earlier, students pay little attention to these instructions and seldom benefit from them. We are working on enriching these and additional characters with:

- explicit models of the educational activities they are associated with, of the emotional states that can influence learning from these activities and of effective collaborative interaction.
- the capability of modeling, from the interaction with the game, the cognitive and meta-cognitive skills of the players, along with their emotional states and the status of the collaborative interaction;
- the capability of making intelligent decisions as to when and how to intervene to improve effective collaboration and learning, without compromising the level of motivation and engagement fueled by the game.

Architecture

Figure 2 sketches our proposed general architecture underlying the functioning of socially intelligent characters in the context of a multi-player, multi-activity educational game. As students engage in the different activities available within the game, their behavior is monitored by the agents currently involved in the interaction, through their Behavior Interpreters. Each Behavior Interpreter specializes in interpreting actions related to a specific student's behavior (e.g., behavior related to game performance, meta-cognitive skills, collaboration and emotional reaction) and updates the corresponding elements in the student model of the player its agent is interacting with.

A Game Actions Interpreter, for instance, processes all the student's game actions within a specific activity, which then can be used to infer information on the student cognitive and meta-cognitive skills.

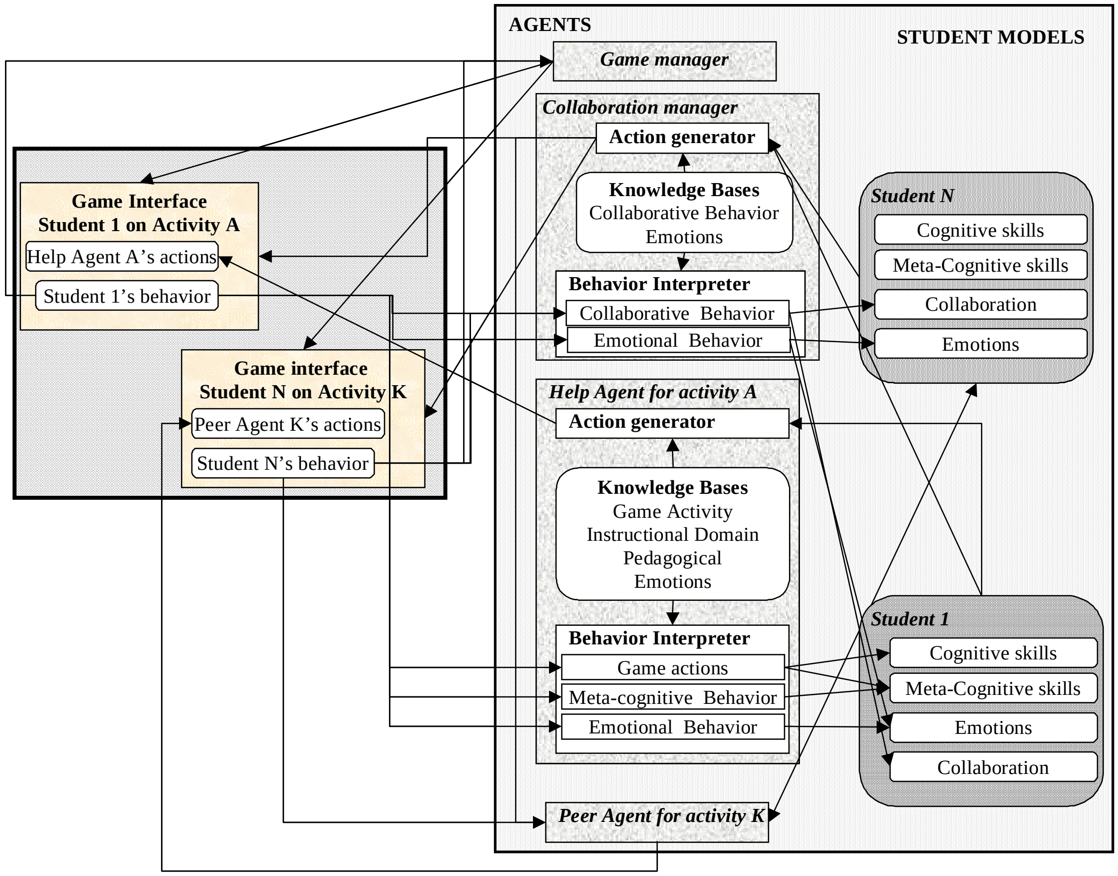

Figure 2: Architecture for multi-player, multi-activity educational games based on socially intelligent agents

A Meta-Cognitive Behavior Interpreter tracks all those student's actions beside game actions that can indicate meta-cognitive activity, such as utterances and eye or mouse movements expressing reflection and exploration, and passes them to the student model as further evidence of the student's meta-cognitive skills.

Each agent then uses the information in the student model and the expertise encoded in its knowledge base (that depends on that agent's role in the game) to generate actions that help the student learn better from the activity she is involved in.

The agents in the architecture include a Game Manager, the Collaboration Manager and agents related to specific game activities (like Help Agent for activity A and Peer Agent for activity K in Figure 2). The Game Manager knows about the structure of the game and guides the students through its activities. The Collaboration Manager is in charge of orchestrating effective collaborative behavior. As shown in Figure 2, its Behavior Interpreter captures and decodes all those students' actions that can indicate collaboration or lack thereof, along with the related emotional reactions. The actions that pertain to the Collaboration Manager include selecting adequate collaboration roles and partners for a student within an activity selected by the student himself or suggested by the Game Manager. The pool of partners from which the Collaboration Manager can select collaboration partners includes both the other players or the artificial agents available for that activity. Thus, if no other player can currently be an adequate partner for a student, because of incompatible cognitive or emotional states, the Collaboration Manager can suggest an artificial agent as a partner (like, in Figure 2, the Peer Agent selected for Student N in activity K).

The artificial agents related to each game activity have expertise that allow them to play specific roles within that activity. So, for instance, a Help agent (like Help Agent for activity A in Figure 2) will have expert knowledge on the activity actions and on the underlying instructional domain, along with knowledge of the emotional states that can influence the benefits of providing help and pedagogical knowledge on how to provide this help effectively. This knowledge may include information on how to trigger constructive learning through the elicitation of meta-cognitive skills such as self-explanation or self-monitoring, in which case the agent will include a Behavior Interpreter to capture behavior related to these skills. Peer agents, on the other hand, will have game and domain knowledge that is incomplete to different extents, so that they can be selected by the Collaboration Manager to make the student play specific collaborative roles in the activity (such as instructor or learning companion).

Student Models

The student models in our architecture are based on the probabilistic reasoning framework of Bayesian networks [17] that allows performing reasoning under uncertainty by relying on the sound foundations of probability theory. One of the main objections to the use of Bayesian networks is the difficulty of assigning accurate network parameters (i.e. prior and conditional probabilities). However, even when it is the case that the parameters cannot be reliably specified by experts or learned from data, providing estimates for them allows the designer to clearly define the assumptions the model must rely upon and to revise the assumptions by trial and error on the model performance. Thus, we believe that Bayesian networks provide an appropriate formalism to model and integrate in a principled way the multiple sources of uncertainty involved in monitoring a student's cognitive and emotional states, and the unfolding of a collaborative interaction.

Modeling cognitive and meta-cognitive skills. Bayesian networks have been extensively used to build user models representing domain knowledge and cognitive [11]. In [3], we have described how to automatically specify the network structure and conditional probabilities to model the relations between problem solving behavior and domain knowledge. In [4], we have started to extend this work to model learning through meta-cognitive skills that trigger constructive thinking, in particular to model learning of instructional material through self-explanation. We plan to adapt this approach to student models for multi-players educational games, to formalize the probabilistic relationships between game activity, meta-cognitive skills and learning of the target instructional knowledge.

Modeling collaboration. A preliminary Bayesian model of effective collaborative interaction has been proposed by [22]. This model is based on the findings that successful collaborative learning can be achieved by making group members adopt different roles during the collaborative process, such as leader, observer, critic, coach. The model attempts to trace the progress of group members through the different collaborative roles by monitoring the actions that they perform on an interface especially designed to reify each role and the actions that pertain to it. As described earlier, we also adopt a role-based approach to model effective collaboration, but we cannot structure and constrain the game interface as [22] did, because this would compromise the level of fun and engagement that students currently experience with Avalanche. Hence, we need to devise alternative ways to capture the collaborative roles that students adopt during the interaction. We plan to start by making the adoption of different collaborative roles one of the mandatory game activities, orchestrated by the Collaboration Manager. This will reduce the collaboration-monitoring problem to the problem of verifying that students effectively perform the role they have been assigned. We are exploring, among other solutions, the use of speech recognition technology to detect speech acts specific to certain roles. However, as the research proceeds, we hope to also achieve a better understanding of how to monitor and support less constrained collaboration.

Modeling emotions. Since emotional engagement is the element that makes educational games attractive to learners, it is fundamental that this variable be accurately monitored and taken into account by any agent that generates actions in the game. Starting from existing research on the structure of emotions [16], we are working to include in the game a probabilistic formalization of relevant emotional states (such as frustration, boredom and excitement) and their dynamics, as they are influenced by game elements, tutorial interventions, and collaborative interaction. The formalization will also include a mini-theory of how the players' emotions can be detected. We seek to formalize this theory from current research on technology-enhanced affect communication, which explores ways to measure important aspects of emotional information from the human body by using, for example, facial expressions, vocal intonation, galvanic skin response and heart rate [18].

Action Generators

The action generator for each agent in the game relies on a decision-theoretic model of decision-making predicting thst agents act so as to maximize the expected utility of their actions [19]. Other researchers have started adopting a decision theoretic approach to regulate the behavior of interactive agents designed to provide unsolicited help, e.g., to develop desktop assistants that can decide when and how to notify a user of events external to her current task [9, 10] and to help a computer tutor select optimal tutorial actions in the context of coached problem solving [15]. Also, [7] proposes to use the formalism of decision theory to develop principled definitions of emotional states of rational agents, to be used in the context of multi-agent system applications.

In our architecture, the function representing an agent's preferences in terms of utility values depends on the role of the agent in the educational structure of the game. So, for instance, the Collaboration Manager will act so as to maximize students learning as well as their collaborative behavior. A Help agent will act to maximize the student's understanding of a specific activity, while an agent in charge of eliciting a specific meta-cognitive skill will select actions that maximize this specific outcome. All the agents will also include in their utility functions the goal of maintaining the student's level of fun and engagement above a given threshold, although the threshold may vary with the role of the agent. The action generators' decision-theoretic models can be represented as influence diagrams [8], an extension of Bayesian networks devised to model rational decision making under uncertainty. By using influence diagrams, we can compactly specify how the different actions available to the Avalanche agents influence the relevant elements in the student model, such as cognitive and emotional states or the role-based formalization of effective collaboration illustrated above. We can also encode the agent's utility function in terms of these states, thus providing each agent with a normative theory of how to intervene in the students' game playing to achieve the best trade-off between engagement and learning.

Conclusions

We have presented a preliminary architecture to improve the effectiveness of collaborative educational games. The architecture relies on the usage of socially intelligent agents that calibrate their interventions by taking into account not only the students' cognitive states, but also their emotional states and the unfolding of collaborative interactions within the game. We propose to rely on Bayesian networks and influence diagrams to provide our agents with a principled framework for making informed decisions on the most effective interventions under the multiple sources of uncertainty involved in modeling interaction and learning in multi-player, multi-activity educational game.

References

1. Aronson, E. and S. Patnoe, Jigsaw Classroom. 1997, NewYork: Longman.
2. Chi, M.T.H., Self-Explaining: The dual process of generating inferences and repairing mental models. *Advances in Instructional Psychology*, in press.
3. Conati, C., A. Gertner, K. VanLehn and M. Druzdzel. On-line student modeling for coached problem solving using Bayesian networks, in *User Modeling: Proceedings of the Sixth International conference, UM97*, A. Jameson, C. Paris, and C. Tasso, Editors. 1997, Spring Wien: New York.
4. Conati, C. and K. VanLehn, Toward Computer-Based Support of Meta-Cognitive Skills: a Computational Framework to Coach Self-Explanation. *Int. Journal of Artificial Intelligence in Education*, To Appear.
5. David, A. and M.P. Ball, The Video Game: A Model for Teacher-Student Collaboration. *Momentum*, 1986. **17**(1): p. 24-26.
6. Dillenbourg, P., *et al.*, The evolution of research on collaborative kearning, in *Learning in Humans and Machines*, Spada and Reimann, Editors. 1996.
7. Gmytrasiewicz, P. Modeling users emotions during interactive entertainment sessions. in *AAAI Spring Symposium on AI and Interactive Entertainment*. 2000. Stanford, CA.
8. Henrion, M., J. Breese, and E. Horvitz, Decision Analysis and Expert Systems, in *AI Magazine*. 1991.
9. Horvitz, E. Principles of Mixed-Initiative User Interfaces. in *CHI '99, ACM SIGCHI Conference on Human Factors in Computing Systems*. 1999. Pittsburgh, PA.
10. Horvitz, E., A. Jacobs, and D. Hovel. Attention-Sensitive Alerting. in *UAI '99, Conf. on Uncertainty and Artificial Intelligence*. 1999: Morgan Kaufmann Publishers.
11. Jameson, A., Numerical uncertainty management in user and student modeling: An overview of systems and issues. *User Modeling and User-Adapted Interaction,* 1995. **5**.
12. Klawe, M. *When* Does The Use Of Computer Games And Other Interactive Multimedia Software Help Students Learn Mathematics? in *NCTM Standards 2000 Technology Conference*. 1998. Arlington, VA.
13. Malone, T.W., What Makes Things Fun to Learn? A Study of Intrinsically Motivating Computer Games., 1980, Xerox Corporation, Palo Alto Research Center.: Palo Alto, California:.
14. Malone, T.W. Heuristics for designing enjoyable user interfaces: lessons from computer games. in *Human Factors in Computer Systems*. 1982. Gaithersburg, MA.
15. Murray, C. and K. VanLehn. DT Tutor: A Decision-Theoretic, Dynamic Approach for Optimal Selection of Tutorial Actions. in *ITS 2000, 6th International Conference*. 2000. Montreal, CA.
16. Ortony, A., G. Clore, and A. Collins, *Cognitive Structure of Emotions*. 1988: Cambridge University Press.
17. Pearl, J., *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. 1988, San Mateo, CA: Morgan-Kaufmann.
18. Picard, R., *Affective Computing*. 1997: MIT Press.
19. Russell, S. and P. Norvig, *Artificial Intelligence: A Modern Approach*. 1995, Los Altos, CA: Morgan-Kaufman.
20. Shute, V.J. and J. Psotka, Intelligent tutoring systems: Past, Present and Future, in *Handbook of Research on Educational Communications and Technology*, D. Jonassen, Editor. 1996, Scholastic Publications.
21. Silvern, S.B., Classroom Use of Videogames., in *Educational Research Quarterly,*. 1986. p. 10-16.
22. Singley, M. and P. Fairwater. *Team Tutoring Systems: Reifying Roles in Problem Solving.* in *CSCL '99*. 1999. Stanford, CA.
23. Soller, A., F. Linton, and B. Goodman. Toward Intelligent Analysis and Support of Collaborative Learning Interaction. in *Ninth International Conference on Artificial Intelligence in Education*. 1999. LeMans, France.
24. VanLehn, K., *Rule acquisition events in the discovery of problem solving strategies.* Cognitive Science, 1991. **15**(1).