Knowledge Acquisition for Adaptive Collaborative Learning Environments

Amy Soller[1] and Alan Lesgold[1,2]

Intelligent Systems Program[1] and School of Education[2]
Learning Research and Development Center
University of Pittsburgh
3939 O'Hara Street, Pittsburgh, PA 15260-5159
soller@pitt.edu, al@pitt.edu

Abstract

Success in collaboratively learning subject matter means both learning the subject matter (collaborating to learn), and learning how to effectively manage the team interaction (learning to collaborate). Supporting on-line learning teams means supporting both these activities. We focus on the problem of assimilating the knowledge needed to address interaction problems that may arise during collaborative learning sessions. This involves gathering knowledge about the types of problems that learning groups might encounter, evaluating methods for identifying situations in which those problems exist, and implementing strategies to facilitate groups learning on-line.

Introduction

Success in collaboratively learning subject matter means both learning the subject matter (collaborating to learn), and learning how to effectively manage the team interaction (learning to collaborate). The knowledge acquisition process for systems that support collaborative learning warrants a closer look in light of this additional complexity. This paper discusses issues in the process of assimilating the knowledge required to support human-human interaction in networked learning environments.

The EPSILON (Encouraging Positive Social Interaction while Learning ON-Line) Project aims to provide empirically grounded, just-in-time support to learning teams by employing artificial intelligence algorithms to dynamically analyze team members' actions and conversation. A system with the expertise to observe, analyze, and support collaborative learning activities must be a subject matter and social interaction expert, as well as an expert in coaching these skills, and communicating effectively with the learning group. This is a tall order for an AI system! In the next section, we discuss the steps the EPSILON Project team has taken to address a few of the issues.

The Knowledge Acquisition Process

A supportive collaborative learning system that knows how to address interaction problems that may arise during collaborative learning sessions requires knowledge about the types of problems that may arise, knowledge about how to identify episodes of interaction in which those problems exist, and knowledge about how to facilitate groups having those problems. The following three subsections discuss each of these knowledge acquisition phases.

Knowledge about How People Interact while Learning

Models of social interaction provide some assistance in identifying the types of problems that may arise during collaborative learning sessions. These models describe characteristics of effective group interaction, such as idea generation and constructive criticism (Jarboe, 1996), multiple perspectives, viewpoints, and representations (Koschmann, Kelson, Feltovich, and Barrows, 1996), social grounding and shared understanding (Teasley and Roschelle, 1993), positive interdependence (Johnson, Johnson, and Holubec, 1990), and peer tutoring (Webb, 1992).

Models such as these have been successfully applied to classroom practice, however they present a view of effective social interaction at a level of abstraction too high to directly implement in a system. Whereas a classroom teacher facilitating a learning group may be guided by these models in deducing situations in which a group needs explanation or encouragement, a software agent playing the role of a group facilitator must be able to analyze a group's interaction based on the students' communication patterns and their problem solving actions. Earlier work done by the first author and her colleagues (Soller, Goodman, Linton, and Gaimari, 1998) attempted to assimilate the social interaction models described above, and others, to form a comprehensive model that could be directly applied by an adaptive collaborative learning environment. This research uncovered the need for a model operating at an even lower level of abstraction and integrating linguistic and social considerations.

Learning to Evaluate the Effectiveness of Collaborative Learning

Knowledge about how students in effective teams interact is useful to a system only if it can apply this knowledge to recognize specific situations that call for intervention. Classroom teachers learn to analyze and assess student interaction through close observance of group interaction, trial and error, and experience. Developing a system that can analyze group conversation, however, poses its own challenges.

Although research in discourse and pragmatics (Grosz and Sidner, 1986; Mann and Thompson, 1986) provides insight into the nature and structure of communication patterns, the dynamic and, to some extent, unpredictable nature of human interaction deems natural language understanding a difficult, interdisciplinary area of research. Cahn and Brennan (1999), however, explain that a system can represent or model a dialog using only the "gist" of successive contributions; a full account of each contribution, verbatim, is not necessary. Previous work has established promising research directions based on approaches that adopt this idea. These approaches make use of structured communication interfaces that require users to make the intention of their conversational contributions explicit (Baker and Lund, 1996; Flores et al., 1988; McManus and Aiken, 1995; Soller, Linton, Goodman, and Lesgold, 1999).

It would be impossible to enumerate and assess the effectiveness of all possible interaction patterns. Simple statistical, categorical, or rule-based approaches may work well for some domains (including individual computer-based training), however the dynamic nature of human communication and interaction accounts for too many variables to effectively apply such approaches to evaluate collaboration. For this reason, we have begun to explore the use of supervised machine learning, and other artificial intelligence approaches for recognizing the onset and characteristics of effective collaborative learning episodes.

A review of research in the area of machine learning applied to linguistic problems revealed a significant amount of work in determining which cue words and phrases identify specific conversational intentions (e.g. DiEugenio, Moore, and Paolucci, 1997; Samuel, Carberry, and Vijay-Shanker, 1998), and using this information to automatically perform speech act tagging (Stolcke et al., 2000), but very little research in examining the problem of identifying desirable sequences of utterances. Katz, Aronis, and Creitz (1999), however, have successfully applied rule learning techniques to construct hierarchies of increasingly general speech acts and speech act categories from tagged tutorial dialogue.

Barros and Verdejo (1999) rate the asynchronous collaboration between pairs of students along 4 dimensions: initiative, creativity, elaboration, and conformity. For example, making a proposal positively influences initiative and negatively influences conformity. These four attributes, along with others such as the mean number of contributions by team members and the length of contributions factor into a fuzzy inference procedure that rates students' collaboration on a scale from "awful" to "very good". This work is seminal in breaking asynchronous collaboration down into its contributing factors, and computing a fuzzy rubric to evaluate the final product of group interaction. A closer look at sequences of specific learning episodes during group problem solving may help in composing rubrics for dynamically evaluating synchronous collaborative activity, enabling a facilitator agent to provide direction at the most appropriate instances.

An exploratory study performed by the EPSILON team considered the possibility of training a neural network to learn sequences of student interaction that yield the most (and least) productive learning opportunities. Given a classification scale of learning effectiveness (very effective to very ineffective), and a set of classified training data (in the form of speech acts, such as Request or Acknowledge), the backpropagation algorithm correctly assessed the effectiveness of 13/15 sequences of new student dialogues. The weights in the network revealed that interaction sequences comprised mainly of acknowledgement ("OK", "Yes", or "No") yield low learning achievement, and that a high degree of argument and disagreement during learning may contribute toward less effective learning. The network weights also showed that sequences of effective student interaction contain a balance of different types of conversational contributions (e.g. Request, Inform, Mediate). These results are consistent with the statistical analysis presented in (Soller, Linton, Goodman, and Lesgold, 1999).

In a similar exploratory study, we trained Hidden Markov Models (HMMs) to (1) identify the student playing the role of knowledge "sharer" during knowledge sharing episodes, and (2) determine the effectiveness of the episode. A *knowledge sharing episode* is defined as a segment of interaction during which one team member shares new knowledge with the group. The episode is considered effective if one or more students learns the new knowledge (as shown by a difference in pre-post test performance), and ineffective otherwise. The 5 node HMM for selecting the student playing the role of knowledge sharer performed at 100% accuracy (it classified the knowledge sharer correctly in 15/15 sequences), and the 6 node HMM for determining the effectiveness of the episode performed at 92% accuracy. The data was taken from transcripts of five groups of three students each collaborating on-line to solve object oriented analysis and design problems. The preliminary results of this study are promising, however more data is needed to confirm these findings.

We are in the process of surveying the potential of various other machine learning approaches, including decision tree and rule learning algorithms, to learn events indicative of effective knowledge sharing and transfer. Success in recognizing and quantitatively evaluating collaborative learning activities that are initiated by knowledge sharing events strengthens the possibility of an intelligent system that can identify and facilitate other types of collaborative learning events.

Knowledge about Coaching Group Learning

A collaborative learning system, having learned how to recognize episodes of interaction requiring facilitation, must then know what to do in those situations to increase the group's productivity. In some cases, it may be appropriate for the system to intervene and offer advice to an individual or to the group. In other cases, the system may choose to participate in the conversation (Chan and Baskin, 1988; Rich and Sidner, 1996), playing the role of a devil's advocate or motivator.

The Collaborative Learning Model (developed by Soller et al., 1998) describes potential indicators of effective collaborative learning teams based on a review of research in educational psychology and computer-supported collaborative learning (Brown and Palincsar, 1989; Jarboe, 1996; Johnson, Johnson, and Holubec, 1990; McManus and Aiken, 1995; Teasley and Roschelle, 1993; Webb, 1992), and empirical data from the study described in (Soller et al., 1998). For each indicator of effective collaborative learning, the model proposes strategies for promoting effective peer interaction. These strategies include assigning students roles to induce multiple perspectives, initiating and facilitating brainstorming sessions to increase idea generation, and allowing students to view and comment on their student models, which describe the system's assessment of their performance. Further work is needed to determine how to associate various coaching strategies to the situations that trigger them.

Conclusions and Discussion Questions

Students in effective collaborative learning teams are both effective collaborators and effective learners. Supporting on-line learning teams means supporting both these activities. In this paper, we have focused on the problem of assimilating the knowledge needed to support student collaboration. This involves gathering knowledge about the types of problems that may arise in learning groups, evaluating methods for identifying situations in which those problems exist, and implementing strategies to facilitate groups learning on-line.

Models of small group learning provide an understanding of how to address specific group needs, but assessing student interaction to determine the best type of support to provide is a difficult task even for human facilitators. The EPSILON Project team is exploring the possibility of training a supervised machine learning algorithm to distinguish between effective and ineffective sequences of student interaction. Neural Networks and Hidden Markov Models have shown promise in this regard.

We would like to pose the following questions, which we will be addressing while analyzing data from our latest experiment, to the attendees at this symposium. Can important aspects of interaction and learning success be recognized without consideration of the dialog content? What artificial intelligence approaches might be useful in learning to recognize sequences of interaction indicative of effectiveness? Is learning how to collaborate much different from learning how to do other things, and must this "skill" be learned in the context of learning other activities? If so, how does this (if at all) change the process of knowledge acquisition for the training of collaborative activities?

Acknowledgements

This research was supported in part by the Department of Education Technology Innovation Challenge Grant Program.

References

Barros, B., and Verdejo, M.F. (1999). An approach to analyse collaboration when shared structured workspaces are used for carrying out group learning processes. In S. Lajoie and M. Vivet (Eds.) *Proceedings of Artificial Intelligence in Education*. IOS Press, 449-456.

Baker, M. and Lund, K. (1996). Flexibly structuring the interaction in a CSCL environment. *Proceedings of the European Conference on Artificial Intelligence in Education (EuroAIED '96),* 401-407.

Cahn, J. and Brennan, S (1999). A psychological model of grounding and repair in dialog. *Proceedings of the 1999 AAAI Fall Symposium: Psychological Models of Communication in Collaborative Systems*, Cape Cod, MA, 25-33.

Chan, T.W. and Baskin, A.B. (1988). Studying with the prince: The computer as a learning companion. *Proceedings of the ITS '88 Conference*, Montreal, Canada, 194-200.

DiEugenio, B., Moore, J. D., and Paolucci, M. (1997). Learning features that predict cue usage, *Proceedings of the 35th Annual Meeting of the Association for Computational Linguistics*, Madrid, Spain.

Flores, F., Graves, M., Hartfield, B., and Winograd, T. (1988). Computer systems and the design of organizational interaction. *ACM Transactions on Office Information Systems,* 6(2), 153-172.

Grosz, B. and Sidner, C. (1986). Attention, intentions, and the structure of discourse. *Computational Linguistics*, 12(3), 175-204.

Jarboe, S. (1996). Procedures for enhancing group decision making. In B. Hirokawa and M. Poole (Eds.), Communication and Group Decision Making (pp. 345-383). Thousand Oaks, CA: Sage Publications.

Johnson, D., Johnson, R., and Holubec, E.J. (1990). Circles of learning: Cooperation in the classroom (3rd ed.). Edina, MN: Interaction Book Company.

Katz, S., Aronis, J., and Creitz, C. (1999). Modelling pedagogical interactions with machine learning. *Proceedings of the Ninth International Conference on AI in Education*, LeMans, France, 543-550.

Koschmann, T., Kelson, A., Feltovich, P., and Barrows, H. (1996). Computer-supported problem-based learning. In T. Koschmann (Ed.), CSCL: Theory and Practice of an Emerging Paradigm (pp. 83-124). Mahwah, NJ: Lawrence Erlbaum.

Mann, W., and Thompson, S. (1986). Rhetorical Structure Theory: Toward a functional theory of text organization. TEXT, 8(3).

McManus, M. and Aiken, R. (1995). Monitoring computer-based problem solving. *Journal of Artificial Intelligence in Education*, 6(4), 307-336.

Rich, C. and Sidner, C. L. (1996). Adding a collaborative agent to graphical user interfaces. In Proceedings of the Ninth Annual Symposium o User Interface Software and Technology (UIST '96). Seattle, WA.

Samuel, K., Carberry, S, and Vijay-Shanker, K. (1998). Dialogue act tagging with transformation-based learning. *Proceedings of the 17th International Conference on Computational Linguistics*. Montreal, Quebec, Canada, 1150-1156.

Soller, A., Goodman, B., Linton, F., and Gaimari, R. (1998). Promoting effective peer interaction in an intelligent collaborative learning environment. *Proceedings of the Fourth International Conference on Intelligent Tutoring Systems (ITS 98)*, San Antonio, TX, 186-195.

Soller, A., Linton, F., Goodman, B., Lesgold, A. (1999). Toward intelligent analysis and support of collaborative learning interaction. *Proceedings of the Ninth International Conference on Artificial Intelligence in Education*, LeMans, France, 75-82.

Stolcke, A., Ries, K., Coccaro, N., Shriberg, E., Bates, R., Jurafsky, D., Taylor, P., Martin, R., Van Ess-Dykema, C., and Meteer, M. (2000). Dialogue act modeling for automatic tagging and recognition of conversational speech. *Computational Linguistics*, 26(3).

Teasley, S. and Roschelle, J. (1993). Constructing a joint problem space: The computer as a tool for sharing knowledge. In S. Lajoie and S. Derry (Eds.), Computers as cognitive tools (pp. 229-257). Hillsdale, NJ: Lawrence Erlbaum.

Webb, N. (1992). Testing a theoretical model of student interaction and learning in small groups. In R. Hertz-Lazarowitz and N. Miller (Eds.), Interaction in Cooperative Groups: The Theoretical Anatomy of Group Learning (pp. 102-119). New York: Cambridge University Press.