KSC-PaL: A Peer Learning Agent that Encourages Students to take the Initiative*

Cynthia Kersey
Lewis University
Romeoville, IL 60446 USA
kerseycy@lewisu.edu

Barbara Di Eugenio
University of Illinois at Chicago
Chicago, IL 60607 USA
bdieugen@cs.uic.edu

Pamela Jordan and Sandra Katz
University of Pittsburgh
Pittsburgh, PA 15260 USA
pjordan+@pitt.edu
katz+@pitt.edu

Abstract

We present an innovative application of dialogue processing concepts to educational technology. In a previous corpus analysis of peer learning dialogues, we found that initiative and initiative shifts are indicative of learning, and of learning-conducive episodes. We have incorporated this finding in KSC-PaL, a peer learning agent. KSC-PaL promotes learning by encouraging shifts in task initiative.

1 Introduction

Collaborative learning has been shown to be an effective mode of learning for potentially all participants (Brown and Palincsar, 1989; Fisher, 1993; Tin, 2003). While collaboration in dialogue has long been researched in computational linguistics (Chu-Carroll and Carberry, 1998; Constantino-González and Suthers, 2000; Jordan and Di Eugenio, 1997; Soller, 2004), the study of peer learning from a computational perspective is still in the early stages.

Previous research has suggested several mechanisms that explain why peer learning is effective. Among them are: self-directed explaining (Chi et al., 1994), other-directed explaining (Ploetzner et al., 1999; Roscoe and Chi, 2007) and Knowledge Co-construction – KCC for short (Hausmann et al., 2004). KCC episodes are defined as portions of the dialogue in which students are jointly constructing a shared meaning of a concept required for problem solving. This last mechanism is the most interesting from a peer learning perspective because it is a truly collaborative construct and also because it is consistent with the widely accepted constructivist view of learning.

In our previous work (Kersey et al., 2009) we derived a model of peer interactions that operationalizes KCC via the notion of initiative shifts in dialogue. This model was based on an extensive corpus analysis in which we found a strong relationship between initiative shifts and KCC episodes. A paired t-test showed that there were significantly more initiative shifts in the annotated KCC episodes compared with the rest of the dialogue ($t(57) = 3.32$, $p = 0.0016$). The moderate effect difference between the two groups (effect size = 0.49) shows that there is a meaningful increase in the number of initiative shifts in KCC episodes compared with problem solving activity outside of the KCC episodes. Additionally, we found moderate correlations of learning with both KCC ($R^2 = 0.14$, $p = 0.02$) and with initiative shifts ($R^2 = 0.20$, $p = 0.00$).

We have incorporated this model in an innovative peer learning agent, KSC-PaL, that is designed to collaborate with a student to solve problems in the domain of computer science data structures.

2 KSC-PaL

KSC-PaL, has at its core the TuTalk System (Jordan et al., 2007), a dialogue management system that supports natural language dialogue in educational applications. In developing KSC-PaL we extended TuTalk in three ways.

The first extension is a user interface (see Figure 1) which manages communication between TuTalk and the student. Students interact with KSC-

*This work is funded by NSF grants 0536968 and 0536959.

Proceedings of the NAACL HLT 2010: Demonstration Session, pages 17–20,
Los Angeles, California, June 2010.

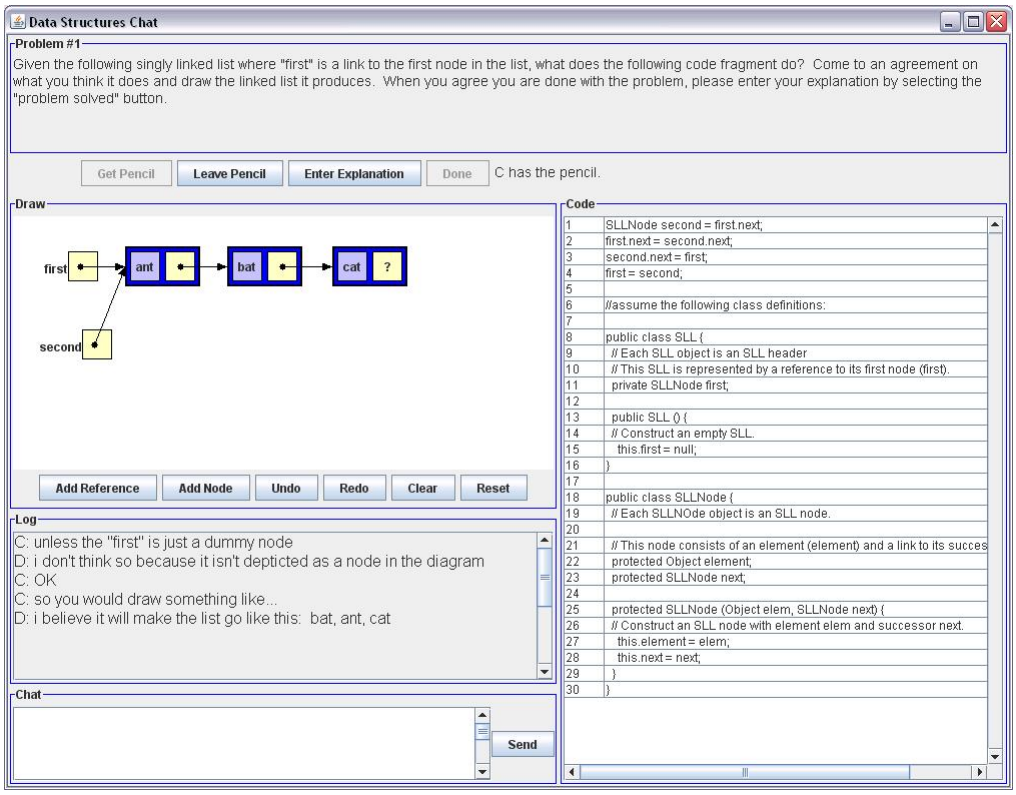

Figure 1: The KSC-PaL interface

PaL using natural language and graphical actions. The student input is processed by the interface and its related modules into an appropriate format and passed to TuTalk. Since TuTalk's interpretation module is not able to appropriately handle all student utterances, a human interpreter assists in this process. The interpreter receives a student utterance along with a list of possible matching concepts from TuTalk (see Figure 4). The interpreter then selects the most likely matching concepts from TuTalk thus assisting in natural language interpretation. If the student utterance doesn't match any of these concepts, a second list of concepts, containing student initiative utterances, are presented to the interpreter. If none of these match then all known concepts are presented to the interpreter for matching. Note that the interpreter has a limited, predetermined set of choices, corresponding to the concepts that TuTalk is aware of. In this way, his/her intervention is circumscribed.

The second addition is the incorporation of a student model that allows the KSC-PaL to track the current state of problem solving and the student's knowledge in order to guide its behavior. TuTalk's student model was replaced with one that incorporates problem solution graphs (Conati et al., 2002). Solution graphs are Bayesian networks where each node represents either an action required to solve the problem or a concept required as part of problem solving. A user's utterances and actions are then matched to these nodes. This provides KSC-PaL with information related to the student's knowledge of problem solving concepts as well as the current topic under discussion.

Thirdly, a planning module was added to TuTalk to make decisions on implementation of problem solving goals and responses to student initiative in order to manage shifts in initiative. The planning module includes an initiative tracker that codes turns with with either student initiative or agent initiative using one classifier for natural language utterances and a separate classifier for drawing and coding actions. Once a turn is classified, it is determined whether a shift in initiative has occurred by compar-

ing the current classification with the classification of the previous turn. An average level of initiative shifts is then computed by dividing the number of initiative shifts by the total number of turns. Based on the initiative level, KSC-PaL encourages initiative shifts by using prompts, hedging, requesting feedback from the student and encouraging student criticism by intentionally making errors in problem solving.

Our evaluation of KSC-PaL (Kersey et al., June 2010) found that students learned using the agent and that KSC-PaL was was successful in encouraging shifts in initiative.

3 Demonstration Outline

We will demonstrate a problem solving episode with KSC-PaL where a user will interact with the agent as a student. Specifically we will show how KSC-PaL attempts to manage the level of initiative and how KSC-PaL reacts to student initiative.

```
1.  Amy: hi there, are you ready
     to start?
2.  Mary: sure
3.  Amy: where to begin?
4.  Mary: i don't know
5.  Amy: maybe we should draw this
6.  Amy: do you want to draw it?
7.  Mary: ok
8.  Amy: so, line 1 ...
9.  Mary: how's that?
10. Amy: hmmm
11. Amy: line 2: first.next=
     second.next
(Mary diagrams the code segment)
12. Mary: there i finished it
13. Amy: looks good
14. Amy: i think we're done drawing
15. Mary: yes
16. Amy: basically, it switches the
    first and second elements
17. Mary: right
18. Amy: ill write the explanation
19. Mary: ok
```

Figure 2: Problem Solving Example

Figure 2 shows an excerpt from a typical problem solving session. Amy, the agent, attempts to maintain a balanced level of initiative either by encouraging the student (Mary) to take initiative, or by taking initiative. In lines 10 and 11, the agent uses a prompt followed by repetition of the problem solving goal to encourage the student to retain the initiative. However, in line 16 the agent takes initiative by proposing a solution to the problem.

KSC-PaL is also able to manage the situation where the student takes initiative in the dialogue in an unanticipated manner, as shown in Figure 3. In line 3, Mary asks a question that is not expected at this point in the problem solving process and the agent responds appropriately in line 4. KSC-PaL is able to respond to a set of student initiations derived from corpus data, such as the one in this excerpt. Additionally, the agent responds in a generic manner to other unanticipated initiations and directs the student back to problem solving.

```
1.  Mary: let's look at the code
2.  Amy: so, i think first and
     second both equal bat
3.  Mary: is second a reference or
     a node?
4.  Amy: second is a reference
```

Figure 3: Student Initiative Example

References

A. L. Brown and A. S. Palincsar, 1989. *Guided, cooperative learning and individual knowledge acquisition*, pages 307–226. Lawrence Erlbaum Associates, Hillsdale, NJ.

M.T.H. Chi, N. De Leeuw, M.H. Chiu, and C. LaVancher. 1994. Eliciting self-explanations improves understanding. *Cognitive Science*, 18(3):439–477.

Jennifer Chu-Carroll and Sandra Carberry. 1998. Collaborative response generation in planning dialogues. *Computational Linguistics*, 24(3):355–400.

Cristina Conati, Abigail Gertner, and Kurt VanLehn. 2002. Using Bayesian networks to manage uncertainty in student modeling. *User Modeling and User-Adapted Interaction*, 12(4):371–417.

María de los Angeles Constantino-González and Daniel D. Suthers. 2000. A coached collaborative learning environment for entity-relationship modeling. *Intelligent Tutoring Systems*, pages 324–333.

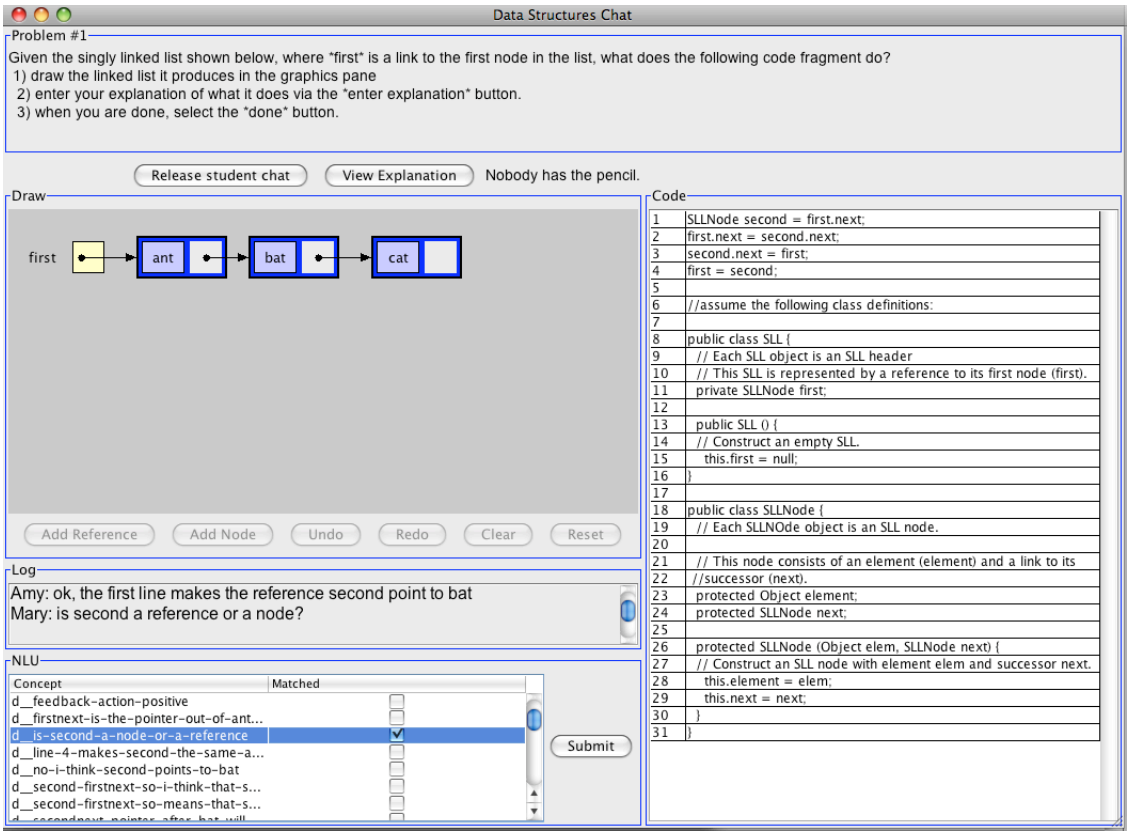

Figure 4: The interface for the human interpreter

E. Fisher. 1993. Distinctive features of pupil-pupil classroom talk and their relationship to learning: How discursive exploration might be encouraged. *Language and Education*, 7:239–257.

Robert G.M. Hausmann, Michelene T.H. Chi, and Marguerite Roy. 2004. Learning from collaborative problem solving: An analysis of three hypothesized mechanisms. In K.D Forbus, D. Gentner, and T. Regier, editors, *26th Annual Conference of the Cognitive Science Society*, pages 547–552, Mahwah, NJ.

Pamela W. Jordan and Barbara Di Eugenio. 1997. Control and initiative in collaborative problem solving dialogues. In *Working Notes of the AAAI Spring Symposium on Computational Models for Mixed Initiative*, pages 81–84, Menlo Park, CA.

Pamela W Jordan, Brian Hall, Michael A. Ringenberg, Yui Cue, and Carolyn Penstein Rosé. 2007. Tools for authoring a dialogue agent that participates in learning studies. In *Artificial Intelligence in Education, AIED 2007*, pages 43–50.

Cynthia Kersey, Barbara Di Eugenio, Pamela Jordan, and Sandra Katz. 2009. KSC-PaL: a peer learning agent that encourages students to take the initiative. In *Proceedings of the Fourth Workshop on Innovative Use of NLP for Building Educational Applications*, pages 55–63. Association for Computational Linguistics.

Cynthia Kersey, Barbara Di Eugenio, Pamela Jordan, and Sandra Katz. June 2010. KSC-PaL: A peer learning agent. In *ITS 2010, The 10th International Conference on Intelligent Tutoring Systems*, Pittsburgh, PA.

R. Ploetzner, P. Dillenbourg, M. Preier, and D. Traum. 1999. Learning by explaining to oneself and to others. *Collaborative learning: Cognitive and computational approaches*, pages 103–121.

Rod D. Roscoe and Michelene T. H. Chi. 2007. Understanding tutor learning: Knowledge-building and knowledge-telling in peer tutors' explanations and questions. *Review of Educational Research*, 77(4):534–574.

Amy Soller. 2004. Computational modeling and analysis of knowledge sharing in collaborative distance learning. *User Modeling and User-Adapted Interaction*, Volume 14(4):351–381, January.

Tan Bee Tin. 2003. Does talking with peers help learning? the role of expertise and talk in convergent group discussion tasks. *Journal of English for Academic Purposes*, 2(1):53–66.