Guidance and Collaboration Strategies in Ill-defined Domains

Toby Dragon
Department of Computer Science
University of Massachusetts -Amherst
dragon@cs.umass.edu

Beverly Park Woolf
Department of Computer Science
University of Massachusetts -Amherst
bev@cs.umass.edu

Abstract. We developed tutors to support critical thinking in four ill-defined domains. Students work with science or art history tutors and receive feedback about their arguments as they sort, filter and categorize data. A coach provides guidance using a set of expert rules and an expert knowledge base, which creates the basis for the content-specific analysis of the student's argument. This paper describes these tutors, their current functionality and our future research to improve both guidance for an individual student and collaboration tools for multiple students.

Keywords: Ill-defined domains, problem-based learning, hypothesis generation, case-based coaching, collaboration, argumentation.

INQUIRY TUTORS FOR ILL-DEFINED DOMAINS

Introduction

A frequently mentioned weakness of education is that, in an effort to get through as much material as possible, many topics are covered briefly and stripped of their contextualizing complexity (Eylon & Linn, 1988; Koschmann et al., 1994). Especially in science education, students are left with a dizzying array of new terms and concepts that have no relationship to their prior knowledge or their everyday experience. Much of the learning that does occur provides the student with only inert knowledge (Eylon & Linn, 1988). For these reasons, teaching these ill-defined domains provides both a major challenge and a major opportunity for intelligent tutoring systems. We propose a set of tools and a working environment that allows students to successfully approach complex problems while still in context. The system supports the students while they structure and relate the knowledge that they gain, and provides feedback about both the domain knowledge and the learning process.

Our system, Rashi, takes an inquiry learning approach to these ill-structured domains. Inquiry learning has been shown to be more successful than typical classroom instruction: it stimulates interest and engages students (White & Frederiksen, 1995; Shute & Glaser, 1990). When students manipulate artifacts themselves and think freely about problems, they become more actively involved and generally become more systematic and scientific in their discovery of laws (White & Frederiksen, 1995). The increased interactivity alone has been shown to increase learning (Shute & Glaser, 1990).

The Rashi tutor currently supports reasoning in four domains; geology, biology, art history and forestry; and has been tested with hundreds of higher education students. Our inquiry tutors for science and art history contain open questions and problems with unfocused and ambiguous solutions. The tutor presents authentic cases in their contextual complexity and invites students to generate questions to investigate the case, gather and analyze data, and generate hypotheses based on inferences made from evidence. We now describe the set of tools provided by Rashi, and how these tools support a student's critical thinking.

Cognitive Tools

Rashi provides both an expansive method of presenting information and a set of tools to help students access and organize information.[1] It is domain independent and multi-disciplinary in that several disciplines share the same infrastructure and generic tools, Figure 1.

Data collection tools include:

[1] These tutors are located at http://ccbit.cs.umass.edu/Rashihome/projects/

- *The Image Explorer*: students click "hotspots" to move to new images, view video clips, and to collect data.
- *The Interview Tool*: students ask questions and receive audio, video, and text answers.
- *The Concept Library*: a hyper-text tool in which students read about content knowledge and collect snippets of text for their notes.
- *The Source Editor*: students access references to information outside of the Rashi system, such as websites and textbooks the instructor finds important. Students can also cite resources they find on their own.

Once data is collected, *critical thinking tools* help students formulate hypotheses, organize evidence and construct arguments, within a central data repository. These tools include:

- *The Inquiry Notebook*: facts collected from data gathering tools are automatically entered into the notebook and students enter facts directly as well.
- *The Argument Editor*: students create hypotheses and inferences and drag data from the notebook to support or refute their claims.

The data collection tools have the ability to present complex and ill-structured topics of study. The critical thinking tools help students to organize knowledge and evaluate real-life questions within these topics without simplifying them or disconnecting them from their context. This helps students not only learn the domain knowledge within the system, but helps prepare students for further study by providing them with a system and set of concepts which allow them to operate in real-world scenarios. If a student needs to develop her understanding beyond the internal representation the tutor provides, she can use outside resources, cited either by herself or the instructor in the Source Editor, and still use the same organizational tools to facilitate learning. This is especially important in these ill-defined domains because one can very seldom assume that all necessary or useful knowledge will be included in the internal representation of the domain.

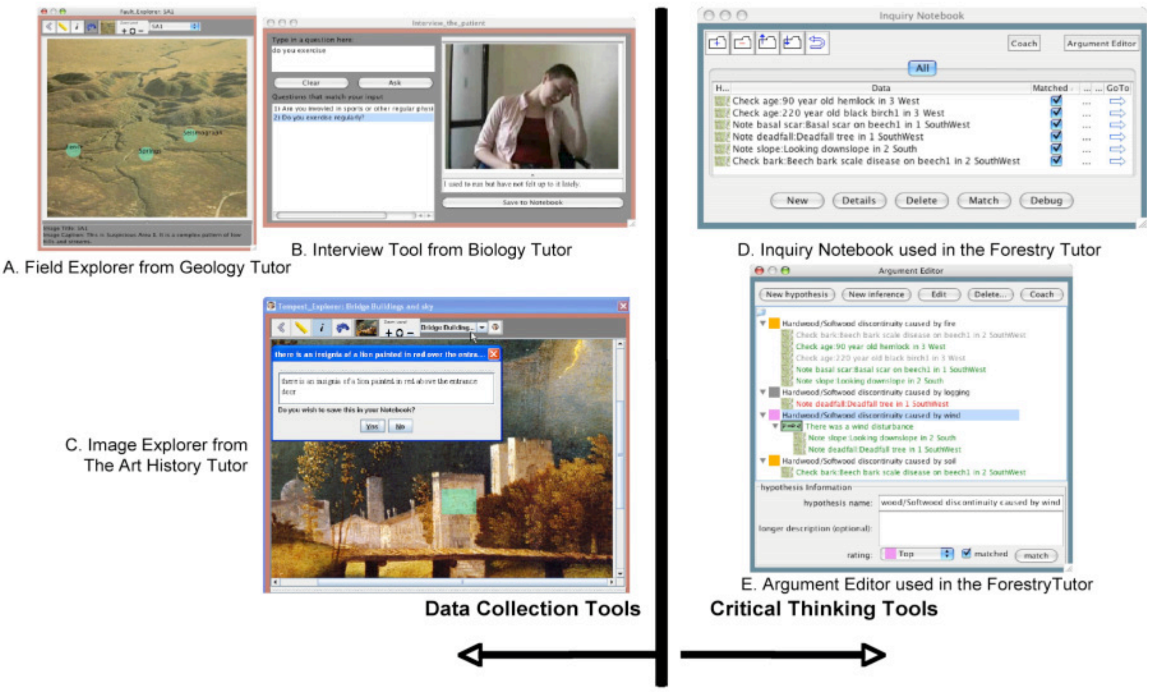

Figure 1. Data Collection and critical thinking tools are available in all the inquiry tutors.

The tools prepare students for further study in ill defined spaces by providing a structure that can be used in real-world scenarios, e.g., the same tools or conceptions can be used with resources outside of the tutoring environment. Different data collection methods (interactive text, still images, video and dynamic maps) create a broad and open-ended space for student exploration and acquaint them with methods commonly used by scientists in the field. This infrastructure has been used in three science domains, biology, geology and forestry (Woolf et al., 2002; 2003; 2005; Dragon et al., 2006).

Another reason the Rashi System is beneficial in these types of domains is that the tutor does not enforce a particular order of student activity; rather it allows students to move opportunistically from one phase to another. For example, in the Biology Tutor, students read a case description and use tools such as the Physical Examination and Laboratory to identify the patient's signs and symptoms. They might use the Interview Tool to ask the patient about her complaints and organize physiological signs, medical history or patient examinations in the Inquiry Notebook. This can be considered a learner-centered approach, because it enables students to explore and query in a variety of ways. "Learner-centered" software increases the educational opportunity for all students by accommodating each learner's individual differences (Bransford et al., 1999). Rashi allows students to discover their own approach to solving cases, and also prepares them for real-life scenarios where one has no instruction as to what the next step should be.

When providing a system that teaches complex domains, one must provide both an expansive method of presenting information as well as useful and understandable tools to help students access and organize the information and their own thoughts. We have seen the set of tools Rashi provides to accomplish this task. Now let us look to how Rashi attempts to provide guidance to the student in these ill-defined domains.

GUIDING STUDENT LEARNING

Providing an appropriate environment to support learning is essential, yet the Rashi environment alone is not enough; students still become confused, lost, and acquire misconceptions while exploring the domains. For these reasons, we have developed and tested a strong coaching component that guides student exploration. The coach uses an expert system, detailing the domain knowledge that the teacher would like to impart. As a student moves through the inquiry cycle, the tutor matches the student's reasoning with the expert knowledge (Dragon et al., 2006) and provides five feedback types:

1. *Hypothesis Feedback*: promotes consideration of multiple hypotheses and offers a list of hypotheses from the expert knowledge base.
2. *Support and Refutation Feedback*: encourages top-down argument construction by urging students to supply supporting or refuting data for their arguments.
3. *Argument Feedback:* encourages bottom-up argument construction by urging students to consider higher-level arguments, possibly offering an argument the student might have missed.
4. *Relationship Feedback*: helps students identify relationships between propositions used in their arguments.
5. *Wrong Relationship Feedback*: identifies contradictions between relationships in the student argument and the relationships in the expert knowledge base, Figure 2.

Feedback includes contextual information about the case being solved. For example, when the tutor informs a student she needs more support for a certain argument, it can also bring her to a location where this support can be found. When correcting a student's relationships, Figure 2 bottom, the tutor can directly address and correct a student's misconceptions. The coach also promotes good inquiry behavior by encouraging students to engage in sound reasoning. For example, if a student randomly collects data and does not explicitly make arguments about the case, the coach will ask her to propose a hypothesis. Once the student makes a hypothesis, the coach urges the student to support or refute it with evidence.

Teachers can adjust the coach in a number of ways: specifying the order in which competing domain knowledge is presented, and promoting a certain order to the inquiry process (Dragon et al., 2006). For example, an instructor might want students to see a crucial piece of evidence before anything else when the student is considering a certain hypothesis. The teacher assigns nodes in the expert knowledge base with an *importance* rating that guides the coach in how to select among nodes that are otherwise equivalent. This allows the instructor to include large amounts of data in the expert knowledge base and still rest assured the student will be pointed toward the most important data first. The instructor can also specify the order in which the different types of feedback are presented. This allows them some control over the ordering of steps in the inquiry process that the coach will present. For example, if the instructor wants the students to collect data before creating relationships, she can make the support and refutation feedback come before relationship feedback to promote this type of behavior. These features allow the instructor who authors the case some ability to fine-tune the behavior of the coach, guiding students through ill-structured spaces in a specified way.

Future Work On Guidance

While the coach does currently engage in these activities, there is still much room for improvement and many open questions as to how best to approach this task. Some of these areas that we will discuss briefly are: When should the coach provide help, should it intervene in student work? How best can the coach traverse the expert knowledge base in order to support the student? What information can be used from the student argument, and how can one balance contradicting information?

The subject of when to offer help is widely debated in the tutoring field. On one hand, students who are engaged should not be interrupted unnecessarily because it can be disruptive to their thought process and seen as an annoyance. On the other hand, research has indicated that students are not incredibly successful at identifying when they need help (Aleven & Koedinger, 2000). This indicates that a system should monitor student behavior in order to intervene at key moments to support learning. Thus far, the Rashi system has only offered help on demand, but we are considering methods of monitoring the student to immediately notice and react to both statements of misconception (such as creating an incorrect relationship) as well as bad inquiry behavior (such as collecting large quantities of random data without a noticeable reason for doing so).

The order in which to traverse the knowledge base when coaching students is a difficult question. Good inquiry requires the consideration of many hypotheses, yet the coach should not distract the student by changing subject while they are engaged. So the real question at hand is how to balance a breadth-first vs. a depth-first approach to traversing the expert knowledge base. Along the same lines, one must decide whether the coach should help the student work down from hypotheses to data, or help them work from data collected up towards hypotheses (top-down vs. bottom-up approach). Research has indicated that a mixture of top-down and bottom-up is most conducive to learning (Krajcik et al., 1998). One possibility we are considering for these questions

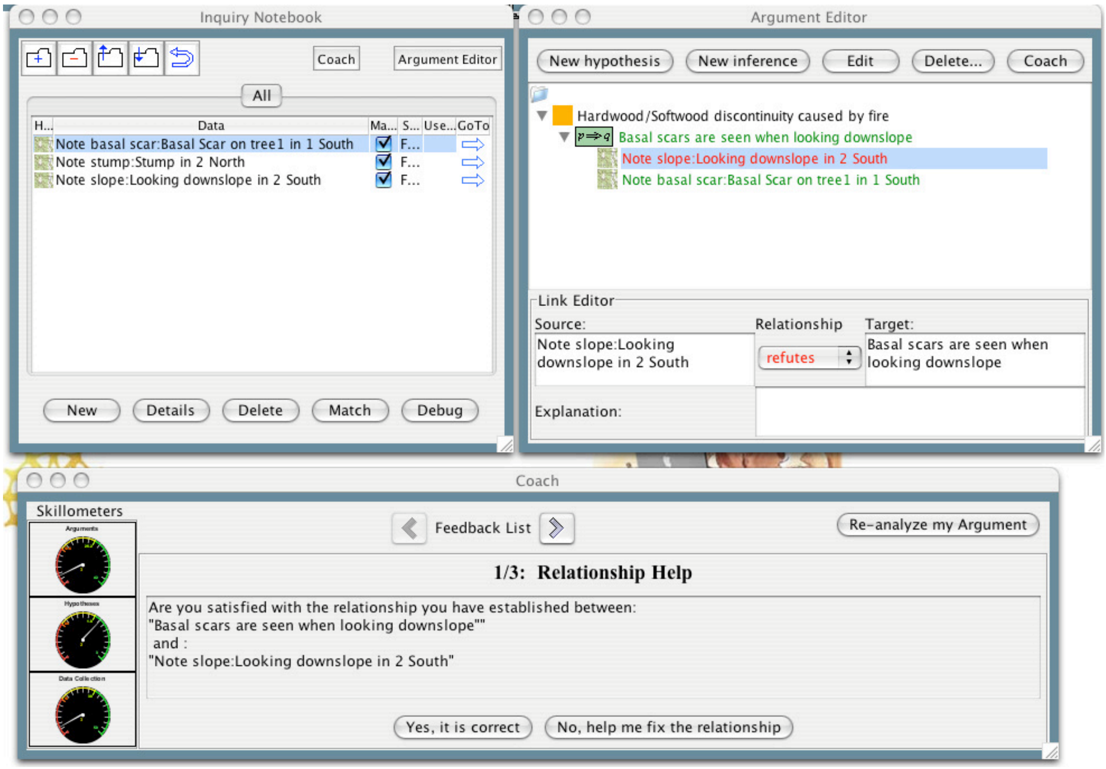

Figure 2. The student receives help from the Forestry Tutor.
Student observations, data and hypotheses are automatically recorded and analyzed in the Inquiry Notebook (top left). This student posits hypotheses in the Argument Editor (top right) and uses a drag and drop feature to move data from the Inquiry Notebook to link supporting or refuting evidence to hypotheses. Here the coach has identified a relationship which the student has set up incorrectly, and is encouraging the student to reconsider.

involves weighting which step is more crucial (e.g. if a hypothesis is almost completely supported, then it should be less important to investigate than a hypothesis that has no support). Another involves tracking the recency with which a student has operated upon relevant material (e.g. if they have recently collected some data, help them work from that data towards a hypothesis).

Finally, we consider the different methods of analyzing student work in these ill-defined inquiry domains, and how to deal with contradictions posed by different analyses. Earlier versions of Rashi used a syntactic coach that analyzed not the content of a student argument, but the structure. For example, if the coach found that a student had more support for an argument that was ruled out than for an argument that was considered the student's top hypothesis, it would ask the student to consider the argument again. This idea has obvious flaws, because one piece of supporting evidence can outweigh any number of other pieces of evidence depending on the meaning of that evidence. For this reason, we moved to an expert knowledge base approach. Yet the expert knowledge base approach requires that every domain have one of these expert knowledge bases, which is time-consuming and difficult to define. So there are obvious advantages and disadvantages to both approaches. We are developing a hybrid approach which attempts to utilize the best of both methods, but this does give the coach competing ideas about how to give advice, and it is not always clear which advice is most useful. The best approach so far seems to be using the expert knowledge base advice when available, and the syntactic advice when the other is missing.

So we can see that Rashi does offer some compelling methods of guidance through ill-defined domains, providing students with both domain knowledge and direction in the inquiry process itself. As we discover better solutions to the problems discussed, the coaching will become more useful and productive. Until then, there is another, equally interesting approach to helping students in ill-defined domains: collaboration.

PROPOSED COLLABORATION TOOLS

There are many difficulties with instructing a student in ill-defined domains. We presented some of these issues and proposed solutions to them, but many other issues remain. So far in our discussion, we have focused on creating an environment that supports students and intelligent help to guide students through their work. This guided help approach makes a major assumption, that the system can understand student work. Since we are still far from having a solid natural language recognition system, and we allow students to enter knowledge from outside the tutor, as well as allowing them to phrase their arguments and hypotheses in their own terms, there will always be information provided by the student that the system does not understand. This is one of many clear limitations of an artificial intelligence approach to teaching in these domains.

On the other hand, perhaps new pedagogy can help us understand student work. For example, collaborative learning is one of the most valuable educational approaches in terms of improved learning outcome (Johnson & Johnson, 2005). Often groups will outperform even the best individuals in the group, producing knowledge that none of its members would have produced by themselves and leading to the generation of new ideas (Ellis et al., 1994). Collaboration encourages students to question processes, ask for advice and monitor each other's reasoning (Slavin, 1990). It is effective especially for students who tend to be left behind in science classrooms, including women and under-represented populations (Ellis et al., 1994). Nearly 600 experimental studies and over 100 corollary studies clearly indicate that cooperation results in higher achievement and greater productivity, more caring, supportive and committed relationships, and greater psychological health, social competence, and self-esteem (Johnson & Johnson, 1989). Other research points to collaboration resulting in a boost of efficiency and accuracy, and problem solving ability (Okada & Simon, 1997).

We intend to develop collaboration tools to support synchronous, symmetric cooperation through the web. A chat system will support students to discuss cases and propose actions as if they were seated in front of the same computer. The tutor will keep track of student actions and discussions and will analyze computer logs to investigate, for example, how collaboration influences learning styles.

We intend to investigate a number of research questions. For example, do group collaborative activities improve individual domain and/or inquiry learning? Does efficient (in the sense of group success) and effective (in the sense of positive group dynamics) collaboration reflect improved individual learning as measured in post-tests? Does collaborative learning result in more effective learning or more positive confidence among women and minorities as compared with individual learning?

Not all collaborative groups make progress in similar settings: it is therefore important to understand groups and to help decide which groups are productive and should remain together, and which groups are not making progress (Tedesco & Rosatelli, 2004). We will develop both a performance model that documents how cognitive events influence completion of a task and a model of the knowledge sharing contributions of the students in the group. Research involves comparing the performance of individuals in different contexts: when students work alone and when working in groups. For this reason, Collaborative Rashi, (C-Rashi), will track all actions performed by students and will align chat exchanges with the logs of the actions taken on the Rashi system, so we can reconstruct the interaction sequence (Soller et al., 2004).

The interface shown in Figure 3 divides the student window into four areas. Area 1 shows the original Human Biology Tutor (which in turn can be divided into frames, as decided by the Rashi server). This is what users see when connected to the original, individual (not collaborative) tutor. Areas 2 and 4 are used for communication among participants; area 2 is student input to the chat while area 4 contains a log of the chat. Area 3 is used to manage the control of the interaction with C-Rashi, and the engine behind it can implement different policies. Verbal exchanges are classified into different phases, Propose, Discuss, Review, incorporating fifteen sentence openers that help analyze effective peer dialogue (Soller & Lesgold, 2003). In addition, three 'quick' buttons ('OK', 'Yes', and 'No') are provided (Area 2) to indicate agreement. Each sentence opener represents a different cognitive process related to the problem solving phase and are simple enough for students to find and select. Students start their conversation using the openers (which also saves them some typing) and then complete the phrase in their own text. The tutor will manage the conversation through the use of topic threads (based on context) that attempt to structure the discussion to reflect the structure of the team's decision processes. In this way, C-Rashi will support learners through the various phases of problem solving, facilitating an extended, in-depth, on-topic discussion and providing a coherent view of the argument. The synchronous, symmetric environment will use a dedicated applet to manage 'mouse control' by introducing the concept of token. A student's clicks on Area 1 in Figure 4 are effective only when she has the token. To prevent the user from clicking on links when she does not have the mouse, a JavaScript function will check the ownership of the token.

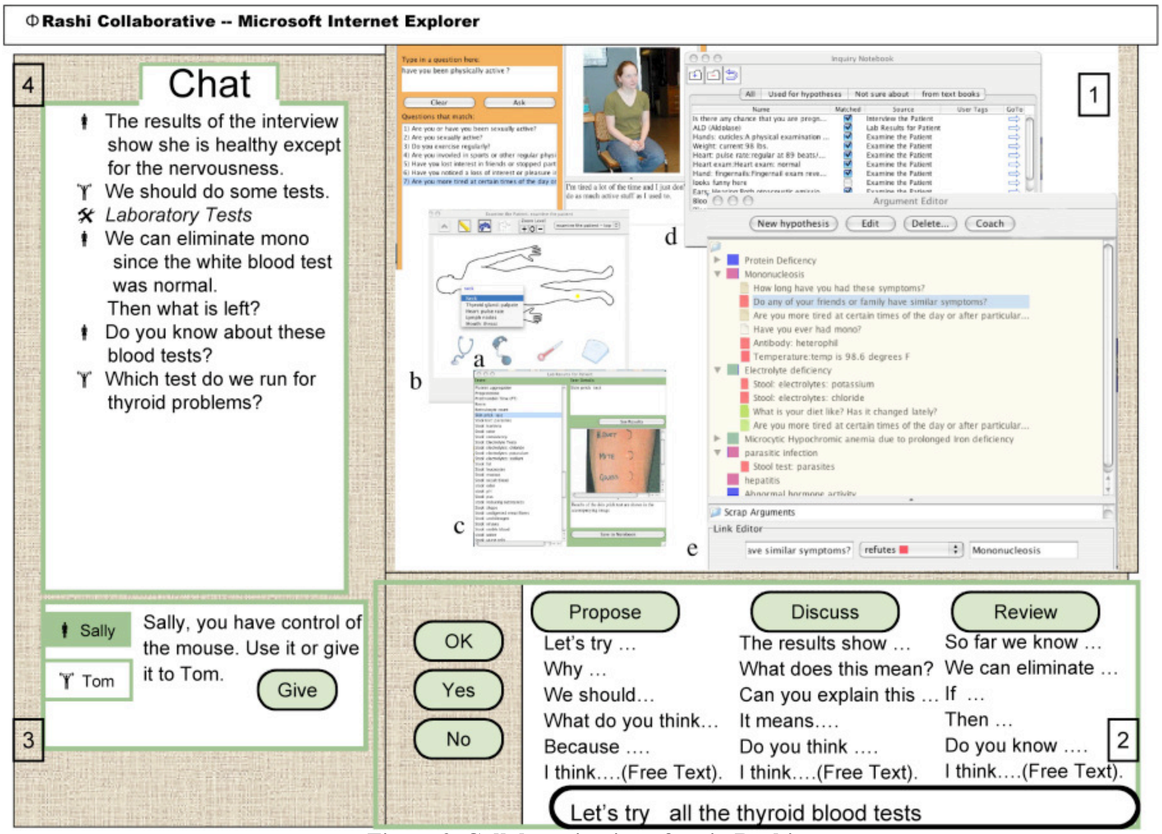

Figure 3. Collaboration interface in Rashi

The Rashi interface is embedded within the collaborative environment allowing participating students to work on an inquiry case (1), communicate with teammates (2), use the shared mouse (3) and review the conversation (4).

Collaboration Theory and New Software Functionality

We base our collaboration environment in part on Karl Smith's widespread and successful 'cooperative learning' model (Smith et al., 2005). This model emphasizes structured or formal collaborative work, and has many elements that are ideally suited to computer-supported learning: positive interdependence (of task, identity, resources, role, reward, and goals), individual and group accountability (challenging and motivating work for every student in a group) and authentic interaction (brainstorming, planning, social and team building skills, and solidarity). Collaborative Rashi will be modified in the following ways to reflect this theory:

- *Forming the team:* We will experiment with various ways to form student teams. Placing students in groups and telling them to work together does not always result in cooperation. We will have teams predefined by the teacher and groups dynamically formed 'on the fly'. We will also experiment with teams selected based on using student models. After collaboration, student polling will ask each participant whether each team member was helpful or not, and this information will be used to improve the student model approach.
- *Monitoring tools and instructor control panel:* The teacher will monitor general trends in student work at individual, class, and community levels, such as the number of hypotheses logged, or whether certain information was recorded. Monitoring tools facilitate accountability within groups and within classes making visible the general level of work a student has done without revealing the solution in progress.
- *Roles (duties) support*: Assigning roles, including task manager, skeptic, accuracy checker, social facilitator and record keeper, has been an essential part of many successful case-based learning methods. Students should have the experience of each different role in a semester. The system will send reminders (for students to discharge their duties, triggered by intelligent rules) and will indicate which tools should be used to accomplish that role (e.g., the skeptic could use the critique feature mentioned above). All students are expected to work collaboratively to solve the problem; roles are extra duties or 'hats' that they put on periodically.

PRIOR EVALUATION OF RASHI

The proposed changes to the Rashi system listed in this report constitute a large investment of time and energy. Yet the project has met with much success, and there is good reason to believe these changes will be fruitful.

The Rashi cases have all been evaluated at Hampshire College and the Universities of Massachusetts and Rhode Island with undergraduates as well as middle school science teachers. The Biology Tutor was evaluated several times in large (500 students) university lecture-based classroom. However, as there was only time to use a few short cases, we consider this evaluation to be a pilot study to test the evaluation instruments. Nevertheless, the results were encouraging: students quickly learned the software and posed open-ended and authentic questions, planned queries and engaged in on-line research. We have also noted significant correlations between a student's inquiry skill level and some of the Rashi use metrics. In particular, there were significant positive correlations between a student's measured inquiry skill level and the number of *hypotheses* posed by that student, the number of *arguments*, the number of *items in the notebook*, the number of *explanations* entered by students, the use of notebook organizing tools and the overall use of Rashi tools. As this is what one would expect, this adds some credence to the validity of the pre-post test instrument. We interpret these results as supporting the usability of the software and its perceived usefulness. Interviews, surveys, essay questions, group discussions, and pre-post essay activities have shown that participants were enthusiastic and impressed with the potential of Rashi as an educational tool. Interactivity was seen as a very positive attribute, with the *Patient Examination* feature in Biology cited as one of the better components. Students' perception of learning the inquiry process was favorable.

Students found the Rashi cognitive tools helpful to organize data and create good arguments. Students were highly positive toward the software, especially toward the tools that allowed them to gather data. They were challenged and very engaged (as anecdotal evidence, one said that if the software were available, she would ask her parents to buy it for her as a Christmas present). At first, students treated the system like a "video game." They were drawn to the *Image Explorer* and other highly visual elements. In conversation with each other, students generated interesting hypotheses and entered data into the Inquiry Notebook, which helped them conduct their analyses and helped faculty analyze their approaches. We observed students going to their texts often and posing multiple hypotheses. Students discussed their reactions with an evaluator and completed surveys. Their overall reactions were very positive, despite certain individual issues. Students did not have difficulty navigating through the tools after the initial explanation and thought it would be fairly easy for a naive user to catch on. Most aspects of the system were rated positively by roughly half of the users.

RELATED WORK AND CONCLUSIONS

Related work in inquiry and argumentation tutors has led to case-based learning environments and tools for gathering, organizing, visualizing, and analyzing information during inquiry (Aleven & Ashley, 1997; Suthers et al., 1997). Some systems support authentic inquiry and knowledge sharing, and several track and analyze student data selections, providing students with space to explore subject matter from microeconomics to medical diagnosis. However, most of these systems do not evaluate a student's hypotheses and have more of a data collecting than an experimentation feel; others are narrowly applicable and do not allow for interactive tutoring. Most are restricted to a single domain since limiting the domain allows the designer to facilitate specific types of interactions (Aleven & Ashley, 1997). Other systems are built with the primary goal of teaching the inquiry process and not the domain.

Rashi maintains a student model that compares student arguments and collected data with the knowledge in expert system, and a domain model that provides a structure with which to model the student. The tutor i) provides a free exploratory space while tracking student arguments, ii) provides intelligent contextual advice and critiques a student's evidence; iii) and remains domain-independent and flexible enough to quickly encode new cases and new domains. In this way, the system provides means for teaching a broad range of ill-defined domains, including science and liberal arts. The current system provides a full set of tools that allow students to engage in inquiry learning with complex, real-world cases. These tools provide structure and organization while still allowing for freedom of exploration and a learner-centered approach. The coaching system guides the student in both domain knowledge and the inquiry process to help those who are lost or confused.

We are currently developing many improvements to the system. Coach improvements include addressing the issues of when to offer help, how to best traverse the knowledge base, and how to handle contradictory goals when analyzing the student argument in different ways. We plan to make the system a collaborative environment that supports both group work and peer and instructor review.

ACKNOWLEDGEMENTS

Research on Rashi was funded in part by 1) the U.S. Department of Education, "Expanding a General Model of Inquiry Learning", Fund for the Improvement of Post Secondary Education, Comprehensive Program, #P116B010483, B. Woolf, P.I., by 2) the National Science Foundation under grant DUE-0127183, "Inquiry Tools for Case-based Courses in Human Biology," M. Bruno, PI and Woolf, Co-PI, by 3) National Science Foundation, CCLI #0340864, "On-line Inquiry Learning in Geology," D. Murray, P.I., B. Woolf co-PI, 4) and by 4) National Science Foundation DUE-0341521 "Reading the Forest Floor: Online Inquiry Learning in Forestry." B. Woolf, P.I., and L. Winship, Co-PI, and 5) National Science Foundation DUE-0341197, "Engaging Undergraduates in On-line Inquiry Learning: A Case-based Cyber Library in Human Biology," M. Bruno PI.

Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the funding agencies.

REFERENCES

Albanese, M. A., & Mitchell, S. (2003). Problem-based learning: a review of literature on its outcomes and implementation issues. *Acad Med, 68*(1), 52-81.

Aleven, V., & Ashley, K. D. (1997). Teaching case-based argumentation through a model and examples: Empirical evaluation of an intelligent learning environment. In B. d. Boulay & R. Mizoguchi (Eds.), *Artificial Intelligence in Education, Proceedings of AI-ED 97 World Conference* (pp. 87-94). Amsterdam: IOS Press.

Aleven, V., & Koedinger, K. R. (2000). *Limitations of student control: Do students know when they need help?* . Paper presented at the 5th International Conference on Intelligent Tutoring Systems.

Alloway, G., Bos, N., Hamel, K., Hammerman, T., Klann, E., Krajcik, J., et al. (1996). Creating an inquiry-learning environment using the World Wide Web, *International Conference of Learning Sciences*.

Arroyo, I., Beal, C. R., Murray, T., Walles, R., & Woolf, B. P. (2004). Web-Based Intelligent Multimedia Tutoring for High Stakes Achievement Tests. In *Intelligent Tutoring Systems, 7th International Conference* (pp. 468-477). Macei0, Alagoas, Brazil: Springer.

Aspy, D. N., Aspy, C. B., & Quimby, P. M. (1993). What doctors can teach teachers about problem-based learning. *Educational Leadership, 50*(7), 22-24.

Bransford, J. D., Brown, A. L., & Cocking, R. R. (Eds.). (1999). *How People Learn: Brain, Mind, Experience, and School*. Washington, DC: National Academy Press.

deVries, E. (2002). Computer-mediated epistemic dialogue: Explanation and argumentatin as vehicles for understanding scientific notions. *Journal of the Learning Sciences, 11*(1), 63-103.

Dragon, T., Woolf, B. P., Marshall, D., & Murray, T. (2006). Coaching within a domain independent inquiry environment, *8th International Conference on Intelligent Tutoring Systems*. Jhongli, Taiwan.

Ellis, S., Klahr, D., & Sieglar, R., (1994, April). The birth, life, and sometimes death of good ideas in collaborative problem solving. Paper presented at the annual meeting of the American Educational Research Association, New Orleans.

Eylon, B.-S., & Linn, M. C. (1988). Learning and instruction: An examination of four research perspectives in science education. *Review of Educational Research, 58*(3), 251-301.

Hakkarainen, K. (2003). Progressive Inquiry in a Computer-Supported Biology Class. *Journal of Research in Science Teaching, 40*(10), 1072-1088.

Johnson, D., and Johnson R., (2005) The Cooperative Learning Center, The University of Minnesota, http://www.co-operation.org/index.html. April 2005.

Koschmann, T. D., Myers, A. C., Feltovich, P. J., & Barrows, H. S. (1994). Using technology to assist in realizing effective learning and instruction: A principled approach to the use of computers in collaborative learning. *Journal of the Learning Sciences, 3*(3), 225-262.

Krajcik, J., Blumenfeld, P. C., Marx, R. W., Bass, K. M., Fredricks, J., & Soloway, E. (1998). Inquiry in project-based science classrooms: initial attempts by middle school students. *The Journal of the Learning Sciences, 7*(3&4), 313-350.

Mennin, S. P., Friedmann, M., Skipper, B., Kalishman, S., & Snyder, J. (1993). Performances on the NBME I, II, and III by medical students in the problem-based learning and conventional tracks at the University of New Mexico. *Academic Medicine, 68*(8), 616-624.

Okada, T., & Simon, H., (1997). Collaborative discovery in a scientific domain. Cognitive Science, 21 (2), 109-146

Pea, R. (1994). Seeing what we build together: Distributed multimedia learning environments for transformative communications. *Journal of the Learning Sciences, 3*(3), 285-299.

Scardamalia, M., & Bereiter, C. (1994). Computer support for knowledge-building communities. *Journal of the Learning Sciences, 3*(3), 265-283.

Shute, V. J., & Glaser, R. (1990). A Large-scale Evaluation of an Intelligent Discovery World: Smithtown. *Interactive Learning Environments, 1*, 51-57.

Slavin, R. E. (1990). Research on cooperative learning: Consensus and controversy. *Educational Leadership 46*, 12: 52-54.

Smith, K.A., Sheppard, S.D., Johnson, D.W., and Johnson, R.T. (2005). "Pedagogies of Engagement: Classroom-based Practices," Journal of Engineering Education, Vol. 94, No. 1, 2005, pp. 87–101.

Soller, A. (2004). Computational modeling and analysis of knowledge sharing in collaborative distance learning. *User Modeling and User-Adapted Interaction: The Journal of Personalization Research*, 14 (4), 351-381.

Soller, A. and Lesgold, A. (2000). 'Knowledge acquisition for adaptive collaborative learning environments' American Association for Artificial Intelligence Fall Symposium, Cape Cod, MA, USA

Suebnukarn, S., & Haddawy, P. (2004). A collaborative intelligent tutoring system for medical problem-based learning, *International Conference on Intelligent User Interfaces*.

Suthers, D. D., Toth, E., & Weiner, A. (1997). An Integrated Approach to Implementing Collaborative Inquiry in the Classroom, *Computer Supported Collaborative Learning (CSCL'97)* (pp. 272-279). Toronto.

Tedesco, P.A., & Rosatelli, M.C. (2004). Helping Groups Become Teams: Techniques for Acquiring and Maintaining Group Models. J. Mostow, & Tedesco, P. (Eds.). *Designing Computational Models of Collaborative Learning Interaction.* ITS 2004 Workshop, Maceio-Alagoas, Brazil.

Vernon, D. T., & Blake, R. L. (1993). Does problem-based learning work?: A meta-analysis of evaluative research. *Academic Medicine, 68*(7), 550-563.

White, B., & Frederiksen, J. R. (1995). *Developing Metacognitive Knowledge and Processes: The key to making scientific inquiry and modeling accessible to all students* (No. CM-95-04). Berkeley, CA: School of Education, University of California, Berkeley.

Woolf, B. P., Marshall, D., Mattingly, M., Lewis, J., Wright, S., Jellison, M & Murray, T. (2003). Tracking student propositions in an inquiry system. In U. Hoppe, F. Berdeho & J. Kay, (Eds.) Artificial Intelligence in Education, Proceedings of AIED 2003, IOS Press, pp. 21-28.

Woolf, B. P., Reid, J., Stillings, N., Bruno, M., Murray, D., Reese, P., Peterfreund, A. & Rath, K. (2002) A General Platform for Inquiry Learning, Proceedings of the 6th Int'l Conference on Intelligent Tutoring Systems, Lecture Notes in Computer Science 2363, 681-697, France.

Woolf, B., Murray, T., Marshall, D., Dragon, T., Kohler, K., Mattingly, M., Bruno, M., Murray, D. & Sammons, J. (2005). Critical Thinking Environments for Science Education, International Conference on AI and Education, 2005, Amsterdam.