Knowledge-based Environments for Teaching and Learning

Beverly Park Woolf, Elliot Soloway, William J. Clancey, Kurt Van Lehn, & Dan Suthers

(1991) AI Magazine, 11(5)74-77.

Second Generation Systems

The Spring Symposium on Knowledge-based Environments for Teaching and Learning focused on the use of technology to facilitate learning, training, teaching, counseling, coaxing and coaching. Sixty participants from academia and industry assessed progress made to date and speculated on new tools for building second generation systems.

Selection of topics and participants was motivated by a desire for ideological breadth and depth. Panel leaders included William J. Clancey and Alan Lesgold (researchers of real-world systems); Kurt VanLehn (champion of cognitive models); Beverly Park Woolf (defender of discourse systems); Elliot Soloway (advocate for alternative environments); and Sarah Douglas (spokesperson for supportive systems).

Human-Computer Interaction

Researchers have moved away from building omniscient tutors capable of detecting all possible errors and misconceptions. Instead, research is now focused on building empathetic partners that choose from among several forms of interaction based on the content of the communication and the needs of the student [Woolf, 1988]. Possible communication styles include didactic explanation, guided discovery learning, coaching or coaxing, and critiquing. Although no one style is preferred, different tutorial applications will be better addressed with a given primary style.

For example, as explained by Dan Suthers and James Lester, didactic explanation is good for communicating a body of declarative knowledge shared by some community (e.g. biologists). In such applications, the student needs to learn the community's terminology, and thus didactic explanation may be more efficient than requiring a student to rediscover the principles of the field on his or her own. On the other hand, the more active nature of discovery learning helps the student "own" the acquired knowledge to a greater extent than can didactic explanation.

The style of interaction varies *within* a tutorial domain as well as across types of domains. For example, Lewis et al. [1990] showed how a (human) tutor changed strategies from script-like to opportunistic when students suggested an activity or showed the need for remediation of a deficiency.

Communication Research Issues

Pressing research issues in human-computer communication were identified both in *artificial intelligence* and in *education*. In *artificial intelligence*, research issues include the representation and control of knowledge. From this perspective, knowledge of didactic explanation might be represented and organized in a system, along with the basic knowledge of a domain. Indexing mechanisms for accessing different perspectives on the topic should be designed using abstractions appropriate for the content selection task.

Choosing and organizing domain knowledge provides the next set of research issues. Control should account for the tutor's ability to dynamically switch strategies according to multiple constraints in a manner sensitive to features that human tutors use in tutorial interactions. The tutor should consider available student modeling/diagnosis when making tutorial decisions based on multiple goals. Further work is required to characterize "relevance" for selecting knowledge for didactic explanation, especially when multiple perspectives on the topic are available. Even when the primary emphasis is on stimulating the student's own creativity and intelligence, the program's design must still be based on solid theory of relevance to select its actions and response. To do so, memory and pragmatic knowledge should be brought to bear on language processing.

Another research issue concerns the characterization of coherence in machine response. Is coherence a property of the "knowledge pool" to be used in generating the next response or a property of the dialogue or both? In choosing content from a multiple granularity knowledge base, how do we ensure that the chosen pool of knowledge is coherent given the dialogue context?

Educational research issues focus on adequately modeling the student and the pedagogical context (see next section), and then identifying how a system might stimulate and facilitate the student's own abilities and creativity.

A separate issue concerns how relevant knowledge should be presented once it has been selected. For example, the tutor might state generalizations, use case examples, or provide analogies. Presentations, whether explanations or examples, must be presented in such a way that the student will be prepared to understand new material and integrate it into an existing conceptual framework, or into one which has been built up in the preceding dialogue. We need to better understand how to choose and coordinate multi-media/modality presentations at the interface media level, e.g. the use of text, diagrams, charts, pictures, animation, and sound.

In summary, despite much work attempting to do so, we still have not figured out how to make dialogue sensitive to dialogue context and to what is known or knowable about the student's "state."

Cognitive Modeling

The cognitive modeling group provided strong advocacy for the use of cognitive modeling in building these systems. They argued for increased use of modeling at three stages of design of knowledge-based systems, primarily (1) development of pedagogical and subject-matter theories, (2) design of instruction, and (3) delivery of instruction. Of these phases, the design of instruction is the one that seems to have achieved the most direct benefit from cognitive modeling, including substantial benefits from modeling subject matter experts. For instance, Anderson et al. [1990] attribute much of the success of their tutors to the cognitive task analysis of experts in Lisp, geometry and algebra.

Work on modeling good teachers and tutors has only just begun (with the exception of a few early classics, such as the work of Stevens and Collins on Socratic tutoring [1977]). VanLehn expects this line of investigation to pay off at least as well, if not better, than the modeling of experts and learners.

Of the three phases of pedagogical work, the actual delivery of the instruction is the area where cognitive modeling has found the least fruitful application. Mostly, this is due to a historical accident. In most systems to date, teacher models have been weaker than expert models and student models. Although a good teacher model might compensate for an impoverished expert or student model, experience has shown that strong expert and student models require a decent teacher model for the system to be effective.

VanLehn underscores the fact that modeling is just good engineering practice, regardless of whether one is building a hydroelectric dam or a science course. With tongue in cheek, he suggests that if students could sue malfeasant instructional developers, cognitive modeling would be much more common since it is so obviously effective.

William J. Clancey, however, was more reserved about the utility of cognitive modeling. While acknowledging that building such models is possible, he questions the relation they have to mechanisms of human learning. For instance, does the model show the student how to interpret and generate domain concepts, or does it simply justify the machine's presentations? Clancey would like to see alternative cognitive models available within a system rather than a single "correct" model used to justify instruction.

Understanding Plans and Goals

In the move away from building all-knowing and all-powerful tutors, researchers have focused on developing environments that implicitly elicit information about student goals and plans. Human dialogue succeeds despite ambiguity and digressions because both participants model the discourse, the subject matter, and the other speaker; and both participants actively work towards success of the discourse.

This suggests that continuing efforts be made to enhance the machine's ability to do its part. Techniques such as plan recognition and learning still play only a small role in current teaching systems. Interfaces were described that inquire about beliefs and high-level thoughts while supporting meta-cognitive activities. Students might choose from a menu of high-level plans, such as a menu item in an Algebra tutor that says "collect all variables to one side of the equation." Such interfaces require more careful analysis and structuring of the task domain and of cognitive structures; they also require mechanisms to support cooperative dialogue and to "understand" student perspectives.

Real-World Applications

William J. Clancey and Alan Lesgold led several discussions on the impact of knowledge-based systems in industry and the military. The clear emergence of new architectures and positive training results have produced the feeling that progress is being made. Indeed, several systems were described which achieve the two-sigma effect [Bloom, 1984], which is the same improvement in learning that results from one-on-one human tutoring over classroom tutoring. Several success stories were described in which students using tutors learned knowledge and skills in one-third to one-half the time it took for a control group to learn the same material [Shute, 1990].

In one special case, students working with an Air Force electronics troubleshooting tutor for only 20 hours gained a proficiency equivalent to that of trainees with 40 months (almost 4 years) on-the-job training [Lesgold, Lajoie, Bunzo & Eggan, 1990]. In another example, students using a Lisp tutor at Carnegie-Mellon University [Anderson, 1990] completed programming exercises in 30% less time than those receiving traditional classroom instruction and scored 43% higher on the final exam. In a third study, students using a microworld environment learned general scientific inquiry skills and principles of basic economics in one-half the time required by students in a classroom setting [Shute, Glaser & Raghavan, 1989].

Given these results, the group asked why more tutors were not being used and why existing systems were not more effective. One reason why industry and the military have not widely adopted these systems relates to the lack of tutoring-specific artificial intelligence development tools, such as shells and frameworks, similar to the shells used to build rapidly expert systems. Tools would facilitate large-scale development; and a simple tool, such as a simulation tied to an expert system or to a lock-step tutor, might be a practical way for a designer to get started on a path of incremental design through feedback from the user. Some researchers suggested that a teacher should interact with a variety of tools, much as a conductor orchestrates a suite of instruments.

Other reasons for the slow adoption of new systems might include the difficulty in reducing cognitive task analysis to engineering practice and in developing new knowledge representations, e.g., qualitative simulation, which are better suited to representing human cognition than those offered by first-generation expert system tools. An additional barrier is the lengthy development cycle required before systems can move from research lab to salable products.

"Hot" Research Issues

Several areas emerged as 'hot' or new research areas. These were discussed throughout the symposium.

Situated Learning

Situated learning (and teaching/ acting/planning) arose frequently as a topic. It was espoused primarily by William J. Clancey, Jeremy Roschelle, and Etienne Wenger all from the Institute for Research on Learning, Palo Alto. Since situations or contexts in which a skill is learned can not be exhaustively or completely described, training systems inevitably predetermine what is relevant. Similarly, conventional Artificial Intelligence models of expertise leave out how experts know what is relevant and how they change their minds. This viewpoint suggests that Artificial Intelligence needs to place increasing emphasis on knowledge representation as an activity within a perceptual space and organized by social interactions. Existing systems omit the social context in which a domain representation is created, justified, and changed. At present, knowledge-based cognitive modeling cannot characterize the work a person does to understand the artifacts with which they interact. One reason why intelligent training systems are not more efficient is because the environment surrounding the industrial task can't be made fully explicit. Alternatively, on-the-job training is cost efficient, in part because there is no need to simulate the training situation.

Computer as mediator

Jeremy Roschelle demonstrated a system that could facilitate discussion among several students and could support student explanations and demonstrations. In such a case, the computer becomes a mediator, a malleable object capable of being pressed into service for both teaching and learning.

Andrea diSessa showed that the goals to be taught by a system are negotiable; his Boxer system is a platform which enables students to discover their own interests and which facilitates their own discoveries. For example, he showed an example in which young students invented the rules of graph construction.

Empowering curriculum designers

Jim Spohrer described a system developed at Apple Computer, Inc. which assists curriculum designers to incorporate multi-media. Oliver Selfridge challenged the group to question the nature of the learning task implicit in their emerging machines.

Qualitative reasoning

Ken Forbus demonstrated that a system could qualitatively model a complex domain, e.g, a steam boiler or a propulsion plant, and that such a representation could be used for teaching. His work on qualitative modeling is now 10 years old and its formalization is nearly ready to provide the reasoning behind qualitative modeling within a teaching environment.

Conclusions

Participants at this symposium represented diverse backgrounds and methodologies; little commonality might have been expected. Yet, some consensus was achieved and new scientific ground broken. For example, agreement was reached on the need for a variety of discourse approaches and improved cognitive models, although no particular solution to achieve widespread use of either was forthcoming.

Several areas were identified as needing further research. Basic research is needed in planning and plan recognition, building natural-language interfaces, and testing architectures, such as blackboards, for teaching systems.

From the viewpoint of communication, the symposium was a real success; discussion was lively and at times controversial. Research appears to be strong in depth, broad in perspective, and motivated by the promise of building more powerful teaching environments with greater knowledge, increased inference capability, and more complex reasoning ability. The field seems to be alive and well.

Acknowledgements

Beverly Park Woolf and Dan Suthers have received partial support from the National Science Foundation under grant number MDR 8751362, and from External Research, Apple Computer, Inc., Cupertino, CA. Kurt VanLehn's research is supported by the Office of Naval Research under contracts N00014-88-K-0086 and N00014-86-K-0678.

References

Anderson, J.R. (1990) Analysis of Student Performance with the LISP tutor. In N. Frederickson et al. (Eds.), *Diagnostic Monitoring of Skill and Knowledge Acquisition,* Lawrence Erlbaum Associates, Hillsdale, NJ.

Anderson, J.R., Boyle, C.F., Corbett, A.T. & Lewis, M.W. (1990) Cognitive modeling and intelligent tutoring. *Artificial Intelligence,* Vol. 42, 7-50.

Bloom, B.S. (1984) The 2-Sigma Problem: The Search for Methods of Group Instruction as Effective as One-to-One Tutoring. *Educational Researcher,* Vol. 13., No. 6, pp. 4-16.

Lesgold, A., Laijoie, S.P., Bunzo, M., and Eggan, G. (1990) A coached practice environment for an electronics troubleshooting job. In J. Larkin, R. Chabay and C. Shefic (Eds.), *Computer Assisted Instruction and Intelligent Tutoring Systems: Establishing Communication and Collaboration,* Lawrence Erlbaum Associates, Hillsdale, NJ.

Lewis, M., McArthur, D., Stasz, C., & Zmuidzinas, M. (1990) Discovery-based Tutoring in Mathematics. In Working Notes American Association of Artificial Intelligence Spring Symposium, AAAI, Palo Alto, CA.

Shute., V., Rose Garden Promises of Intelligent Tutoring Systems: Blossom or Thorn? (1990) Paper presented at the Space Operations, Applications and Research (SOAR) Symposium, also contact V. Shute, AFHRL, Brooks Air Force Base, TX 78235-5601.

Shute, V.J., Glaser, R., and Raghavan, K. (1989) Inference and discovery in an exploratory laboratory. In P.L. Ackerman, R.J., Sternberg, and R. Glaser (Eds.), Learning and Individual Differences, W.H. Freeman, New York, pp. 279-326.

Stevens, A., & Collins, A. (1977) The Goal Structure of a Socratic Tutor. In Proceedings of the Association for Computing Machinery Annual Conference. Also available as BBN Report No. 3518 from Bolt Beranek and Newman Inc., Cambridge, Mass., 02138.

Woolf, B. (1988) Intelligent Tutoring Systems: A Survey. In H. Shrobe and the American Association for Artificial Intelligence (Eds.), Exploring Artificial Intelligence, Morgan Kaufmann Publishers, Palo Alto, CA, pp. 1-43.

About the Authors

Beverly Park Woolf is an assistant professor in the Department of Computer and Information Science at the University of Massachusetts. Her main research focus is in human-machine communication, specifically in knowledge representation and control issues for building explanation, tutoring, advisory, user modeling and knowledge acquisition systems. She

holds a Ph.D. in Computer Science and a Ed.D. in Education, both from the University of Massachusetts.

Elliot Soloway is an Associate Professor of Electrical Engineering and Computer Science at the University of Michigan, Ann Arbor, MI. He directs the "Highly-Interactive Computing Environments" project located in the AI Lab.

William J. Clancey is a Senior Research Scientist at the Institute for Research on Learning, an independent, not-for-profit organization. His current interests are relating AI programming to traditional scientific modeling, studying computer systems in the workplace, and re-examining the relation of cognitive science theories to the processes of human memory and learning.

Kurt VanLehn is an associate professor in the computer science department and a senior scientist at the Learning Research and Development Center, both at the University of Pittsburgh. Dr. VanLehn's main research interest is in cognitive simulation and its applications to education. He holds a Ph.D. degree in Computer Science from MIT.

Dan Suthers holds a M.S. in computer science from the University of Massachusetts, where he is currently completing his Ph.D. His research is concerned with the design of computer-based media for knowledge communication, including planning explanations in an interactive context, representation and use of multiple perspectives and granularities in explanation, and multimedia interfaces.

Thoughts and Afterthoughts on the 1988 Workshop on Principles of Hybrid Reasoning

Alan M. Frisch and Anthony G. Cohn

The 1988 Workshop on Principles of Hybrid Reasoning, a one-day AAAI-sponsored workshop, was held in St. Paul, Minnesota on August 21, 1988, in conjunction with the National Conference on Artificial Intelligence. This article reports on the workshop and presents some of our afterthoughts based upon prolonged discussion of the issues that arose during the workshop. To a certain extent this article can serve as a survey of research on hybrid reasoning; to aid in this purpose we include numerous citations to the literature. All references can be found in the bibliography by Alan Frisch and Richard Scherl that accompanies this report.

Researchers in Artificial Intelligence recently have been taking an increasing interest in hybrid representation and reasoning systems—systems that consist of two or more integrated subsystems, each of which may employ distinct representation languages and inference systems. Though a number of such systems have been designed, studied, constructed, and put into use, little effort has been devoted to comparing the systems or searching for common principles underlying them.

The workshop addressed this need by bringing together a small number of leading researchers on hybrid reasoning for a day of intensive interaction. The workshop was organized by Alan Frisch (Workshop Chair), Ron Brachman, and Rich Thomason.

Each participant was invited to submit a short paper that best characterized their work on hybrid reasoning. The submissions were collected into a proceedings distributed to all participants prior to the workshop. As the submissions included previously-published papers as well as early drafts of work in progress, it was agreed at the workshop that the proceedings would be distributed no further. However, since most of the draft papers have subsequently appeared in published form. it is now possible to give a virtual proceedings. In the bibliography that accompanies this article published versions of the submitted papers are indicated with an asterisk.

Overview of the Workshop

The workshop program comprised seven invited talks, two moderated discussion sessions, one dinner, and two coffee breaks. We outline the talks and discussion sessions below and leave the dinner and coffee breaks as an exercise for the reader. We concluded with an informal late-night discussion on whether the workshop had been useful to the participants and whether we would like another one. The overall feeling was that the workshop had been useful and another would be desirable, but it should be either longer or devoted to discussion of selected position papers on a specific topic.

Characterization of Hybrid Knowledge Representation and Reasoning Systems

Peter Patel-Schneider began his presentation by categorizing various types of hybrid knowledge representation and reasoning systems. A system can qualify as hybrid by employing multiple representations or by employing multiple reasoning methods, thus suggesting a characterization of hybrid systems along these two dimensions. Along the representation dimension, a system can have multiple redundant representations of the same knowledge in different media—as in the vivid reasoning system (which we shall call VIVID) presented by Brachman and Etherington (Etherington *et al.* 1989) and the multiple reasoners at a single layer of the CAKE architecture (Rich 1985)—or it can have different representations for different kinds of knowl-