

Constructed vs. Received Representations for Learning about Scientific Controversy: Implications for Learning and Coaching

Violetta Cavalli-Sforza (violetta@cs.cmu.edu)

Language Technologies Institute, Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15213 USA

Abstract

The development of a graphical representation for performing a task can potentially yield a greater understanding of the task domain, but it is itself a demanding task that can distract from the primary one of learning the domain. In this research, we investigated the impact of constructing versus receiving a graphical representation on learning and coaching the analysis of scientific arguments. Subjects studied instructional materials and used the Belvedere graphical interface¹ to analyze texts drawn from an actual scientific debate. One group of subjects used a box-and-arrow representation, augmented with text, whose primitive elements had preassigned meanings tailored to the domain of instruction. In the other group, subjects used the graphical elements as they wished, thereby creating their own representation.

Our results support the following conclusions. From the perspective of learning target concepts, developing one's own representation may not hurt those students who gain a sufficient understanding of the possibilities of abstract representation, although there are costs in time on task and in the quality of the diagrams produced. The risks are much greater for less able students because, if they develop a representation that is inadequate for expressing the concepts targeted by instruction, they will use those concepts less or not at all. From the perspective of coaching students, a predefined representation has a significant advantage. If it is appropriately expressive for the concepts it is designed to represent, it provides a common language and clearer shared meaning between the student and the coach, enabling the coach to understand students' analysis more easily and to evaluate it more effectively against a model of the ideal analysis.

Background

The initial goal of our research was to develop tools for helping students in middle school through early university to engage in critical thinking. We chose scientific controversies as a domain of application, because understanding the various assumptions and reasons underlying scientific debate and change is an important part of appreciating and communicating the nature of the scientific enterprise, even to non-scientists. In contrast, until recently, school science curricula have exhibited a strong bias towards portraying scientific endeavor solely as hands-on *inquiry* and on *final-form* science, that is, the knowledge currently accepted by the scientific community (Duschl, 1990).

As part of our tool base, we developed graphical argumentation formalisms that would enable students to determine and represent how the variety of information brought to bear

in a scientific debate fits together. We investigated the use of coaching strategies to help students develop correct analyses of textual arguments and to generate their own arguments (Paolucci *et al.*, 1996; Toth97 *et al.*; Cavalli-Sforza, 1998). Some of the resulting tools and materials were made available in selected schools and over the Internet (Suthers *et al.*, 1997). Related work, focusing on belief change more than on argument structure, is described by Schank (1995).

Constructed vs. Received Graphical Representations

Although there have been several efforts to develop environments for supporting policy and design discussions (Conklin & Begemann, 1988; Fischer, & McCall, 1989; Stefik *et al.*, 1987; Tatar *et al.*, 1991), authoring of arguments and argumentative text (Neuwirth & Kaufer, 1989; Schuler & Smith, 1990; Smolensky *et al.*, 1987; Streitz *et al.*, 1989), and exploration of the structure of information (Lowe, 1985; Marshall *et al.*, 1991), to our knowledge, there has been little empirical investigation of the use of diagrammatic representations in argument-like tasks and its effect on learning and coaching.

Subjects working with ConvinceMe (Schank, 1995), benefited from comparing their own beliefs and argument networks to the program's, relative to subjects working exclusively with text. The active construction, rather than the passive examination, of diagrams improves performance for verifying the validity of syllogisms when combined with a precise method for developing the diagram from the problem description (Grossen & Carnine, 1990).

In contrast, developing one's own procedure for problem solving may lead to learning that generalizes more easily to different problem types, but may also lead to "buggy" algorithms. Diagrams help solve analytical reasoning problems similar to those in the GRE, but only for subjects who show an initial aptitude for that type of reasoning (Cox *et al.*, 1994), indicating that there may be individual differences in the ability to use diagrams.

Finally, diagrammatic representations can assist in understanding computer programs, but only if the diagram supports the type of reasoning required by the task; even so, performance may be slower than with text representations (Green & Petre, 1992).

Although these results don't unanimously support the advantages of graphical representations, we believed that a diagrammatic representation of argument structure was an important part of any set of tools for helping students understand and engage in scientific argument, since, by virtue of its ex-

¹This system was a conceptual predecessor of the system described by Suthers *et al.* (1997).

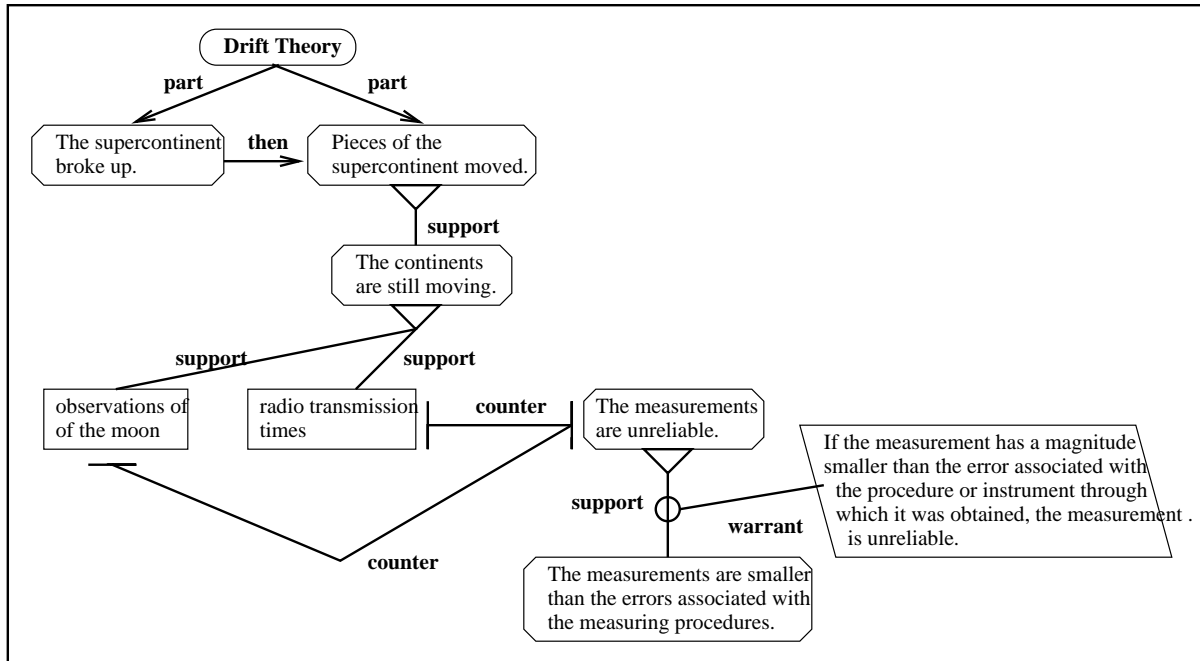


Figure 1: Graphical representation of an argument and an attack on its evidence

plicitness, it would help students ferret out from the texts the information and assumptions required by the argument.

The advantages of a box-and-arrow representation for displaying arguments are discussed in Cavalli-Sforza (1998). Very briefly, in a representation such as the one shown in Figure 1 (loosely based on Toulmin's (Toulmin *et al.*, 1984) model of argument steps), the boxes and arrows correspond to the main components of argument structure: different types of statements (positions and data) and relationships between them (supporting and contradictory, or other negative relationships). The graphical depiction of argument structure makes clear how different information is used to create an argument in favor of a claim (a rule, hypothesis or theory) or to refute that claim or the argumentation surrounding it. The presence or absence of relationships between statements also indicates opportunities for providing further support or refutation. For example, it is possible to refute a claim by arguing directly against it (making or strengthening an argument in favor of a contradictory claim), by arguing against any of the data or warrants used in supporting a claim, or by "undercutting" the support, that is, arguing that the data does not really support that claim.

The predefined graphical representation shown in Figure 1 was designed to have a good "cognitive fit" with the task, in the sense of focusing subjects' attention on the basic concepts targeted by the instructional materials and providing the kind of processing advantages that might be expected from graphical representations. There was no clearly comparable alternative representation. We tried, for example, tabular formats for representing arguments, but they lacked the richness required to represent the range of information present in an argument. In the end, we decided to let some subjects develop their own representation after providing two suggestions: a

box-and-arrow type representation for which they could determine the meaning of the shapes and links provided, and labeled-text schemata (see Figure 2). We hypothesized that the effort of developing or extending a representation might lead to carry out deeper processing of the target concepts of argument analysis.

Coaching

A further reason for using graphical representations was to provide an explicit and shared language for describing argument structure. Such an external representation facilitates interaction between student and coach by providing concrete targets of discussion during a coaching interaction. We hypothesized that the a box-and-arrow graphical representation whose primitives encode key concepts of argumentation provides a more effective channel for coaching subjects in applying those concepts than a graphical representation developed by the subjects themselves, especially for computer coaches (Paolucci *et al.*, 1996; Toth *et al.*, 1997). In this study, however, the coach was the experimenter and coaching advice was provided based on a general mental model of the argument contained in the texts being analyzed. Because there was usually more than one way of representing, or even interpreting, the content of the text, this approach was deemed more flexible and appropriate than comparing a subjects' work to a specific "expert" diagram. For similar reasons, we applied a generally non-interventionist coaching philosophy. The coach waited for the subject to produce at least a partial draft of the analysis before beginning to comment on it, unless the subject appeared to need or specifically requested assistance.

Two classes of strategies were used in providing assistance. *Scaffolding strategies* were applied when the subject seemed

<p>PREMISES: Space is very cold. The earth was very hot when it formed.</p> <p>WARRANT: The laws of heat transfer: An object at a different temperature from its surroundings will eventually come to a common temperature with them. The rate of cooling or warming of the object is approximately proportional to the difference in temperature.</p> <p>CONCLUSION: The young earth began to cool by losing heat to space.</p>	<p>CLAIM: The inside of the earth is still very hot.</p> <p>WARRANT: The temperature throughout an object is at least as great as that of material exiting from the object.</p> <p>REASONS: Hot lava comes out of volcanoes.</p>
--	--

Figure 2: Two examples of a labeled-text schema used for a causal inference and an argument step

unable to proceed with a task and included: 1) reminding the subject of the goal and progress to date; 2) asking a question about a step in the solution; 3) suggesting a high-level plan for the subject to follow; 4) providing a limited number of choices for the subject to explore; 5) giving specific hints about what path(s) to follow; 6) suggesting (some of) the crucial steps in a solution; 7) performing some parts of the plan; 8) leading the subject through each step of a plan by asking a series of questions; 9) performing the task for the subject. *Coaching strategies* were applied to comment on the subject's work after a (partial) analysis had been produced; they included: 1) signaling a potential problem; 2) suggesting information to consider; 3) criticizing; 4) correcting; 5) explaining; 6) arguing.

The application of scaffolding and coaching strategies was interwoven. For example, a coaching strategy might be used to evaluate the subject's analysis, which would lead to the subject attempting to improve the analysis. This might call for a scaffolding strategy to help the subject carry out the improvement. Within the two general classes of strategies, individual strategies differ in the amount and type of knowledge required, as well as in the degree of coach engagement in the task. The coach always attempted to begin with low knowledge and low engagement strategies, which require the subject to do more work, and to progress to more demanding strategies only when the subject needed more assistance or feedback.

The Experiment

Four subjects, non-science majors at the University of Pittsburgh, participated in an extended experiment which included the following sessions.

- **Pre-Test and Post-Test Sessions.** In both sessions, subjects gave definitions of terms relevant to scientific argument and answered questions on an article concerning the debate about the Cretaceous-Tertiary ("dinosaur") extinctions.
- **Belvedere Sessions.** Subjects worked with the Belvedere environment for six sessions. The first four were training sessions. Subjects began with a tutorial in which they learned to use the Belvedere environment's graphical and

text capabilities. Then they read instructional materials discussing concepts of scientific argumentation, and short domain texts describing two competing theories about the formation of mountains continents and ocean basins. These texts were drawn from the debate between the supporters of Continental Drift theory and the Contraction theory. The texts were chosen to exemplify different aspects of scientific argument. Some texts described the theories and showed how they explained the data, others provided different types of arguments in favor of each theory (e.g. based on data, based on analogy), others yet were dialectical texts debating the merits of the two theories relative to the data and the characteristics of the explanation (e.g. uniqueness, breadth, parsimony, internal consistency, etc.). In the last two sessions on Belvedere, subjects were given similar texts concerning a third theory, the Expansion theory, and were asked to perform comparable analyses.

Two subjects worked in each of two experimental conditions. The instructional materials, which contained the same conceptual content, introduced a related set of concepts and showed how they could be applied in analyzing a domain text. In the **FIXED** condition, subjects saw a box-and-arrow diagram representing the sample text analysis (Figure 1); in the **FREE** condition, the same analysis was carried out via one or more labeled-text schemata such as those in Figure 2. Subsequently, subjects performed similar analysis on parallel texts, equivalent in structure to the ones used in the instruction, and were coached by the experimenter according to the approach previously described. Coaching remained available to subjects in the last two sessions as well, but they no longer received the explicit guidance of the instructional materials.

- **Review Sessions.** Three sessions were spent reviewing argument concepts through a comparison of the theories relative to different criteria, further oral analysis of domain texts, and a summary of the debate.

Altogether, each subject spent between 28 and 38.5 hours in the experiment. The Belvedere sessions, ranging between 16.5 and 19.5 hours per subject, were videotaped. The videotapes were transcribed and indexed to the evolution of the diagrams produced by subjects. The review sessions were audiotaped.

Data Analysis Methodology

Questionnaire Data

The pre-test and post-test questionnaires required subjects to provide definitions for several of the terms introduced in the instructional materials. These included terms describing the status of scientific statements (e.g. theory, analogy, model), argument terms (e.g. claim, warrant, counterargument, undercut), and criteria for evaluating scientific explanations (e.g. explanatory breadth, explanatory parsimony, inconsistency and conservativeness of an explanation, sufficiency of the proposed mechanism). In total there were 37 terms, of which 27 appeared in both pre-test and post-test questionnaires; 10 appeared only in the post-test questionnaire because subjects were not expected to be familiar with these terms prior to the experiment.

For each term in the questionnaire, subjects' answers were compared to a target definition, such as was used in the instructional materials. Two raters assigned scores independently. The scores were then discussed and occasionally modified. The resulting post-discussion inter-rater correlations were 0.83 (N = 108) for items on both pre-test and post-test, and 0.86 (N = 40) for items only on the post-test.

Interaction Data

The protocols for roughly one third of the many graphs that subjects drew were segmented into units of interaction between the subject and the coach, and were analyzed quantitatively and qualitatively as follows.

For each diagram, the raw protocol of the interaction was synthesized into an *interaction summary*. First, the protocol was segmented into groups of related utterances by the subject and/or the coach. Each *utterance group* represents a self-contained thought by one person and abstracts away interruptions and simultaneous speech by the other person that were not extending the original thought in a significant way. The utterance group was then summarized into a *summary utterance*, a single statement that reflects the essential content of the utterance group. One or two summary utterances, taken together, form one *interaction unit*. An interaction unit can include:

- a summary utterance by the subject and one that represents the coach's response,
- a summary utterance by the coach with a verbal response or action by the subject, or
- a summary utterance by the coach alone, for example a comment praising the subject's work.

Several interaction units may be grouped into a single *interaction group* to indicate that they all pertain to a specific issue or topic of discussion. Interaction groups may be nested to show the overall structure of the interaction. The top-level group has a heading describing the main topic of the interaction, with nested interaction group headings giving the various subtopics addressed.

Each interaction unit was also classified on the basis of the primary **content** as well as the communicative **form** of the interaction. The content categories include: rhetorical

structure, representation, scientific domain, graphical interface, and other. An interaction unit dealing primarily with the interrelationship of two statements in the domain text was coded as rhetorical structure; an interaction unit dealing primarily with the appropriate link, link name, or linkage pattern to use to represent that relationship in the graph was coded as representation.

The communicative form categories represent different communicative actions used by the coach to help the subject improve his or her graphical analysis of a text. They included, among others: review, remind, check subject's understanding of text, explain, give an example, ask for justification, check coach's understanding of subject's intentions, check coach's understanding of subject's work, ask a leading question about the text/diagram, hint at the characteristics of a solution, suggest possible actions, tell what to do, correct by giving the solution, argue, draw attention to a problem, criticize, praise.

An example of an interaction unit composed of two summary utterances is the following, which was analyzed as having content "rhetorical structure" and form "hint". Although the subject is speaking in terms of graphical links (e.g. *and*), the coach is hinting that simple conjunction is not the correct relationship among the statements.

S: I'll use an *and* to show the relation between items 1, 2, 3.

C: You should be more specific than *and*. It has to do with time sequence.

For a subset of the protocols, where complex or problematic interactions had occurred, we also examined the interaction qualitatively. We analyzed the application of coaching strategies as embodied by the specific coaching actions, and their degree of success or failure in terms of bringing about the desired results.

Results and Discussion

Use of Graphical Representation

A detailed analysis of the diagrams produced by FIXED condition subjects showed that the predefined argument-specific graphical representation was an appropriate representation for the task. Both subjects used most links and shapes without difficulty; errors were usually due to an error in analyzing the domain content or the argument structure of the text. The representational problems subjects encountered were often localized, for example, confusion over the direction of a link due to naming ambiguities, and omission of links, which subjects created when needed.² Subjects were also able to combine primitive links in correct and creative ways. Finally, subject diagrams were often quite good even before the experimenter intervened, and were relatively easy to read even when they became crowded.

Whereas in the FIXED representation condition both subjects appeared to be largely successful in applying the concepts they had learned, subjects in the FREE representation condition differed significantly in this regard. One subject used a box-and-arrow type formalism to encode many

²There were, however, some interesting misuses of devices for representing conjunction.

of the target relational concepts, which enabled him to analyze domain texts at an appropriate level of detail and to represent complex structural relationships, especially for argument diagrams. For causal diagrams (theory description diagrams), the subject had trouble abstracting away from domain-specific relationships towards more general temporal, causal, and explanatory relationships; consequently these diagrams were difficult to understand. Graphical primitives were often used inconsistently across both types of diagrams. The other subject did not find a way of using the box-and-arrow graphics to represent either type of content abstractly. This subject's causal diagrams were essentially analogical renderings of the physical situation; the argument diagrams exclusively employed labeled-text schemata, which emphasized the role played by different statements over their relationship, and displayed little structural detail. Since this subject did not extend the schemata to represent dialectical relationships, the diagrams did not show the relationships of the contents of the text to other parts of the debate any more explicitly than the original text.

Learning of Target Concepts

We had hypothesized that, by visually reinforcing the terminology and the target concepts, the FIXED representation would help subjects learn those concepts better than subjects in the FREE condition. However, subjects' scores for the 37 questionnaire items revealed little or no difference. Based on absolute performance scores in the Post-Test, subjects fell into a stronger and a weaker group, each containing a subject from each condition. Within the same group the FIXED condition subject was slightly ahead. Improvement scores showed the weaker subject of the FREE condition (Free-2), who started with the lowest scores of any subject, consistently improving more than the weaker subject in the FIXED condition (Fixed-2). Of the stronger subjects, Fixed-1 improved substantially more than the Free-1.

Aside from the small number of subjects, several factors may have contributed to the inconclusiveness of these results. The questionnaire was not an adequate means of assessing learning, since it tested subjects' ability to verbalize concepts, whereas the training was geared towards applying them. Examining the diagrams subjects drew, especially diagrams drawn in later sessions, we found that representation was related to subjects' ability to use the target concepts in their diagrams, if not always to their ability to supply proper definitions. Subjects in the FIXED condition consistently distinguished the status of statements in their diagrams (as observation, explanation, plain statement, etc.) because the representation required them to do so. FREE condition subjects seldom, if ever, did. Only subjects in the FIXED condition and Free-1, who developed a sufficiently expressive set of graphical primitives, consistently represented the relational concepts present in the text, especially multi-step support and dialectical relationships. For the more complex concepts that were realized in box-and-arrow diagrams with distinctive linkage patterns (e.g. analogy/model, parsimony of an explanation), there was a link between subjects' use of the concept in the diagram and their ability to give good definitions. The crucial factor underlying these results appears to be whether the representation is relation-centered, as box-and

arrow representations are, or role-centered as the labeled-text schema adopted by Subject Free-2 is.

Coaching

The results of the analysis of interaction data supports the hypothesis that the FIXED representation condition was a more effective vehicle for coaching in several ways.

Amount of Coaching. Subjects in the FIXED condition required less coaching overall, as measured by the number of interactions units with coaching content.

Target of Coaching. Subjects in the FREE condition required a significantly greater proportion of coaching about representation itself, at the expense of verifying that the diagram encoded the text correctly and completely. This interaction, negotiating the meaning and assessing the adequacy of the representation, did not clearly benefit subjects' understanding of the structure and contents of the text.

Directiveness of Coaching. Subjects in the FIXED condition could be coached using less directive strategies, allowing the subjects to do more of the work. Most problems in FIXED condition diagrams were localized and could be fixed with a brief interaction; often the coach only needed to point out that there was a problem or to hint at the nature of the problem. In contrast, the experimenter often needed to coach FREE condition subjects in greater detail.

Ease and Acceptance of Coaching. From the coach's perspective, interaction with FIXED condition subjects was easier, more systematic and more effective than with FREE condition subjects. FIXED condition subject diagrams, especially causal diagrams, were easier for the coach to understand and to compare to her own model(s) of the correct analysis, allowing problems areas to be identified and an agenda for coaching to be quickly established. Problems requiring some global restructuring of the diagram could be addressed by applying coaching and scaffolding strategies more systematically, reducing the problem to subproblems that were smaller and easier to correct. This was partially attributable to subjects having a stock of diagramming primitives that enabled them to put the advice into effect more readily. In contrast, subjects in the FREE representation condition experienced significantly more difficulty in implementing the experimenter's advice. Subject Free-1 actually resisted accepting advice, especially on problems related to representation, preferring instead to keep on experimenting with alternative representations of his own invention. The limitations of the representations used/developed by the other subject contributed to this subject's difficulty in understanding and implementing coaching advice.

Summary and Conclusions

We investigated whether a predefined graphical representation whose primitives encode key concepts of scientific argument is more effective for learning those concepts than a graphical representation learners develop themselves. Although, we did not find evidence that such a predefined representation improves ability to *define* target concepts, there was evidence that the representation subjects ultimately adopt impacts whether they are able to *use* those concepts in their

diagrams. The risk of allowing learners to construct their own graphical representation is that they may not succeed in developing a representation that is sufficiently expressive, making the diagram construction task an ineffective vehicle for learning and coaching. Even if it is sufficiently expressive, there is a cost in time and clarity of the diagrams and coaching is considerably more effortful. In contrast, a carefully designed task-appropriate representation is easily adopted and produces diagrams that encode the concepts targeted by instruction clearly and correctly. Moreover, such a representation facilitates coaching interaction over the instructional content of the task. Unlike a representation under development, which requires constant negotiation over its meaning, it provides a well-defined language for communication between the learner and the coach.

Acknowledgments

This research was funded under the National Science Foundation's Applications of Advanced Technology program under the title "Tools for Thinking About Complex Issues in Science and Public Policy," Grant MDR-9155715.

We gratefully acknowledge the assistance of Dr. Alan Lesgold, who supervised this research, and of other members of the research group at the Learning Research and Development center, University of Pittsburgh: John Connelly, Massimo Paolucci, Dan Suthers, and Arlene Weiner.

References

- Cavalli-Sforza, V. (1998). *Constructed vs. received graphical representations for learning about scientific controversies: Implications for learning and coaching*. Doctoral dissertation, Intelligent Systems Program, University of Pittsburgh.
- Conklin, J., & Begeman, M. L. (1988). gIBIS: A hypertext tool for exploratory policy discussion. *ACM Transactions on Office Information Systems*, 6(4), 303–331.
- Cox, R., Stenning, K., & Oberlander, J. (1994). Graphical effects in learning logic: reasoning, representation and individual differences. *Proceedings of the Sixteenth Annual Conference of the Cognitive Science Society* (pp. 237–242). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Duschl, R. A. (1990). *Restructuring science education. The importance of theories and their development*. New York, NY: Teachers' College Press, Columbia University.
- Fischer, G., McCall, R., & Morch, A. (1989). JANUS: Integrating hypertext with a knowledge-based design environment. *Hypertext '89 Proceedings* (pp. 105–117). Baltimore, MD: Association for Computing Machinery.
- Green, T. R. G., & Petre, M. (1992). When visual programs are harder to read than textual programs. In G. C. van der Veer, M.J. Tauber, S. Bagnarola, & M. Antavolits (Eds.), *Human-Computer Interaction: Tasks and Organization*. Rome: CUD.
- Grossen, B., & Carnine D. (1990). Diagramming a logic strategy: Effects on difficult problem types and transfer. *Learning Disability Quarterly*, 13, 168–182.
- Lowe, D.G. (1985). Co-operative structuring of information: the representation of reasoning and debate. *International Journal of Man-Machine Studies*, 23, 97–111.
- Marshall, C. C., Halasz, F. G., Rogers, R. A., & Janssen Jr., W. C. (1991). Aquanet: A hypertext tool to hold your knowledge in place. *Hypertext '91 Proceedings*. Baltimore, MD: Association for Computing Machinery.
- Neuwirth, C. M. & Kaufer, D. S. (1989). The role of external representations in the writing process: implications for the design of hypertext-based writing tools. *Hypertext '89 Proceedings* (pp. 319–342). Baltimore, MD: Association for Computing Machinery.
- Paolucci, M., Suthers, D., & Weiner, A. (1996). Automated advice-giving strategies for scientific inquiry. In C. Frasson, G. Gauthier, & A. Lesgold (Eds.) *Intelligent Tutoring Systems, Third International Conference, ITS '96. Lecture Notes in Computer Science*. New York, NY: Springer.
- Schank, P. (1995). *Computational tools for modeling and aiding reasoning: Assessing and applying the Theory of Explanatory Coherence*. Doctoral dissertation, University of California, Berkeley.
- Schuler, W. & Smith, J. B. (1990). Author's argumentation assistant (AAA): A hypertext-based authoring tool for argumentative tasks. In A. Rizk, N. Streitz, and J. André (Eds.), *Hypertext: Concepts, systems, and applications*. Cambridge University Press, Electronic Publishing Series.
- Smolensky, P. et al. (1987). *Computer-aided reasoned discourse, or, how to argue with a computer*. (Tech. Rep. CU-CS-358 87). University of Colorado, Department of Computer Science.
- Stefik, M., Foster, G., Bobrow, D. G., Kahn, K. Lanning S., & Suchman L. (1987). Beyond the chalkboard: Computer support for collaboration and problem solving in meetings. *Communications of the ACM*, 30(1), 32–47.
- Streitz, N. A., Hannemann, J., & Thüring, M., 1989. From ideas and arguments to hyperdocuments: travelling through activity spaces. *Hypertext '89 Proceedings* (pp. 343–364).
- Suthers, D., Erdosne Toth, E., & Weiner, A. (1997). An Integrated Approach to Implementing Collaborative Inquiry in the Classroom. *Proceedings of Computer Supported Collaborative Learning*.
- Tatar, D. G., Foster, G., & Bobrow, D. G., (1991). Design for conversation: Lessons from Cognoter. *International Journal of Man-Machine Studies*, 34, 185–209.
- Toth, J. A., Suthers, D., & Weiner, A. (1997). Providing expert advice in the domain of collaborative scientific inquiry. *Proceedings of Artificial Intelligence in Education* (pp. 302–308).
- Toulmin, S, Rieke, R., & Janik A. (1984). *An introduction to reasoning*, 2nd edition. New York, NY: Macmillan.