

Representational Support for Collaborative Inquiry¹

Dan Suthers

Department of Information and Computer Sciences

University of Hawai'i

1680 East West Road, POST 303A

Honolulu, HI 96822

suthers@hawaii.edu

Abstract

Empirical work with “Belvedere,” a software environment for the construction of diagrammatic representations of evidential relations, is summarized, leading to the hypothesis that variation in features of representational tools can have a significant effect on the learners’ discourse and on learning outcomes. For example, by manipulating the concepts used by a toolkit, it is possible to manipulate the distinctions attended to by learners. Once learners have constructed some representations, their learning interactions appear to be further guided by the objects and relationships (expressed or potential) that these representations make salient. These kinds of design considerations are critical for collaborative learning software, yet are insufficiently studied. This paper describes the work that led to this position, sketches a theoretical analysis of the roles of constraint and salience in the effect of representational bias on collaborative learning discourse, and describes initial studies now underway that varies textual, diagrammatic and tabular representations, and investigates their effects on the quality of discourse between learners.

1. Introduction

Decades of research into cognitive and social aspects of learning has developed a clear picture of the importance of learners’ active involvement in the expression, examination, and manipulation of their own knowledge (e.g., [8, 42, 50]), as well as the equal importance of guidance provided by social processes and mentorship [5, 33, 54, 68]. Recently these findings have been reflected in software technology for learning: systems are now providing learners with the means to construct and manipulate their own solutions while being guided by the software and interacting with other learners. The present

work is within this spirit, providing representational tools in support of collaborative learning.

Representational tools for data manipulation and knowledge mapping range from basic office tools such as spreadsheets and outliners to “knowledge mapping” software. Such tools provide representational guidance that help learners see patterns, express abstractions in concrete form, and discover new relationships [15, 32]. These representational tools can function as *cognitive* tools that lead learners into certain knowledge-building interactions [11, 23, 31]. The present work is based on the hypothesis that properly designed representational tools guide collaborative as well as individual learning interactions. Specifically, as learner-constructed external representations become part of the collaborators’ shared context, the distinctions and relationships that are made salient by these representations may influence their interactions in ways that influence learning outcomes.

Within classroom and laboratory settings, *collaborative learning* has been shown to correlate with greater learning, increased productivity, more time on task, transfer of knowledge to related tasks, higher motivation, and heightened sense of competence [20, 46, 52, 54]. Similarly, collaborative use of instructional software can be at least as effective as individual use [21, 24, 67]. However, collaborative learning does not work for all learners, and the results of instructional outcome studies are mixed. We cannot expect learning gains just because learners are sitting together [6]. Thus, a goal of computer supported collaborative learning (CSCL) systems [29, 40] – and of this work – is to improve the effectiveness of collaborative learning as an instructional format: i.e., to support peer interactions in a manner that increases learning gains. In the K-12 environment, collaborative learning with computers can also be a necessity due to resource limitations. In postsecondary distance education, electronic forms of collaborative learning can help reduce the isolation of telecommuting

¹Copyright 1999 IEEE. Published in the Proceedings of the Hawai’i International Conference on System Sciences, January 5-8, 1999, Maui, Hawai’i.

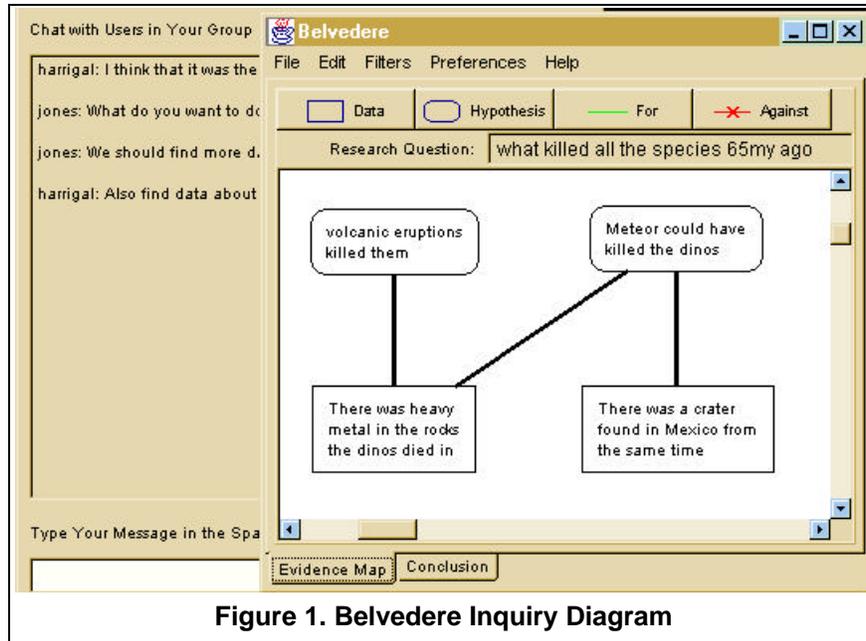


Figure 1. Belvedere Inquiry Diagram

learners and increase the interactivity of the distance learning experience [1, 22].

The hypothesis of this paper – that representational tools can influence collaborative learning discourse in educationally significant ways – became apparent to the author while working with a networked software system, designed and deployed by the author with colleagues,² that provides learners with shared workspaces for coordinating and recording their collaboration in scientific inquiry. During classroom and laboratory use of “Belvedere,” the author has observed that by manipulating the concepts used by a toolkit, it is possible to manipulate the distinctions attended to by learners. Furthermore, once learners have constructed some representations, their learning interactions appear to be further guided by the objects and relationships (expressed or potential) that these representations make salient. Based on these observations, the author is undertaking a systematic study of how variation in features of the representational tools provided by these environments can have a significant effect on the learners’ knowledge-building discourse and on learning outcomes.

This paper begins with a description of the Belvedere software environment. Then, examples of the kinds of interactions that led to the representational bias hypothesis are provided, followed by a theoretical account and an outline of the research program that is presently underway.

2. Software for Collaborative Inquiry

The “Belvedere” software is a networked software system that provides learners with shared workspaces for coordinating and recording their collaboration in scientific inquiry. The version described in this paper, Belvedere 2.x, is a complete redesign and reimplement of Belvedere 1.x, previously reported in [62, 63]. Belvedere’s core functionality is a shared workspace for constructing “inquiry diagrams,” which relate data and hypotheses by evidential relations (consistency and inconsistency). The software also includes an artificial intelligence “advisor” [39, 64], a “chat” facility for unstructured discussions, and facilities for integrated use with Web browsers. The diagramming window is shown in Figure 1, with an additional window (left side) for a “chat” facility. The default “palette” (the horizontal row of icons near the top of Figure 1) makes salient the most crucial distinctions we want learners to acquire in order to conduct scientific inquiry.

Left to right, the icons are “data” for empirical statements, “hypothesis” for theoretical statements, then links representing “for” and “against” evidential relations. Learners use the palette by clicking on an icon, typing some text (in the case of statements) and optionally setting other attributes, and then clicking in the diagram to place the statement or create the link. The palette is configurable: other categories and relations can be added, such as “unspecified” statements about which learners disagree or are uncertain, “principle” for law-like

² See acknowledgments at end.

statements, and a link for conjunction. An icon for an automated “coach” can also be added.

Other features of Belvedere, briefly noted, include the following. Users can set different “belief levels” for the statements and relations, and display these as line thickness with a “filter.” References to external objects can be sent from other applications directly into the Belvedere workspace. Java applets have been embedded in the Web-based curricular materials, enabling learners to send references to these pages into the workspace with a click of a button. The feasibility of extending this mechanism to other kinds of documents, such as MS Word and Excel documents, has been demonstrated, and it is possible to reinvoke these applications in a platform independent manner. Thus Belvedere can be used as a conceptual organizer for use of various tools during an inquiry. Koedinger, Suthers & Forbus [28] have also demonstrated coordinated integration of Belvedere with a tutoring agent and a simulation engine.

2.1 Software Implementation

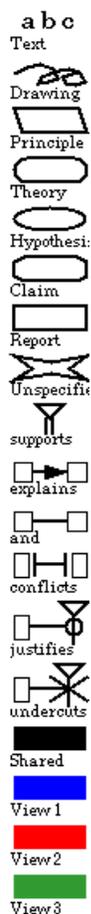
The Belvedere application is written in Java, and is available for MacOS, Windows ‘95, NT, and Solaris. It is deployed as a client within a networked architecture that is designed to provide affordable access to intelligent collaborative educational functionality on a variety of desktop user platforms. Belvedere 2.1 includes a lightweight, NT or Unix-based server consisting of a SQL database (currently msql or postgres) for shared workspaces and a Java “Connection Manager” that provides for “what you see is what I see” updating of clients sharing the same workspace. Each Belvedere 2.1 Java application (or, more recently, applet) includes the evidence mapping facility, the built-in advisor, and a JDBC interface to the database and Connection Manager. Belvedere applications can also run stand-alone, saving workspaces to the local disk but not providing distance collaboration functionality. See [60] for a discussion of the more complex CGI-based architecture of Belvedere 2.0 and other aspects of the design.

2.2 Classroom Implementation

In coordination with the software development, the author’s colleagues³ developed a comprehensive method for implementing Belvedere-supported collaborative inquiry in the classroom. The approach includes student activity plans worked out in collaboration with teachers. Students work in teams to investigate real-world “science challenge problems,”⁴ designed with attention to National Science Education Standards to match and enrich the

curriculum. A science challenge problem presents a phenomenon to be explained, along with indices to relevant resources. The teams plan their investigation, perform hands-on experiments, analyze their results, and report their conclusions to others. Investigatory roles are rotated between hands-on experiments, table-top data analysis, and computer-based literature review and use of simulations and analytic tools as well as Belvedere. Assessment rubrics are given to the students at the beginning of their project as criteria to guide their activities. The rubrics guide peer review, and help the teacher assess nontraditional learning objectives. See [61] for further information on this integrated approach to classroom implementation, as well as discussion of a third-party evaluation. The present paper focuses on representational issues that arose from informal observations during use in the classroom and laboratory sessions, which have resulted in a new line of work.

3. Representations and Discourse



Belvedere 1.0 was initially used with students aged 12-15 working alone or in pairs in our lab, as well as by students working in small groups in a 10th grade biology classroom. Belvedere 2.x is under use by 9th and 10th grade science classes in Department of Defense Dependent Schools overseas. During this time we learned some important lessons about the role of external representations in collaborative learning.

Belvedere 1.0 was designed under the assumption that a visual representation language (augmented with automated advice-giving) can help students learn the nuances of scientific argumentation, provided that (a) the language is *capable of capturing all of these nuances*, and (b) *students express their arguments in the language*. Guided by (a), Belvedere 1.0 was provided with a rich palette of statement types and relationships. One of the Belvedere 1.0 palettes is shown to the left.

3.1 Locus of Discourse

As indicated by (b) above, we expected students to express all of their significant argumentation in the diagrams using primitives such as these. However, we found that much relevant argumentation was “external,” arguing *from* the representations rather than arguing *in* the representations. Faced with a decision concerning some manipulation of the representations, students would begin to discuss

³ Primarily Eva Toth and Arlene Weiner.

⁴ Available at <http://lilt.ics.hawaii.edu/belvedere/materials/>.

substantial issues until they reached tentative agreement concerning how to change the representation. In the process, statements and relations we would have liked students to represent went unexpressed.

Our initial frustration soon gave way to an understanding that this is an opportunity: proper design of manipulable representations can guide students into useful learning interactions. Thus, we downplayed the following potential roles of the representations: (1) as a medium *through* which communication takes place, (2) as a complete record of the argumentation process, and (3) as a medium for expressing formal models – in favor of their role in stimulating and guiding collaborative learning discourse.

3.2 Discussion Initiated by Categorical Choices

Belvedere requires all knowledge units (statements and relations) to be categorized at the time of creation. We often observed that learners who were using Belvedere initiated discussion of the appropriate categorical primitive for a given knowledge unit when they are about to represent that unit [59]. Although this is not surprising, it is a potentially powerful guide to learning, provided that it happens at the right time, and that discussion focuses on the underlying concepts rather than the interface widget to select. For example, consider the following interaction in which students were working with a version of Belvedere, that required all statements to be categorized as either “data” or “claim.” (The example is from a videotape of students in a 10th grade science class.)

S1: So data, right? This would be data.

S2: I think so.

S1: Or a claim. I don’t know if it would be claim or data.

S2: Claim. They have no real hard evidence. Go ahead, claim. I mean who cares? who cares what they say? Claim.

The choice forced by the tool led to a peer coaching interaction on a distinction that was critically important for how they subsequently handled the statement. The last comment of S2 shows that the relevant epistemological concepts were being discussed, not merely which toolbar icon to press or which representational shape to use.

It is not always useful to confront learners with choices, even if they may become important at some point in the development of expertise. For example, in other interactions with a more complex version of this tool that provided more categories, we saw students’ discussion sidetrack on subtle distinctions that were not important at their stage of learning:

S_M: “So what would that be...”

S_E: “Uhh...”

S_M: “An **ob--**”

S_E: “A **claim?**”

S_E consults sheet of paper in front of her; [pause]
“How about a **law?** scientific color?”

S_M: “Do you want to say a **warran--** uhh, no.”

S_E?: “Wait, what’s a **warrant?** I just read that; why some things...”

S_M: “[sigh] Oh dear.”

S_E: “Kind of like a **law**, like ...” [pause]

S_M: “Yeah, but there are exceptions, I think, because it can’t travel, like, complete distances.”

It was not necessary for these students to be struggling with all of these concepts at the outset of their learning experience, although S_M’s final utterance was a nice touch. These observations led us to simplify Belvedere’s representational framework to focus on the most essential distinction needed concerning the epistemological source of statements: empirical (“data”) versus hypothetical (“hypothesis”). Further simplifications, described in sections 3.3 and 3.4 below, were motivated by similar observations concerning the use of relations (links).

Interactions such as these suggested that (1) by manipulating the primitive concepts and relations provided by a representational formalism, it is possible to manipulate *what* distinctions are attended to by learners; (2) by manipulating when the software requires that a category be chosen, it is possible to manipulate *when* these distinctions are attended to. Combining this control with a cognitive analysis of a given learning task, we believe it is possible to design interface interaction sequences that draw learners’ attention to the right distinctions at the right time [25, 27].

3.3 Ontological Clarity

The original set of argumentation relations included evidential, logical, causal, and rhetorical relations as well as the various classifications of statements exemplified above. In exchanges similar to the previous example, we observed students’ confusion about which relation to use. Sometimes more than one applied. We felt that the ontologically mixed set of relation categories confused students about what they were trying to achieve with the diagrams, and did not help them focus on learning key distinctions. In order to encourage greater clarity, we decided to focus on evidential reasoning, and specifically on the most essential relational distinction for evidence based inquiry: whether two statements are consistent or inconsistent.

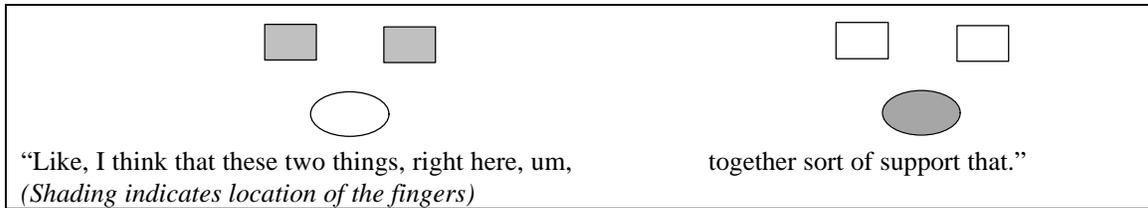


Figure 2. Gesturing to express a relationship between adjacent units.

3.4 Artifactual versus Fundamental Distinctions

Furthermore, we eliminated directionality from Belvedere's link representations of relations. At one time there were at least three versions of the "consistency" relation: "predicts," "explains" (drawn from hypotheses to data) and "supports" (drawn from data to hypotheses). Early versions of our automated advisor attempted to reason about and even enforce these semantics. However, we found that use of these links was inconsistent and sometimes differed from the intended semantics, consistent with other research on hypermedia link categories [34, 53]. The use of "predicts," "explains," and "supports" links was also misguided because these links are "surface" level discourse relations that do not encourage learners to think in terms of the more fundamental consistency relationships. Whether a hypothesis predicts or explains a datum is an artifact of the chronology of the datum with respect to statement of the hypothesis. Whether one uses "supports" (from datum to hypothesis) or one of the other two links (from hypothesis to datum) is an artifact of the focus of the discourse process by which the diagram is being constructed. Hence we eliminated these in favor of a single non-directional relation that expresses the more fundamental notion of evidential consistency.

3.5 Discussion Guided by Salience and Task

Consideration of ways in which subjects interacted with the representations led us to appreciate subtle ways in which external representations may guide discourse. For example, Figure 2 outlines a diagram state in which three statements were clustered near each other, with no links drawn between the statements. One student pointed to two statements simultaneously with two fingers of one hand, and drew them together as she gestured towards the third statement, saying "Like, I think that these two things, right here, um, together sort of support that" (Figure 2, from a videotape of an early laboratory study of Belvedere in which it was possible to make N-to-one links).

This event was originally taken merely as an example of how external representations facilitate the expression of complex ideas. However, this observation applies to any external representation. Reconsideration of this

example led to the hypotheses that several features of the representational system in use made the student's utterance more likely. First, elaboration on these particular statements is more likely because they (instead of others) are expressed as objects of perception in the representation. Second, this event is more likely to occur in a representational environment that provides a primitive for connecting statements with a support relation (particularly N-to-1) than in one that did not -- the students perceive their task as one of linking things together. Third, it may have been easier to recognize the relationship between the three statements because they happened to be spatially nearby each other. In this example, proximity was determined by the users rather than intrinsic to the representational toolkit. However, a representational tool could constrain proximity based on potential relationships between knowledge units (although Belvedere does not currently do this).

3.6 Representations in CSCL Systems for Collaborative Critical Inquiry

Observations such as these led the author to reexamine other software systems in use for the collaborative learning of "critical inquiry" and "scientific argumentation" skills, and identify the need for a series of systematic studies. Several major representational approaches to CSCL for critical inquiry are summarized below, in preparation for discussion of their implications for discourse. (The purpose of this discussion is to characterize major representational approaches in CSCL systems rather than to provide an adequate review of the systems themselves.)

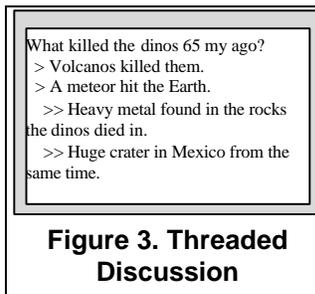


Figure 3. Threaded Discussion

Hypertext/hypermedia systems include CLARE [66]; CSILE [48, 49], the Collaboratory Notebook [14, 38], Web-Camile and Web-SMILE [17]. (Seminal systems include gIBIS [12] and NoteCards [18], which were not developed for educational applications.) These systems all have in common a hyperlinking of different comments relevant to an issue,

usually with categorization of the hyperlinks or their targets with labels such as “answer, argument, problem, solution, comment,” etc. There is wide variation in this category: some take the form of a threaded discussion or other tree structure that may be viewed in summary form (see Figure 3 for a characterization), while others support construction of graphs of “nodes” or “cards” through which one navigates, viewing one card at a time. For the purpose of illustration, threaded discussions (Figure 3) will stand for this class of representations. This choice was made to provide a simple reference point for comparison, and is not meant to detract from the richness of mature systems such as CSILE or its successor, Knowledge Forum, which uses several of the representational approaches discussed herein.

Argument mapping environments, a variation on concept mapping [36], include Belvedere [61, 62, 63], ConvinceMe [44], and Euclid [55]. All of these utilize node-link graphs representing rhetorical, logical, or evidential relationships between assertions (usually categorized as “hypothesis” versus “data” or “evidence”). As characterized in Figure 4, the entire graph is viewed and manipulated at once, distinguishing these systems from hypermedia environments in which one normally views and manipulates one node of the graph at a time.

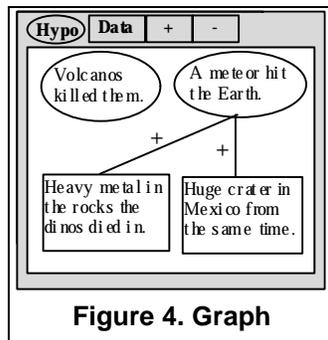


Figure 4. Graph

SenseMaker [2], a component of the KIE system [3], exemplifies an intermediate approach. Statements are organized in a 2-dimensional space and viewed all at once, as in argument graphs (see Figure 5). However, SenseMaker uses containment rather than links to represent the relationship of evidential support: an empirical statement is placed inside the box of the theory it supports. SenseMaker also uses containment to represent decomposition of a theory into hypotheses, a facility that was tried in early versions of Belvedere as well.

Finally, another representation is an evidence or criteria matrix (Figure 6). Such matrices organize hypotheses (or solutions) along one axis, and empirical

evidence (or criteria) along another, with matches between the two being expressed symbolically in the cells of the matrix. Puntambekar *et al.* [43] experimented with such a representation in a paper-based collaboration tool.

	+	-		
Data \ Hypo	Volcanic	Meteor		
Heavy metal in the rocks.	+	+		
Huge crater in Mexico.		+		

Figure 6. Matrix

Examining the figures above, the differences in representational formalisms provided by existing CSCL software for critical inquiry is striking. Yet more striking is the fact that there appear to be no systematic studies *comparing* the effects of external representations on collaborative learning discourse, although a number of valuable studies have been conducted on software utilizing single representational formalisms. (At a recent conference on computer supported collaborative learning, private communications with several designers of systems exemplified above corroborated this need. The author and other designers were not aware of such a study, and had all chosen designs based on informed intuition. Exceptions include [16, 69].) Given that these representations define the fundamental character of software intended to guide collaborative learning, a systematic comparison is overdue. The question is *not* “who’s system is better?” but rather “what kinds of interactions, and therefore learning, does each representational formalism encourage?” It may well be the case that all of the above representations are useful, albeit for different learning and problem solving phases or task domains.

3.7 External Representations in Individual and Collaborative Contexts

Substantial research has been and continues to be conducted concerning the role of external representations in individual problem solving, generally showing that the kind of external representation used to depict a problem may determine the ease with which the problem is solved [7, 30, 32, 35, 70]. Some representations unnecessarily restrict the solver’s search space (e.g., the 9-dots problem [19]). Alternatively, some forms of representation may have constraints built into them so that the representation actually enhances problem solving success [26, 57]. One might ask whether this research is sufficient to predict the effects of representations in collaborative learning. A related but distinct line of work undertaken in collaborative learning contexts is needed for several reasons. The interaction of the cognitive processes of several agents is different than the reasoning of a single

agent [37, 41, 47, 51], so may be affected by external representations in different ways. In particular, shared external representations can be used to coordinate distributed work, and will serve this function different ways according to their representational biases. Also, the mere presence of representations in a shared context with collaborating agents may change each individual's cognitive processes. One person can ignore discrepancies between thought and external representations, but an individual working in a group must constantly refer back to the shared external representation while coordinating activities with others. Thus it is conceivable that external representations have a greater effect on individual cognition in a social context than they do when working alone.⁵ Finally, much prior work on the role of external representations in individual problem solving have used well defined problems. Further study is needed on ill structured, open ended problems such as those typical of scientific inquiry.

4. Theoretical Discussion

The major hypothesis resulting from the foregoing work is that variation in features of representational tools used by learners working in small groups can have a significant effect on the learners' knowledge-building discourse and on learning outcomes. The claim is not merely that learners will talk about features of the software tool being used. Rather, with proper design of representational tools, this effect will be observable in terms of learners' talk about and use of *subject matter concepts and skills*. The author has begun investigations concerned with the conditions under which effects on learning interactions and outcomes are obtained: what features have what kind of effect? This section develops an initial theory of how representations guide learning interactions, and applies this analysis to make specific predictions concerning the effects of selected features of representational tools. The discussion begins with some definitions.

Representational tools are software interfaces in which users construct, examine, and manipulate external representations of their knowledge. The present work is concerned with symbolic as opposed to analogical representations. A formalism/artifact distinction [58] is critical to the present work: A representational tool is a software implementation of a **representational formalism** that provides a set of primitive elements out of which representations can be constructed. (For example, in Figure 4 the representational formalism is the collection of primitives for making hypothesis and data

statements and "+" and "-" links, along with rules for their use.) The software developer chooses the representational formalism and instantiates it as a representational tool, while the user of the tool constructs particular **representational artifacts** in the tool. (For example, in Figure 4 the representational artifact is the particular diagram of evidence for competing explanations of mass extinctions.)

Learning interactions include interactions between learners and the representations, between learners and other learners, and between learners and mentors such as teachers or pedagogical software agents. The present work focuses on interactions between learners and other learners, specifically verbal and gestural interactions termed **collaborative learning discourse**.

Each given representational formalism manifests a particular **representational bias**, expressing certain aspects of one's knowledge better than others [65]. The phrase **knowledge unit** will be used to refer generically to components of knowledge one might wish to represent, such as hypotheses, statements of fact, concepts, relationships, rules, etc. Representational bias manifests in two major ways:

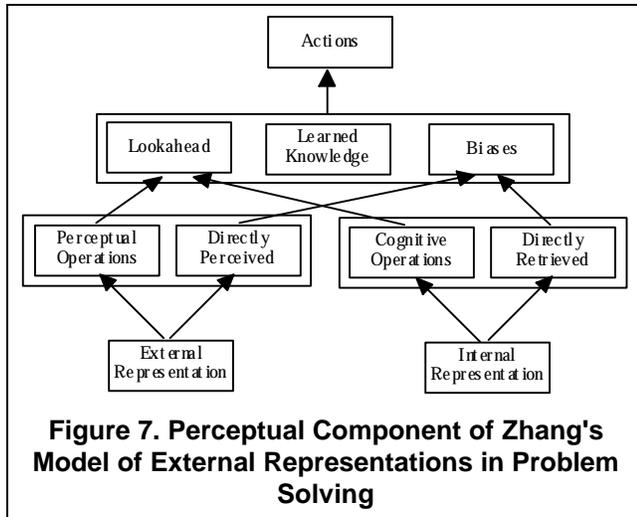
- ◆ **Constraints:** limits on expressiveness, and constraints on the sequence in which knowledge units can be expressed [45, 57].
- ◆ **Salience:** how the representation facilitates processing of certain knowledge units, possibly at the expense of others [32].

Representational tools mediate collaborative learning interactions by providing learners with the means to articulate emerging knowledge in a persistent medium, inspectable by all participants, where the knowledge then becomes part of the shared context. Representational bias *constrains* which knowledge can be expressed in the shared context, and makes some of that knowledge more *salient* and hence a likely topic of discussion.

4.1 A Perceptual/Logical Analysis

Zhang [70] distinguishes *cognitive* and *perceptual* operators in reasoning with representations (Figure 7). Cognitive operations operate on internal representations; while perceptual operations operate on external representations. According to Zhang, the latter perceptual operations take place without making an internal copy of the representation (although internal representations may change as a result of these operations). The author's theoretical outlook is highly sympathetic with Zhang's account. Expressed in terms of Zhang's framework, the present analysis is concerned primarily with perceptual operations on external representations rather than cognitive operations on internal representations. This is

⁵ Micki Chi, personal communication.



because the proposed work is concerned with how representations that reside in learners' perceptually shared context mediate collaborative learning interactions. While it is the case that cognitive operations on internal representations will influence interactions in the social realm, CSCL system builders do not design internal representations -- they design tools for constructing external representations. These external representations are accessed by perceptual operations, so it is the perceptual features of a representational formalism that are of greatest interest for CSCL systems.

Stenning and Oberlander [57] distinguish constraints inherent in the logical properties of a representational formalism from constraints arising from the architecture of the agent using the representational formalism. This corresponds roughly to the present author's distinction between "constraints" and "salience," if one considers primarily the perceptual architecture. "Constraints" are logical and semantic features of the representational formalism; while "salience" is not, being better understood in terms of Zhang's distinction between obtaining information by "direct perception" versus application of perceptual operators (Figure 7): information that is recoverable from a representation is salient to the extent to which it is recoverable by automatic perceptual processing rather than through a controlled sequence of perceptual operators. (The author does not hold the view that no computation is required for perception. "Direct perception" requires computation, albeit highly automatic and requiring no executive control. Recovery of certain information from a representation may require controlled application of multiple direct perceptions. For example, examining a graph, one's perception of the color of a node in a graph is more direct than one's perception of whether this node is connected by arcs to another specified node. Visual

search – a sequence of direct perceptions – is required to make the latter judgement. For our purposes, the important point is that the work required to retrieve information from a representation can vary as the representational system changes.)

The discussion now turns to the identification of dimensions along which different representational formalisms vary, and predictions that a given kind of learning interaction will increase along that same dimension.

4.2 Representational Formalisms Bias Learners Towards Particular Ontologies

The first major hypothesis claims that important guidance for learning interactions comes from ways in which a representational formalism *limits* what can be represented [45, 57]. A representational formalism provides a set of primitive elements out of which representational artifacts are constructed. These primitive elements constitute an "ontology" of categories and structures for organizing the task domain. Learners will see their task in part as one of making acceptable representational artifacts out of these primitives. Thus, they will search for possible new instances of the primitive elements, and hence (according to this hypothesis) will be biased to think about the task domain in terms of the underlying ontology. This point is illustrated by the examples from our Belvedere experience, given previously.

4.3 Salient Knowledge Units Receive More Elaboration

This hypothesis states that learners will be more likely to attend to, and hence elaborate on, the knowledge units that are perceptually salient in their shared representational workspace than those that are either not salient or for which a representational proxy has not been created. This is for reasons of *Reminding*: the visual presence of the knowledge unit in the shared representational context serves as a reminder of its existence and any work that may need to be done with it; and of *Ease of Reference*: it is easier to refer to a knowledge unit that has a visual manifestation, so learners will find it easier to express their subsequent thoughts about this unit than about those that require complex verbal descriptions [10]. These claims apply to any visually shared representations. However, to the extent that two representational formalisms differ in kinds of knowledge units they make salient, these functions of reminding and ease of reference will encourage elaboration on different kinds of knowledge units. The ability to manipulate learners' elaborations is important because substantial psychological research shows that

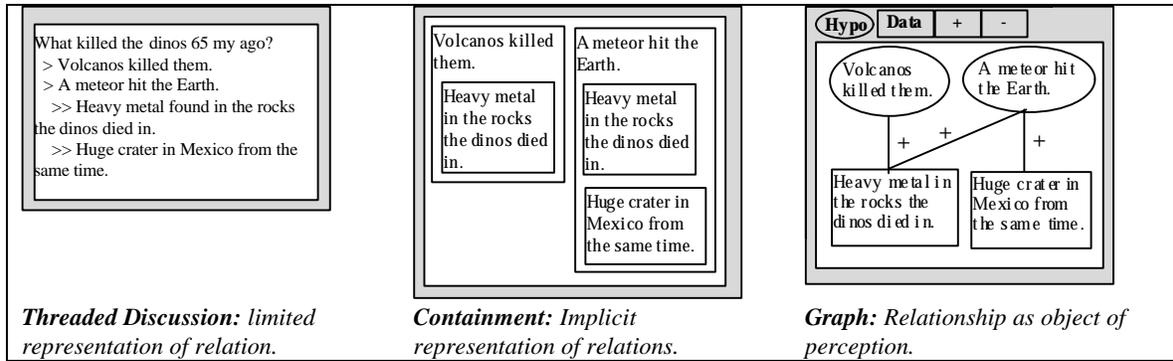


Figure 8. Example of Elaboration Hypothesis

elaboration leads to positive learning outcomes, including memory for the knowledge unit and understanding of its significance (e.g., [9, 13, 56]).

For example, consider the three representations of a relationship between four statements shown in Figure 8. The relationship is one of evidential support. The middle formalism uses containment to represent evidential support, while the right-hand formalism uses an arc. It becomes easier to perceive and refer to the *relationship* as an object in its own right as one moves from left to right in the figure. Hence the present hypothesis claims that relationships will receive more elaboration in the rightmost representational formalism. (The opposite prediction could be made in situations where learners see their task as one of putting knowledge units “in their place” in the representational environment. Once a unit is put in its place, learners may feel it can be safely ignored as they move on to other units not yet placed.)

4.4 Salience of Missing Knowledge Units Guides Search for New Knowledge

Some representational formalisms provide structures for organizing knowledge units, in addition to primitives for construction of individual knowledge units. *Unfilled “fields”* in these organizing structures, if perceptually

salient, can make *missing* knowledge units as salient as those that are present. If the representational formalism provides structures with predetermined fields that need to be filled with knowledge units, the present hypothesis predicts that learners will try to fill these fields. For example, a two dimensional matrix has cells that are intrinsic to the structure of the matrix: they are there whether or not they are filled with content. Learners using a matrix will look for knowledge units to fill the cells. The ability to manipulate learners’ awareness of missing knowledge could be a useful form of scaffolding for metacognitive competence [4].

For example, Figure 9 shows artifacts from three representational formalisms that differ in salience of missing evidential relationships. In the textual representation, no particular relationships are salient as missing: no particular prediction about search for new knowledge units can be made. In the graph representation, the lack of connectivity of the volcanic hypothesis to the rest of the graph is salient. However, once some connection is made to one data item, the hypothesis will appear connected, so one might predict that only one relationship involving each object will be sought. In the matrix representation, all undetermined relationships are salient as empty cells. The present

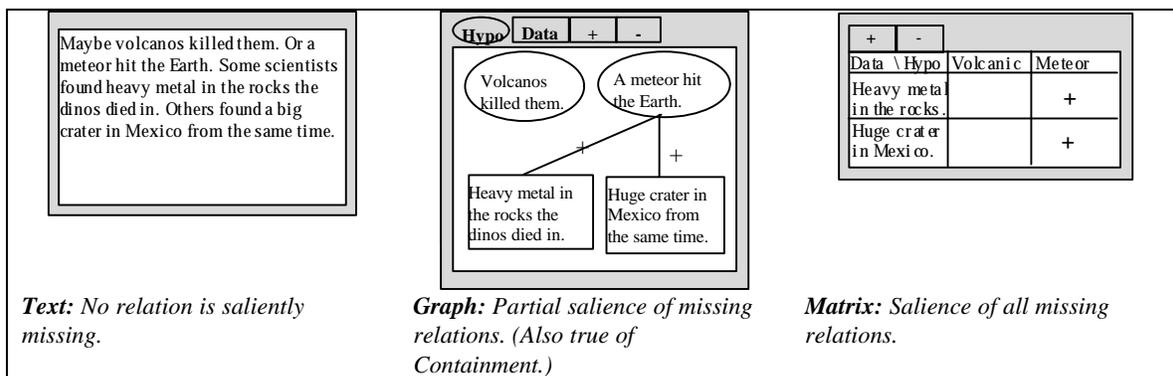


Figure 9. Example of Salient Absence Hypothesis

hypothesis predicts that learners will be more likely to discuss all possible relationships between objects when using matrices.

Other variations on representational formalisms with corresponding predictions are possible; however the discussion must end here for space considerations. The discussion now turns to studies now underway designed to test these predictions.

5. A Systematic Investigation of the Effects of Representational Bias on Discourse

The author has begun studies that test the effects of these formalisms on collaborative discourse and learning. At this writing, pilot studies have been run and are under analysis, and a proposal for in-depth study has been funded. Four representational formalisms are being compared in a proximal collaborative learning configuration. The four formalisms, which are characterizations of systems being deployed today, were chosen to maximize predicted differences along certain dimensions. These formalisms are **threaded discussions** (Figure 3), **graphs** (Figure 4), **containment** (Figure 5), and **matrices** (Figure 6). These formalisms intentionally differ on more than one feature (see Table 1). The research strategy is to maximize the opportunity to observe significantly different effects on learners' discourse and on learning outcomes. These results will then inform well-motivated selection of studies that vary

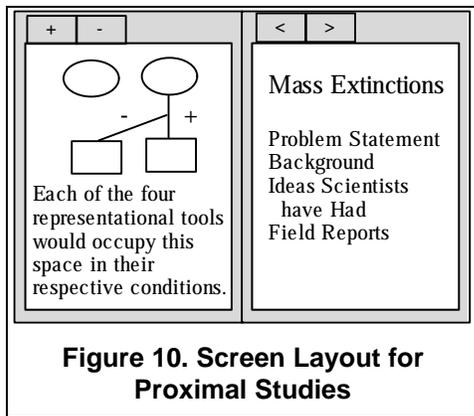
one feature at a time as needed to disambiguate alternate representational explanations for the results. This expedient approach is necessary in order to explore the large space of experimental comparisons within the time scale on which collaborative technology is being adapted.

Subjects are using a version of Belvedere modified for the experiments. They are presented with a "science challenge problem" in a web-browser. A science challenge problem presents a phenomenon to be explained, along with indices to relevant resources. It is important that these are relatively ill-structured problems: at any given point many possible knowledge units may reasonably be considered. This provides the necessary degrees of freedom within which representational bias can work.

The computer screen is divided in half as shown in Figure 10. The left hand side contains the representational tool -- any one of Threaded Discussion, Containment, Graph (shown), or Matrix. The right hand side contains a web browser open to the entry page for the science challenge materials. Students seated in front of the same monitor are asked to read the problem statement in the web browser on the right. They are then be asked to identify hypotheses that provide candidate explanations of the phenomenon posed, and evaluate these hypotheses on the basis of laboratory studies and field reports obtained through the hypertext interface. They are asked to use the representational tool on the left to record the information they find and how it bears on the problem.

	<i>Threaded</i>	<i>Containers</i>	<i>Graphs</i>	<i>Matrices</i>
Organization of Inquiry Activity	Discussion topics are posted, followed by chronologically organized replies.	Hypotheses are first recorded as boxes in the workspace. Empirical observations are sorted by placing them in boxes.	Hypotheses and empirical observations are recorded at any time as shapes placed in the workspace. Evidential relations are recorded by linking shapes together.	Hypotheses and empirical observations are recorded at any time by creating new columns & rows (respectively). Evidential relations are recorded by placing symbols in empty cells.
Ontology	Implicit: ◆ statements ◆ reply chronology	Implicit: ◆ hypothesis ◆ empirical observation ◆ consistency	Explicit: ◆ hypothesis ◆ empirical observation ◆ consistency ◆ inconsistency	Explicit: ◆ hypothesis ◆ empirical observation ◆ consistency ◆ inconsistency
Saliency of Known Relations	Implicit in Context: ◆ reply chronology	Implicit in Context: ◆ consistency	Explicit Object: ◆ consistency ◆ inconsistency	Explicit Object: ◆ consistency ◆ inconsistency
Saliency of Missing Relations	No saliency.	Lack of <i>some</i> consistency relation for a hypothesis.	Lack of <i>some</i> consistency or inconsistency relation for a statement.	Relations for <i>all combinations</i> of hypothesis and empirical observation.

Table 1. Features of Selected Representational Formalisms



Analysis will be based on transcripts of subjects' spoken discourse, gestures, and modifications to the interface; as well as measures of learning outcomes. Based on the features of representational formalisms shown in Table 1 and the discussion of section 4, the following predictions are made. The symbols ">" and ">>" indicate that the discourse phenomenon at the beginning of the list (concept use, elaboration, or search) will occur at a significantly greater rate in the treatment condition(s) on the left of the symbol than in those on the right. The double symbol ">>" indicates greater confidence in the prediction.

Concept Use (section 4.2): Graph, Matrix >>

Container, Threaded Discussion. The Graph and Matrix representations *require* that one categorize statements and relations. This will initiate discussion of the proper choice, possibly including peer coaching on the underlying concepts. The Container and Threaded Discussion representations provide only implicit categorization. There may be some talk of where to put things, but this talk is less likely to be expressed in terms of the underlying concepts.

Elaboration on Relations (section 4.3): Graph > Matrix

>> Container > Threaded Discussion. Graphs and Matrices make relations explicit as objects that can be pointed to and perceived, while this is not the case in the other two representations. More tenuously, the ability to link relations to other relations in Graphs may increase elaboration on relations relative to Matrices, if this facility is used. The appearance of one statement inside another's container constitutes a more specific assertion than contiguity of statements in a Threaded Discussion. Hence subjects are more likely to talk about whether a statement has been placed correctly in the Container representation.

Search for Missing Relations (section 0): Matrix >>
Graph, Container >> Threaded Discussion. In the Matrix representation, there is an empty field for *every* undetermined relationship, prompting subjects to consider all of them. In Graphs or the Container representations, salience of the lack of *some* relationship goes away as soon as a link is drawn to the statement in question or another is placed in its container, respectively. Threaded Discussion does not specifically direct search towards missing relationships.

Follow-up studies will test the generality of selected results in a distance learning environment. The method will be similar except that learners will be working in different rooms using a "chat" facility as a medium of discourse instead of speech.

6. Summary

Prior experience with Belvedere suggests that variation in features of the representational tools provided by such technology can have a significant effect on the learners' knowledge-building discourse and on learning outcomes. The paper sketched a theoretical analysis of the role of constraints and salience in representational bias, and outlined a systematic investigation being undertaken by the author. Work underway will inform the design of future software learning environments and provide a better theoretical understanding of the role of representational bias in guiding learning processes.

7. Acknowledgments

The author is grateful to Alan Lesgold, who initiated the Belvedere project, for his mentorship; Violetta Cavalli-Sforza for work on Belvedere 1.0's argumentation language; Eva Toth and Arlene Weiner for substantial contributions to Belvedere's curriculum and classroom implementation; Dan Jones for work on the client-server architecture and collaborative database; Kim Harrigal for work on the client; Massimo Paolucci and Joe Toth for work on the coach; John Connelly and Sandy Katz for evaluation of the coach; Cynthia Liefeld for assistance with ongoing pilot studies; and Micki Chi, Marth Crosby, and John Levine for discussions concerning the role of representations in learning, visual search, and social aspects of learning, respectively. Work on Belvedere and development of these ideas was funded by DoDEA's Presidential Technology Initiative and by DARPA's Computer Aided Education and Training Initiative. At this writing, the research has been funded by NSF's Learning and Intelligent Systems program.

8. References Cited

- [1] Abrami, P.C. & Bures, E.M. (1996). Computer-supported collaborative learning and distance education, Reviews of lead article. *The American Journal of Distance Education* 10(2): 37-42.
- [2] Bell, P. (1997). Using argument representations to make thinking visible for individuals and groups. In *Proceedings of the Computer Supported Collaborative Learning Conference '97*, pp. 10-19. University of Toronto, December 10-14, 1997.
- [3] Bell, P., Davis, E. A., & Linn, M. C. (1995). The knowledge integration environment: Theory and design. In *Proceedings of the Computer Supported Collaborative Learning Conference '95*, pp. 14-21. Mahwah, NJ: LEA.
- [4] Brown, A.L., Bransford, J. D., Ferrara, R.A., & Campione, J.C. (1983). Learning, remembering, and understanding. In *J.H.F.E.M. Markman (Ed.), Handbook of Child Psychology, Volume 3: Cognitive Development* pp. 77-166. New York: Wiley, 1983.
- [5] Brown, A. L. & Campione, J. C. (1994). Guided discovery in a community of learners. In *K. McGilly (Ed), Classroom Lessons: Integrating Cognitive Theory and Practice*. Cambridge: MIT Press, 1994. pp. 229-270.
- [6] Brown, A.L., & Palincsar, A.S. (1989). Guided, cooperative learning and individual knowledge acquisition. In L. Resnick (Ed.), *Knowing, Learning, and Instruction: Essays in Honor of Robert Glaser*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- [7] Chandrasekaran, B.I., Narayanan, N., & Iwasaki, Y. (1993). Reasoning with diagrammatic representations: A report on the spring symposium. *AI Magazine* 14(2): 49-56. 1993.
- [8] Chi, M. & Bassok, J. (1989). Learning from examples via self-explanations. In *L. Resnick (Ed.) Knowing, learning and instruction: Essays in Honor of Robert Glaser*, pp. 251-282. Hillsdale, NJ: Lawrence Erlbaum Associates.
- [9] Chi, M.T.H., Bassok, M., Lewis, M., Reimann, P., & Glaser, R. (1989). Self-explanations: How students study and use examples in learning to solve problems. *Cognitive Science*, 13:145-182, 1989.
- [10] Clark, H.H. & Brennan, S.E. (1991). Grounding in Communication. In L.B. Resnick, J.M. Levine and S.D. Teasley (eds.), *Perspectives on Socially Shared Cognition*, American Psychological Association, 1991, pp. 127-149.
- [11] Collins, A. & Ferguson, W. (1993). Epistemic Forms and Epistemic Games: Structures and Strategies to Guide Inquiry. *Educational Psychologist*, 28(1), 25-42.
- [12] Conklin, J. & Begeman, M.L. (1987). gIBIS: A hypertext tool for team design deliberation. In *Hypertext'97 Proceedings*, Chapel Hill, NC, pp 247-252. New York: ACM.
- [13] Craik, F. I. M., & Lockhart, R. S. (1972). Levels of processing: A framework for memory research. *Journal of Verbal Learning and Verbal Behavior*. 11: 671-684.
- [14] Edelson, D. & O'Neill, D. (1994). The CoVis Collaboratory Notebook: Computer support for scientific inquiry. Paper presented at the *Annual Meeting of the American Educational Research Association*, New Orleans.
- [15] Goldenberg, E. P. (1995). Multiple Representations: A Vehicle for Understanding Understanding. In *D. Perkins, J. Schwartz, M. West, & M. Wiske (Eds) Software Goes to School: Teaching for Understanding with New Technologies*, pp. 155-171. New York: Oxford University Press 1995.
- [16] Guzdial, M. (1997). Information ecology of collaborations in educational settings: Influence of tool. *Proc. 2nd Int. Conf. on Computer Supported Collaborative Learning (CSCL'97)*, Toronto, December 10-14, 1997. pp. 91-100.
- [17] Guzdial, M., Hmelo, C., Hubscher, R., Nagel, K., Newstetter, W., Puntambekar, S., Shabo, A., Turns, J., & Kolodner, J. L. (1997). Integrating and guiding collaboration: Lessons learned in Computer-Supported Collaborative Learning research at Georgia Tech. *Proc. 2nd Int. Conf. on Computer Supported Collaborative Learning (CSCL'97)*, Toronto, December 10-14, 1997. pp. 91-100.
- [18] Harp, B. & Neches, R. (1988). Notecards: An everyday tool for aiding in complex tasks. *Proc. Architectures for Intelligent Interfaces Workshop*, Monterey, California, Marcy 1988. ACM/SIGCHI.
- [19] Hayes, J. R. (1989). *The complete problem solver*. Hillsdale, NJ, Lawrence Erlbaum Associates.
- [20] Johnson, D. & Johnson, R. (1989). *Cooperation and Competition: Theory and Research*. Interaction Book Company, 1989.
- [21] Johnson, R.T., Johnson, D.W., & Stanne, M.B. (1985). Effects of cooperative, competitive, and individualistic goal structures on computer-assisted instruction. *Journal of Educational Psychology*, 77(6), 668-677.
- [22] Jonassen, D.H., Davidson, M., Collins, M., Campbell, J., & Bannan Haag, B. (1995). Constructivism and computer-mediated communication in distance education. *Journal of Distance Education* 9(2): 7-26.

- [23] Jonassen, D.H. & Reeves, T.C. (1996). Learning with technology: Using computers as cognitive tools. In D. H. Jonassen (Ed.), *Handbook of Research on Educational Communications and Technology*. pp. 693-719. New York: Simon and Schuster.
- [24] Justen, III, J.E., Waldrop, P.B., & Adams, II, (1990, July). Effects of paired versus individual user computer-assisted instruction and type of feedback on student achievement. *Educational Technology*, 30(7), 51-53.
- [25] Kaput, J. (1995). Creating Cybernetic and Psychological Ramps from the Concrete to the Abstract: Examples from Multiplicative Structures. In D. Perkins, J. Schwartz, M. West, & M. Wiske (Eds) *Software Goes to School: Teaching for Understanding with New Technologies*, pp. 130-154. New York: Oxford University Press 1995.
- [26] Klahr, D. and M. Robinson (1981). "Formal assessment of problem-solving and planning processing in school children." *Cognitive Psychology* 13: 113-148.
- [27] Koedinger, K. (1991). On the design of novel notations and actions to facilitate thinking and learning. *Proceedings of the International Conference on the Learning Sciences*, pp. 266-273. Charlottesville, VA: Association for the Advancement of Computing in Education. 1991.
- [28] Koedinger, K. R., Suthers, D. D., & Forbus, K. D. (1998). *Component-Based Construction of a Science Learning Space*. Paper presented at the Proceedings of Intelligent Tutoring Systems 98, San Antonio, TX.
- [29] Koschmann, T. D. (1994). Toward a theory of computer support for collaborative learning. *The Journal of the Learning Sciences*, 1994 3(3): 219-225.
- [30] Kotovsky, K. and H. A. Simon (1990). "What makes some problems really hard: Explorations in the problem space of difficulty." *Cognitive Psychology* 22: 143-183.
- [31] Lajoie, S. P., & Derry, S. J. (Eds.). (1993). *Computers as Cognitive Tools*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- [32] Larkin, J. H. & Simon, H. A. (1987). Why a diagram is (sometimes) worth ten thousand words. *Cognitive Science* 11(1): 65-99. 1987.
- [33] Lave, J. & Wenger, E. (1991). *Situated Learning: Legitimate Peripheral Participation*. Cambridge: Cambridge University Press.
- [34] Marshall, C. C., & Rogers, R. A. (1992). Two Years before the Mist: Experiences with Aquanet. Paper presented at the *Proceedings of the Fourth ACM Conference on Hypertext*, Milano, November 30-December 4, pp. 53-62.
- [35] McGuiness, C. (1986). "Problem representation: The effects of spatial arrays." *Memory & Cognition* 14(3): 270-280.
- [36] Novak, J. (1990). Concept mapping: A useful tool for science education. *Journal of Research in Science Teaching* 27(10): 937-49.
- [37] Okada, T. & Simon, H. A. (1997). Collaborative discovery in a scientific domain. *Cognitive Science* 21(2): 109-146.
- [38] O'Neill, D. K., & Gomez, L. M. (1994). The collaboratory notebook: A distributed knowledge-building environment for project-enhanced learning. In *Proceedings of Ed-Media '94*, Vancouver, BC.
- [39] Paolucci, M., Suthers, D., & Weiner, A. (1996). Automated advice-giving strategies for scientific inquiry. *Intelligent Tutoring Systems, 3rd International Conference*, Montreal, June 12-14, 1996.
- [40] Pea, R. (1994). Seeing what we build together: Distributed multimedia learning environments for transformative communications. *Journal of the Learning Sciences*, 3(3): 285-299.
- [41] Perkins, D.N. (1993). Person-plus: A distributed view of thinking and learning. In G. Salomon (Ed). *Distributed cognitions: Psychological and Educational Considerations* pp. 88-111. Cambridge: Cambridge University Press.
- [42] Perkins, D. N., Crismond, D, Simmons, R., & Unger, C. (1995). Inside Understanding. In D. Perkins, J. Schwartz, M. West, & M. Wiske (Eds) *Software Goes to School: Teaching for Understanding with New Technologies*, pp. 70-87. New York: Oxford University Press 1995.
- [43] Puntambekar, S., Nagel, K., Hübscher, R., Guzdial, M., & Kolodner, J. (1997). Intra-group and Intergroup: An exploration of Learning with Complementary Collaboration Tools. In *Proceedings of the Computer Supported Collaborative Learning Conference '97*, pp. 207-214. University of Toronto, December 10-14, 1997.
- [44] Ranney, M., Schank, P., & Diehl, C. (1995). Competence versus performance in critical reasoning: Reducing the gap by using Convince Me. *Psychology Teaching Review*, 1995, 4(2).
- [45] Reader, W. (Unpublished). Structuring Argument: The Role of Constraint in the Explication of Scientific Argument, manuscript dated November 1997.
- [46] Rysavy, D.M. & Sales, G.C. (1991). Cooperative learning in computer-based instruction. *Educational Technology Research & Development*, 39(2), 70-79.

- [47] Salomon, G. (1993). No distribution without individuals' cognition: A dynamic interactional view. In G. Salomon (Ed). *Distributed cognitions: Psychological and Educational Considerations* pp. 111-138. Cambridge: Cambridge University Press.
- [48] Scardamalia, M., & Bereiter, C. (1991). Higher levels of agency for children in knowledge building: A challenge for the design of new knowledge media. *The Journal of the Learning Sciences* 1(1), 37--68.
- [49] Scardamalia, M., Bereiter, C., Brett, C., Burtis, P.J., Calhoun, C., & Smith Lea, N. (1992). Educational applications of a networked communal database. *Interactive Learning Environments*, 2(1), 45-71.
- [50] Scardamalia, M., Bereiter, C., & Lamon, M. (1994). The CSILE project: Trying to bring the classroom into World 3. In K. McGilly (Ed), *Classroom Lessons: Integrating Cognitive Theory and Practice*. Cambridge: MIT Press, 1994. pp. 201-228.
- [51] Schoen, D. (1992). Designing as reflective conversation with the materials of a design situation. *Knowledge-Based Systems Journal - Special Issue on AI and Design*, Vol. 5, No. 1, 1992, pp. 3-14.
- [52] Sharan, S. (1980). Cooperative learning in small groups: Recent methods and effects on achievement, attitudes, and ethnic relations. *Review of Educational Research*, 50, 241-272
- [53] Shipman, F. M., & McCall, R. (1994). Supporting Knowledge-Base Evolution with Incremental Formalization Design Evaluation. *Proceedings of ACM CHI'94 Conference on Human Factors in Computing Systems*, Boston, April 24-28, pp. 285-291.
- [54] Slavin, R. E. (1990). *Cooperative learning: Theory, research, and practice*. Englewood Cliffs, NJ: Prentice-Hall.
- [55] Smolensky, P., Fox, B., King, R., & Lewis, C. (1987). Computer-aided reasoned discourse, or, how to argue with a computer. In R. Guindon (Ed.), *Cognitive science and its applications for human-computer interaction* (pp. 109-162). Hillsdale, NJ: Erlbaum.
- [56] Stein, B. S., & Bransford, J. D. (1979). Constraints on effective elaboration: Effects of precision and subject generation. *Journal of Verbal Learning and Verbal Behavior* 18:769-777.
- [57] Stenning, K. & Oberlander, J. (1995). A cognitive theory of graphical and linguistic reasoning: Logic and implementation. *Cognitive Science* 19(1): 97-140. 1995.
- [58] Stenning, K. & Yule, P. (1997). Image and language in human reasoning: A syllogistic illustration. *Cognitive Psychology* 34: 109-159.
- [59] Suthers, D. (1995). Designing for internal vs. external discourse in groupware for developing critical discussion skills. *CHI'95 Research Symposium*. Denver, May 1995.
- [60] Suthers, D. & Jones, D. (1997). An architecture for intelligent collaborative educational systems. *AI-Ed 97, the 8th World Conference on Artificial Intelligence in Education*, Kobe, August 20-22, 1997.
- [61] Suthers, D., Toth, E., and Weiner, A. (1997). An Integrated Approach to Implementing Collaborative Inquiry in the Classroom. *Proc. 2nd Int. Conf. on Computer Supported Collaborative Learning (CSCL'97)*, Toronto, December 10-14, 1997. pp. 272-279.
- [62] Suthers, D. and Weiner, A. (1995). Groupware for developing critical discussion skills. *CSCL '95, Computer Supported Cooperative Learning*, Bloomington, Indiana, October 17-20, 1995.
- [63] Suthers, D., Weiner, A., Connelly, A. and Paolucci, M. (1995). Belvedere: Engaging students in critical discussion of science and public policy issues. *AI-Ed 95, the 7th World Conference on Artificial Intelligence in Education*, August 16-19, 1995, Washington DC
- [64] Toth, J., Suthers, D., and Weiner, A. (1997). Providing expert advice in the domain of collaborative scientific inquiry. *8th World Conference on Artificial Intelligence in Education (AIED'97)*, Kobe, August, 1997.
- [65] Utgoff, P. (1986). Shift of bias for inductive concept learning. In R. Michalski, J. Carbonell, T. Mitchell (Eds.) *Machine Learning: An Artificial Intelligence Approach, Volume II*, Los Altos: Morgan Kaufmann 1986, pp. 107-148.
- [66] Wan, D., & Johnson, P. M. (1994). Experiences with CLARE: a Computer-Supported Collaborative Learning Environment. *International Journal of Human-Computer Studies*, October, 1994.
- [67] Webb, N. (1989). Peer interaction and learning in small groups. *International Journal of Education Research*, 13:21-40, 1989.
- [68] Webb, N. & Palincsar, A. (1996). Group processes in the classroom. In *Handbook of Educational Psychology*, D. Berliner & R. Calfee, Eds. Simon & Schuster Macmillian, New York 1996.
- [69] Wojahn, P. G., Neuwirth, C. M., & Bullock, B. (1998). Effects of interfaces for annotation on communication in a collaborative task. *Conference on Human Factors in Computing Systems (CHI 98)*, 18-23 April, Los Angeles, pp. 456-463.
- [70] Zhang, J. (1997). The nature of external representations in problem solving. *Cognitive Science*, 21(2): 179-217, 1997.