
Perspectives on Collaborative Learning

After the exploration of computer support for personal and small-group perspectives described in chapters 4 and 5, I tried to push as hard as I could a model of threaded discussion with perspectives. I developed a Web-based tool called WebGuide, designed to mediate and structure collaborative learning. This software defined a flexible system of perspectives on a shared knowledge construction space. WebGuide provides an electronic and persistent workspace for individuals and teams to develop and share distinctive points of view on a topic. The software and associated usage practices were designed by trials in a middle school classroom and in an advanced graduate seminar. Experience in these use-situations raised a range of questions concerning theoretical and practical issues, which drove further research. This chapter is a reflection on what was collaboratively learned about how software artifacts can mediate learning and shared cognition.

This chapter's multi-faceted discussion of the design of the WebGuide collaboration system reflects on the intricate relationship between theory, design and usage evaluation. It demonstrates the interplay of considerations required in designing support for the intertwining of personal and group perspectives, arguably that aspect of collaboration most in need of computational support. This design study reflects on the emergence of abstract theory from practical implementation issues.

WebGuide was probably my most intensive software development effort. I tried to create a system that would support group cognition in the sense of collaborative knowledge building. Threaded discussion seemed to be an appropriate medium, but its use always tended to be limited to the exchange of personal opinions, at best. I developed a computational system to support the sharing of interpretive perspectives, but it turned out to be too complicated to use fluidly, despite repeated attempts to make its interface more intuitive. Theoretical reflections related to WebGuide led me to bring in communication specialists and to undertake the analyses of part II of this book and the reflections of part III.

The main section of this chapter was written for the April 1999 AERA Conference. A year later, it was peer-reviewed in the online Journal of Interactive Media in Education; materials from the interaction with the reviewers have been appended to the paper.

1. Introductory Narrative

For some years now I have been interested in how to personalize the delivery of information from knowledge repositories to people based on their preferred *perspectives* on the information (Stahl, 1995, 1996). For instance, designers often critique an evolving design artifact from alternative technical points of view; different designers have

different personal concerns and styles, requiring considerations based upon access to different rules of thumb, rationale, constraints, standards and other forms of domain knowledge. Computer design environments should support these important interpretive perspectives. I am now primarily interested in applying similar mechanisms of perspectival computer support within contexts of collaborative learning.

In 1997, Ted Habermann—an information architect at the National Oceanic and Atmospheric Administration (NOAA) who makes geophysical data available to school children over the Web—suggested to me that we try to develop some computer support for a project at his son’s middle school. Dan Kowal, the environmental sciences teacher at the Logan School for Creative Learning in Denver, was planning a year-long investigation of alternative perspectives on the issue of “acid mine drainage” (AMD)—the pollution of drinking water supplies by heavy metals washed out of old gold mines. The fact that Dan and I were interested in “perspectives” from different perspectives seemed to provide a basis for fruitful collaboration. Ted obtained NSF funding for the project and we all spent the summer of 1998 planning the course and its perspectives-based software. Each of us brought in colleagues and worked to create a Java application (WebGuide), a set of auxiliary web pages and to put together a group of adult mentors representing different perspectives on AMD and a course curriculum.

The class started in September and the software was deployed in October. The students in Dan’s class were aware of the experimental nature of the software they were using and were encouraged to critique it and enter their ideas into WebGuide. Feedback from these twelve-year-old students provided initial experience with the usability of WebGuide and resulted in a re-implementation of the interface and optimization of the algorithms over the school’s Christmas vacation.

In January 1999, I organized an interdisciplinary seminar of doctoral students from cognitive, educational and computational sciences to study theoretical texts that might provide insight into how to support collaborative learning with perspectives-based software. The seminar used WebGuide as a major medium for communication and reflection, including reflection on our use of the software. This provided a second source of experience and raised a number of issues that needed to be addressed in software redesign.

In this chapter I would like to begin a reflection on the issues that have arisen through our WebGuide experiences because I think they are critical to the ability to support collaborative learning with computer-based environments. The potential for computer mediation of collaboration seems extraordinary, but our experience warns us that the practical barriers are also enormous. Certainly, our experiences are not unique, and similar projects at the universities of Toronto, Michigan, Berkeley, Northwestern, Vanderbilt, Georgia Tech, etc. have run into significant obstacles for years. Indeed, we observed many of these issues in a seminar in the year prior to the implementation of WebGuide (dePaula, 1998; Koschmann & Stahl, 1998). However, I believe that perspectives-based software addresses or transforms some of the issues and raises some of its own.

Let me describe how computer support for perspectives has evolved in WebGuide. I will first discuss the preliminary implementation as used in Dan’s middle school environmental course and explain how perspectives are supported in that version. A number of design issues led to an extended attempt to bring theory to the aid of reflection

on practice. This included the graduate seminar that used a revised version of WebGuide. Finally, following the original part of this chapter is a condensed version of the dialog that took place between the *Journal of Interactive Media in Education (JIME)* reviewers and me, where responses from winter 2000 and spring 2001 bring in reflections from subsequent design iterations.

2. Practice I: Environmental Perspectives

An early implementation of WebGuide was in use in Dan's classroom at the Logan School. For the previous five years, his class had researched the environmental damage done to mountain streams by the Gamble Gulch mine site. Then they investigated the social issue of acid mine drainage from various perspectives: environmental, governmental, mine-owner and local landowner. Working in teams corresponding to each of these perspectives, they articulated the position of their perspective on a set of shared

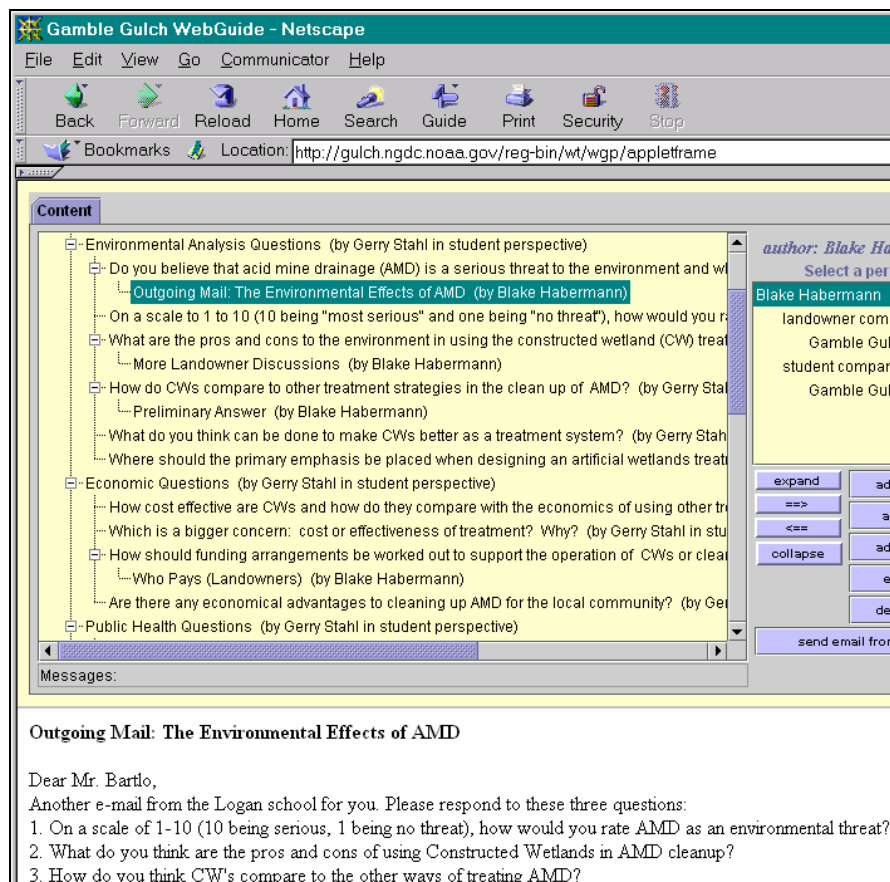


Figure 6-1. The Gamble Gulch version of *WebGuide* viewed in a Web browser. The top part is a Java applet displaying an outline view of note titles. The content of the selected note is displayed in an HTML frame below. To the right are buttons for navigating the outline and changing the content in the shared knowledge space. The view shown is from the personal perspective of one student.

questions.

The “Gamble Gulch” application of WebGuide served as the medium through which the students collaboratively researched these issues with their mentors and with teammates. Each student and mentor had their personal display perspective, and their display perspectives each inherited from one of the content-based team perspectives (environmental protection, governmental regulation, etc.), depending upon which intellectual perspective they were working on constructing.

Figure 6-1 shows one student’s (Blake) personal perspective on the class discourse. The tree of discussion threads was “seeded” with question categories, such as “Environmental Analysis Questions.” Within these categories, the teacher and I posted specific questions for the students to explore, like, “Do you believe that AMD is a serious threat to the environment?” Here, Blake has sent an email to a mentor asking for information related to this question. Email interactions happen through WebGuide and are retained as notes in its display perspectives. When replies are sent back, they are automatically posted to the discussion outline under the original email. When someone clicks on a title, the contents of that note are displayed in an HTML frame below the applet (as is the body of the student’s email in figure 6-1).

Figure 6-1 goes approximately here

Blake is working in his personal perspective, which inherits from the Class, Student team and Landowner team perspectives (see the dashed red arrows in figure 6-2). Note that the display of his personal perspective (in figure 6-1) includes notes that Dan and I entered in the Student perspective to structure the work of all the students. Blake can add, edit and delete ideas in his perspective, as well as sending email in it. Because he is a member of the landowner team and the student group as well as the class, he can browse

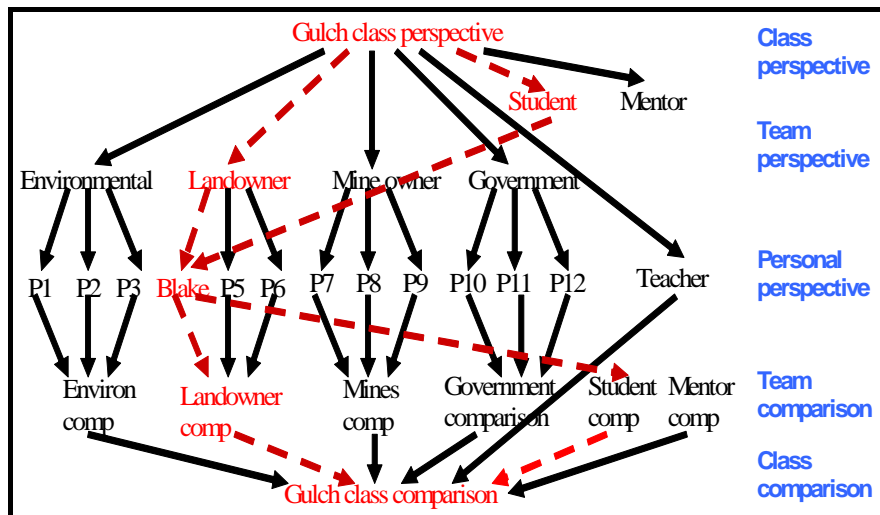


Figure 6-2. The web of perspectives in Gamble Gulch. Information is automatically inherited downward in the diagram. Blake’s perspective includes all the notes entered in the Gulch class, Landowner and Student perspectives. His notes also show up in the Landowner, Student and Gulch class comparison perspectives.

ideas in the Student comparison, the Landowner comparison and the Gamble Gulch class comparison perspectives (see list of perspectives accessible to him on the right of figure 6-1).

Figure 6-2 goes approximately here

For this application, the teacher has decided that perspective comparing and negotiation will take place in live classroom discussions, rather than in WebGuide. After a team or the whole class reaches a consensus, the teacher will enter the statements that they have agreed upon into the team or class perspective.

The goal of the year-long course is not only to negotiate within teams to construct the various positions, but also to negotiate among the positions to reach consensus or to clarify differences. Dan designed this class—with its use of WebGuide—to teach students that knowledge is perspectival, that different people construct different views, and that compilations of facts and arguments differ depending upon the social situation from which they arise. He hopes that his students will not only learn to evaluate statements as deriving from different perspectives, but also learn to negotiate the intertwining of perspectives to the extent that this is possible.

3. Computer Support of Perspectives

The term “perspectives” is over-loaded with meanings; this frequently produces confusion even when it is intended to tacitly exploit in one domain aspects of the perspectives metaphor from a different domain. It may be helpful at this point to distinguish three types of perspectives: literal, figurative and computational.

- *Literal perspectives* are optical or perceptual orientations: one sees objects from the specific angle or vantage point of the physical location of one’s eyes.
- *Figurative perspectives* take metaphorical license and refer to, for instance, different ways of conceptualizing a theme, as in adopting a skeptical view of a conversational claim.
- *Computational perspectives* are the result of software mechanisms that classify elements in a database for selective display. In WebGuide, for example, if I enter a note in my personal perspective then that note will be displayed whenever my perspective is displayed but not when someone else’s personal perspective is displayed.

WebGuide implements a system of computational (i.e., computer-supported, automated) perspectives designed to utilize the perspective metaphor in order to support characteristics of collaboration and collaborative learning. It is unique in a number of ways that distinguish it from other software systems that may use the term “perspectives”:

- Other systems refer to different *representations* of information as perspectives. They might have a graphical and a textual view of the *same* data. In WebGuide, different data is literally displayed in different perspectives while using the same representation—hierarchically structured titles of textual notes.

- In WebGuide, the perspectives mechanism is neither a simple tagging of data nor a database view, but is a dynamic computation that takes into account a *web of inheritance* among perspectives. Thus, Blake's perspective includes not only information that he entered in his perspective, but also information inherited from the Class, Student and Landowner perspectives.
- The web of perspectives can be extended by users *interactively*, and the inheritance of information is always computed on-the-fly, based on the current configuration of this web.
- The information in a perspective has a user-maintained structure in which each note has one or more "parent" notes and may have "child" notes, creating a *web of notes* within each perspective. The order of child notes displayed under a parent note is user defined and maintained so that WebGuide can be used to *organize ideas* within outline structures.

The computational perspectives mechanism we have been exploring incorporates the following features:

- Individual community members have access to their own information source. This is called their *personal perspective*. It consists of notes from a shared central information repository that are tagged for display within that particular perspective (or in any perspective inherited from that perspective).
- Notes can be created, edited, rearranged, linked together or deleted by users within their own personal perspective without affecting the work of others.
- Another student, Annie, can integrate a note from Blake's perspective into her own personal perspective by creating a *link* or *virtual copy* of the note. If Blake modifies the original note, then it changes in Annie's perspective as well. However, if Annie modifies the note, a new note is actually created for her, so that Blake's perspective is not changed. This arrangement generally makes sense because Annie wants to view (or inherit) Blake's note, even if it evolves. However, Blake should not be affected by the actions of someone who copied one of his notes.
- Alternatively, Annie can *physically copy* the contents of a note from Blake's perspective. In this case, the copies are not linked to each other in any way. Since Annie and Blake are viewing physically distinct notes now, either can make changes without affecting the other's perspective.
- There is an inheritance web of perspectives; descendent perspectives inherit the contents of their ancestor perspectives. Changes (additions, edits, deletions) in the ancestor are seen in descendent perspectives, but not vice versa. New perspectives can be created by users. Perspectives can inherit from existing perspectives. Thus, a team comparison perspective can be created that inherits and displays the contents of the perspectives of the team members. A hierarchy of team, sub-team, personal and comparison perspectives can be built to match the needs of a particular community (figure 2, above).

This model of computational perspectives has the important advantage of letting team members inherit the content of their team's perspective and other information sources without having to generate it from scratch. They can then experiment with this content on their own without worrying about affecting what others see.

4. Types of Perspectives

WebGuide provides several levels of perspectives (see figure 2) within a web of perspective inheritance to help students compile their individual and joint research:

- The *class perspective* is created by the teacher to start each team off with an initial structure and some suggested topics. It typically establishes a framework for classroom activities and defines a space used to instantiate the goal of collecting the products of collaborative intellectual work.
- The *team perspective* contains notes that have been accepted by a team. This perspective can be pivotal; it gradually collects the products of the team effort.
- The student's *personal perspective* is an individual's work space. It inherits a view of everything in the student's team's perspective. Thus, it displays the owner's own work within the context of notes proposed or negotiated by the team and class—as modified by the student. Students can each modify (add, edit, delete, rearrange, link) their virtual copies of team notes in their personal perspectives. They can also create completely new material there. This computational perspective provides a personal workspace in which a student can construct his or her own figurative perspective on shared knowledge. Other people can view the student's personal perspective, but they cannot modify it.
- The *comparison perspective* combines all the personal perspectives of team members and the team perspective, so that anyone can compare all the work that is going on in the team. It inherits from personal perspectives and, indirectly, from the team and class perspectives. Students can go here to get ideas and copy notes into their own personal perspective or propose items for the team perspective.

Of course, there is not really a duplication of information in the community memory. The perspectives mechanism merely displays the information differently in the different perspectival views, in accordance with the relations of inheritance.

5. Issues for Perspectives

The first issues to hit home when we deployed WebGuide were the problems of response time and screen real estate. The student computers were slower, had smaller monitors, lacked good Internet connections and were further from the server than the computers of the developers. We were, of course, already familiar with these issues from other Web applications, but one never knows quite how things will work out and how they will be accepted until one tests them under classroom conditions.

A pre-release prototype of WebGuide used dynamic HTML pages. This meant that each time a student expanded a different part of the outline of titles it was necessary to wait for a new page to be sent across the Internet. The dynamic HTML pages also greatly constrained the interface functionality. However, when we moved to a Java applet, we had to wait several minutes to download the applet code to each student computer. Furthermore, it entailed running all the perspectives computations on the slow student computers. In order to reduce the download time significantly, we first rewrote the interface using standard Java Swing classes that can be stored on the student machines. Then we split the applet into a client (the interface) and a server (the perspectives

computations and database access). By downloading only the client part to the classroom, we not only reduced the download time further, but also ran the time-consuming computations on our faster server computers.

Such technical problems can be solved relatively easily by optimizing algorithms or by adjusting tradeoffs based on local conditions. Issues of social practice are much more intransigent. There seem to be two major problems for software for threaded discussions and collaborative knowledge construction like WebGuide:

1. Lack of convergence among the ideas developed in the supported discussions.
2. Avoidance of system use in favor of email, face-to-face conversation or inaction.

WebGuide introduces its computational perspectives mechanism as a structural feature to facilitate the articulation of convergent ideas, and it even incorporates email. In attempting to address the problems posed above, it raises a new set of issues:

3. Is the perspectives metaphor a natural one (or can it be made natural) so that people will use computational perspectives to construct their figurative perspectives?
4. Can the web of perspectives be represented in a convenient and understandable format?

In our trials of WebGuide we have tried to create learning situations that would encourage the use of the software, yet we have observed low levels of usage and under-utilization of the system's full functionality. This raises the following additional issues:

5. How can learning situations be structured to take better advantage of the presumed advantages of the software?
6. How can the system's various capabilities be distinguished, such as its support for threaded discussions and for perspective-making?

In order to answer questions of this magnitude it was necessary to gather more experience, to be more closely involved in the daily usage of the system and to develop a deeper theoretical understanding of collaborative learning and of computer mediation. Having defined these goals, I announced a seminar on the topic of "computer mediation of collaborative learning," open to interested researchers from a number of disciplines—primarily education, cognitive psychology and computer science. The goal of the seminar was explicitly stated to be an experiment in the use of WebGuide to construct knowledge collaboratively, based on careful reading of selected texts. The texts traced the notion of *computer mediation* (Boland & Tenkasi, 1995; Caron, 1998; Hewitt, Scardamalia, & Webb, 1998; Scardamalia & Bereiter, 1996) back to *situated learning theory* (Bruner, 1990; Cole, 1996; Lave, 1991; Lave & Wenger, 1991; Lave, 1996)—and from there back to the notion of *mediated consciousness* in Vygotsky (1930/1978) and its roots in Hegel (Habermas, 1971; Hegel, 1807/1967; Koyeve, 1947/1969) and Marx (1844/1967; 1845/1967; 1867/1976).

In section 8 of this chapter I will comment on our current understanding of the six issues listed above. But first it is necessary to describe the ways in which the seminar attempts to make use of WebGuide and the conceptualization of the theory of computer mediation that is arising in the seminar.

6. Practice II: Theoretical Perspectives

The seminar on computer mediation of collaborative learning is designed to use WebGuide in several ways:

- *As the primary communication medium for internal collaboration.* The seminar takes place largely on-line. Limited class time is used for people to get to know each other, to motivate the readings, to introduce themes that will be followed up on-line, and to discuss how to use WebGuide within the seminar.
- *As an example collaboration support system to analyze.* Highly theoretical readings on mediation and collaboration are made more concrete by discussing them in terms of what they mean in a system like WebGuide. The advantage of using a locally-developed prototype like WebGuide as our example is that we not only know how it works in detail, but we can modify its functionality or appearance to try out suggestions that arise in the seminar.
- *As an electronic workspace for members to construct their individual and shared ideas.* Ideas entered into WebGuide persist there, where they can be revisited and annotated at any time. Ideas that arise early in the seminar will still be available in full detail later so that they can be related to new readings and insights. The record of discussions over a semester or a year will document how perspectives developed and interacted.
- *As a glossary and reference library.* This application of WebGuide is seeded with a list of terms that are likely to prove important to the seminar and with the titles of seminar readings. Seminar members can develop their own definitions of these terms, modifying them based on successive readings in which the terms recur in different contexts and based on definitions offered by other members. Similarly, the different readings are discussed extensively within WebGuide. This includes people giving their summaries of important points and asking for help interpreting obscure passages. People can comment on each other's entries and also revise their own. Of course, new terms and references can easily be added by anyone.
- *As a brainstorming arena for papers.* The application has already been seeded with themes that might make interesting research papers drawing on seminar readings and goals. WebGuide allows people to link notes to these themes from anywhere in the information environment and to organize notes under the themes. Thus, both individuals and groups can use this to compile, structure and refine ideas that may grow into publishable papers. Collaborative writing is a notoriously difficult process that generally ends up being dominated by one participant's perspective or being divided up into loosely connected sections, each representing somewhat different perspectives. WebGuide may facilitate a more truly collaborative approach to organizing ideas on a coherent theme.
- *As a bug report mechanism or feature request facility.* Seminar participants can communicate problems they find in the software as well as propose ideas they have for new features. By having these reports and proposals shared within the WebGuide medium, they are communicated to other seminar participants, who can then be aware of the bugs (and their fixes) and can join the discussion of suggestions.

The seminar version of WebGuide incorporates a built-in permissions system that structures the social practices surrounding the use of the system. Seminar participants each have their own personal perspective in which they can manipulate notes however they like without affecting the views in other perspectives. They can add quick discussion notes or other kinds of statements. They can edit or delete anything within their personal perspective. They can also make multiple copies or links (virtual copies) from notes in their personal perspective to other notes there. Anyone is free to browse in any perspective. However, if one is not in one's own perspective then one cannot add, edit or delete notes there (as in figure 6-3). To manipulate notes freely, one must first copy or link the note into one's own personal perspective. The copy or link can optionally include copying (or virtual copying) all the notes below the selected note in the tree as well. These rules are enforced by the user interface, which checks whether or not someone is in their personal perspective and only allows the legal actions.

Figure 6-3 goes approximately here

Students in the class can form sub-groups either within or across their different disciplines. They develop ideas in their personal perspectives. They debate the ideas of other people by finding notes of interest in the class comparison perspective (or in a

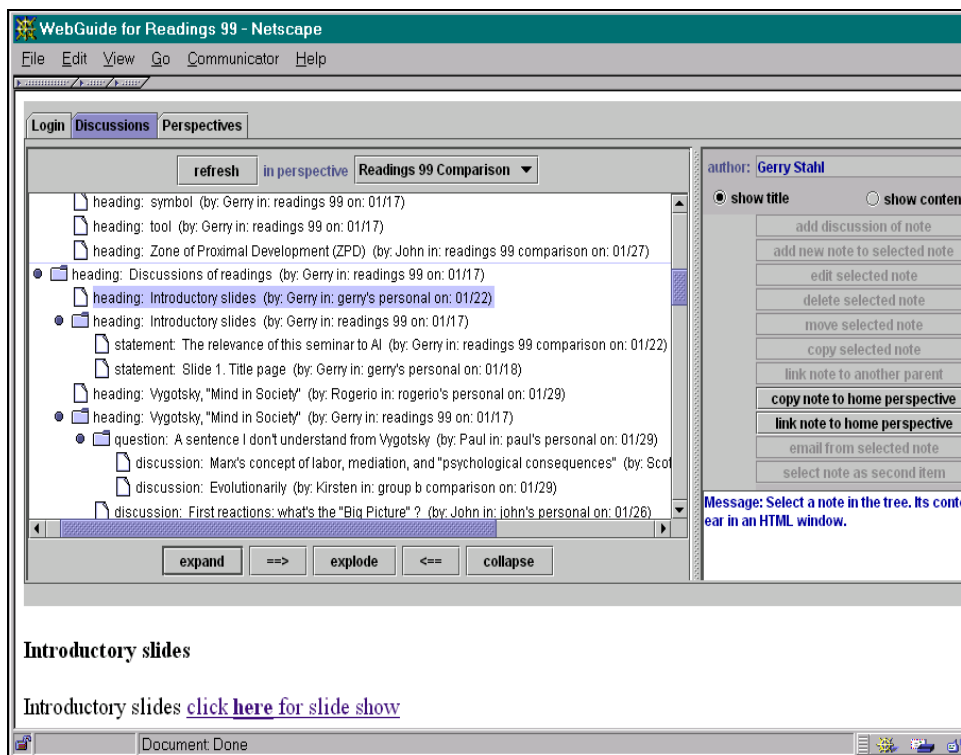


Figure 6-3. The version of *WebGuide* used in the seminar. Note that some of the control buttons on the right are not functional when the logged-in author is not working in his own personal perspective. This enforces certain social practices. Also note that many headings have been inserted to structure the discussion space.

subgroup comparison perspective) and copying these notes into their own personal perspective, where they can comment on them. The clash of perspectives is visible in the comparison perspectives, while the personal perspectives allow for complete expression and organization of a single perspective. This supports the “taking” of other people’s perspectives and the use of shared ideas in the “making” of one’s own perspectives (Boland & Tenkasi, 1995).

The seminar application of WebGuide stresses the use of perspectives for structuring collaborative efforts to build shared knowledge. The goal of the seminar is to evolve theoretical views on computer mediation—and to do so within a medium that supports the sharing of tentative positions and documents the development of ideas and collaboration over time. A major hypothesis investigated by the seminar is that software environments with perspectives—like WebGuide—can provide powerful tools for coordinated intellectual work and collaborative learning. It explores how the use of a shared persistent knowledge construction space can support more complex discussions than ephemeral face-to-face conversation. Many of the desires and concerns in this chapter originated as seminar notes in WebGuide. In particular, the seminar’s focus on theory has actually problematized our understanding of the role of theory.

7. Theory in Practice

Our initial application of WebGuide in the middle school environmental course raised a number of issues that led us to seek theoretical understanding through a seminar, which is serving as a second application of WebGuide. We have begun to see our research differently as a result of the theories we are incorporating into our discussions within the seminar. One thing that has changed is the relation we see of this theory to our research practice.

In my paper proposal to The American Educational Research Association (AERA)—the first draft of this chapter—written prior to our recent explorations, I described our approach by following the narrative order implied by conventional wisdom about the relation of theory to practice. After stating the goal or purpose of the work, I provided a theoretical framework, followed by sections on techniques, evidence, conclusions and educational/scientific import. The assumption here was that when one had a problem one turned first to theory for the solution and then “applied” the theory to some situation—either the problem situation or an experimental test context. After designing the solution based on the pre-existing theory and applying it to the test situation, one gathered evaluative data and analyzed it to measure success. The evaluation then implies whether or not the solution has a general import.

But such an approach is in keeping neither with our current experience nor with our emerging theory. We started last summer with an opportunity to explore some vague notions we had about something we called “perspectives.” We experimented with ever-evolving techniques through a complex collaborative process involving many people, each with their own concerns, understanding and insights. As part of this process some of us turned to theory—but the selection of theoretical texts and our interpretations of them were determined by the processes and issues we observed in our practical strivings.

In this draft of the chapter—still not considered a static final document, but a recapitulation from one particular moment in an on-going process—I am trying to narrate a story about how theory and practice have been co-mingled in our research. We began with an idea for a concrete classroom curriculum and worked on designing tools and structures to support the practical needs of that curriculum. Once we had a working software prototype that could be used over the Web, we deployed it in the middle school classroom. We immediately confronted the problems of response speed and monitor screen real estate that we had been worried about from the start. Students started asking for new functionality and it became clear that they were not using the implemented functions the way they were designed to be used. A dance commenced between the technicians, the educators, the students, the curriculum and the software; as we circled each other, we all changed and became more compatible.

There was no point in trying to evaluate the success of our experiment by gathering data under controlled conditions. It was clear that we needed to figure out how to make things work better, not to measure precisely how well they were (or were not) already working. Beyond the relatively clear technical usability issues there were deeper questions of how software can mediate interpersonal and cognitive relations within collaboration (Hewitt *et al.*, 1998). This led us to look for a theory of computer mediation—and for that matter a theory of collaborative learning—in the graduate seminar. Of course, it turned out that there are no adequate theories on these topics sitting on the bookshelf for us to simply apply. Rather, we had to undertake the construction of such theory, building upon hints strewn about in texts from many disciplines and guided by the problematic, in which we are involved first hand.

Trusting in our intuition that software like WebGuide could facilitate group theory building, we set out to use WebGuide in our theoretical investigations, and thereby further drive the development of the software through additional practical experience even as we were developing theoretical justifications for our design. In reflecting on our experience, I have tried to organize this draft of the chapter in accordance with a non-traditional theory about the relation of theory and practice—an understanding of this relationship more in keeping not only with our practice but with our hermeneutic, dialectical, socially situated activity theory.

Thus, we started out from our vague, only partially articulated background understanding of perspectives as an interesting and promising concept for learning and for computer support (see chapter 4). We set up a real-world situation in which we could explore what happens when this idea is implemented. In this situation we nurtured a process of “structural coupling” (Maturana & Varela, 1987) in which the different actors evolve toward a workable synthesis or homeostasis. Rapid prototyping cycles and participatory design sessions help facilitate this process. As breakdowns in the intention of our design are recognized, we engage in reflection-in-action (Schön, 1983) to make our tacit pre-understanding explicit, to understand what has happened and to project corrective actions. This process of explication raises broad issues and calls for theory. But despite the generality of the issues, the theory is not understood in a completely abstract way, but in terms of its relevance to our situation and to the specific barriers we have uncovered in that concrete situation.

Theory—like everyday thought—often arises after the fact (or well into the complex process of practical investigations) in order to justify situations that would otherwise be

too messy to comprehend and remember. Then, at the first chance it gets, theory reverses the order of things and presents itself as a guiding *a priori* principle. As Hegel (1807/1967) said, “the owl of Minerva flies only at night”: the wisdom of theory arrives on the scene only after the practical events of the day (which theory retroactively captures in concepts) have been put to bed. *Theory* is a cherished way to capture an understanding of what has been learned, even if it distorts the picture by claiming that the practice out of which theory arose was a simple application of the theory’s pre-existing abstract principles.

But, as is pointed out by the analyses of mediated cognition that our seminar is studying, there are other *artifacts* in which experience can be captured, preserved and transmitted (Cole, 1996). *Narrative* is one (Bruner, 1990). In this chapter, I have tried to project a voice which does not redefine the temporality of the experience I am reporting.

Sculpture is another way in which people impose meaningful form on nature and, as Hegel would say, externalize their consciousness through the mediation of wood, clay, plaster or stone, sharing it with others and preserving it as part of their culture’s spirit. Sculptures like that in figure 6-4 are such artifacts. They create spaces that project their own perspectives while at the same time being perceived from observational vantage points. Of course, Moore’s sculptures are not the result of some primordial experience of self-consciousness interacting with unmediated nature. They are late twentieth century explorations of form and material. Here, organic three-dimensional forms are showcased to contrast with socially prevalent two-dimensional representations and with the geometric shapes produced by machinery. The characteristics of the materials of nature are brought forth, in contrast to the plastic substances that retreat from our consciousness as commodities. Also, the pragmatic representational function of symbolic objects is sublimated in the study of their abstracted physical forms and materiality. In negating the commonplace characteristics of signs—which point away from themselves—the non-representational sculptures obtrusively confront their creator and viewers with the nature of the artifact as intentionally formed material object.



Figure 6-4. Henry Moore, Three Piece Sculpture: Vertebrae, 1968-69, bronze, Hirschhorn Sculpture Garden, Washington, DC. Photo by G. Stahl, 2004.

Figure 6-4 goes approximately here

Polished *software* is a very different way of objectifying experience. Buried in the source code and affordances of a software artifact are countless lessons and insights—not only those of the particular software developer, but of the traditions (congealed labor) of our technological world upon which that developer built (Marx, 1867/1976). This is as true of the current version of WebGuide as it is of any software application. So the software application is such an artifact; one that mediates classroom collaboration. But WebGuide strives to preserve insights explicitly as well, within the notes displayed in its perspectives and within their organization, including their organization into personal and group perspectives. The discussions that evolve within this medium are also artifacts, captured and organized by the perspectives.

Perhaps when we understand better how to use WebGuide in collaborative learning contexts it will maintain the knowledge that people construct through it in a way that preserves (in the sense of Hegel's synthesis or *aufheben*) the construction *process* as well as the resultant theory. Then we may have a type of artifact or a medium *that does not reify and alienate the process by which it developed*—that permits one to reconstruct the origin of collaborative insights without laboriously deconstructing artifacts that are harder than stone. Eventually, collaborative practice and software design may co-evolve to the point where they can integrate the insights of multiple perspectives into group views that do not obliterate the insights of conflicting perspectives into the multifaceted nature of truth.

8. Issues for Mediation

We conclude this chapter with an attempt to sort out what we are collaboratively learning through our use of WebGuide. The six issues for perspectives-based software like WebGuide that arose during the middle school application (section 5) appeared in the graduate seminar's usage of the software as well—and were articulated by seminar participants in their notes in WebGuide. These are important and complex issues that other researchers have raised as well. They are not problems that we have solved, but rather foci for future work. They define central goals for our redesign of WebGuide and goals for structuring the mediation of collaborative practices.

Here is a summary of our current understanding of these issues, based on our two practical experiences and our reflections on the theory of computer mediation of collaborative learning:

8.1 Divergence Among Ideas

In his review of computer mediated collaborative learning, dePaula (1998) identified divergence of ideas to be a common problem. He argued that the tree structure imposed by standard threaded discussion support was inappropriate for collaboration. The idea of a threaded discussion is that one contribution or note leads to another, so that each new idea is connected to its “parent” in order to preserve this connection. The problem is that

there is often no effective way to bring several ideas together in a summary or synthesis because that would require a particular note to be tied to several parent notes—something that is typically not supported by discussion software. The result is that discussions proceed along ever diverging lines as they branch out, and there is no systematic way to promote convergence (Hewitt, 1997). It seems clear, however, that collaboration requires both divergence (e.g., during brainstorming) and convergence (e.g., during negotiation and consensus).

WebGuide tries to avoid this common structural problem of threaded discussion media at three levels:

- The note linking mechanism in WebGuide allows notes to be *linked* to multiple parents, so that they can act to bring together and summarize otherwise divergent ideas. As in threaded discussions, every note is situated in the workspace by being identified and displayed as the child of some other note. However, WebGuide allows multiple parents, so that the web of notes is not restricted to a tree.
- Similarly, the graph of perspectives allows for *multiple inheritance*, so that “comparison” perspectives can be defined that aggregate or converge the contents of multiple perspectives. The Logan School application was seeded with comparison perspectives corresponding to the class and subgroup perspectives, so that the overall perspectives graph has a structure in which the inheritance of notes first diverges from the class to the subgroup and then the personal perspectives, and then converges through the subgroup comparison perspectives to the class comparison perspective, as shown in figure 2. The web of perspectives forms a directed acyclical graph rather than a strict hierarchy.
- Another effective way to encourage a well-structured discussion is to seed the workspace with a set of headings to scaffold the discourse. By introducing carefully conceived *headings* high in the perspective inheritance network, a facilitator (such as a teacher) can define an arrangement of topics that will be shared by the participants and will encourage them to arrange related ideas close to each other.

Although WebGuide provided these three convergence mechanisms in both of our usage situations, most participants were not adept at using any of them. This is probably related to the other issues below and is something that needs to be explored further in the future.

8.2 Avoidance of System Use

Media competition poses a barrier to acceptance of new communication software. People are naturally hesitant to adopt yet another communication technology. In a world inundated with pagers, cell phones, voicemail, email, fax, etc. people are forced to limit their media or be overwhelmed. They must calculate how much of a burden the new medium will impose in terms of learning how to use it, acquiring the equipment, checking regularly for incoming messages and letting people know that they are communicating through it. Clearly, a *critical mass* of adoption by one’s communication partners is necessary as well.

In a classroom context, some of these problems are minimized: all one’s partners are required to use WebGuide and the hardware is made available. Yet, it is not so simple. The Logan School students have to communicate with mentors who may not have

Internet access or the proper hardware. Communication with classmates is much easier face-to-face than typing everything (knowing it has to be carefully done for grading). In the graduate seminar, most participants do not have convenient access to the necessary equipment and have to go out of their way to a special lab. This means that they are lucky to communicate through WebGuide once a week, and therefore cannot enter into lively on-going interchanges.

We will have to make WebGuide more accessible by increasing the number of platforms/browsers that it can run on and making it work over slow modems from home. Further, we need to improve its look-and-feel to increase people's comfort level in wanting to use it: speed up response time, allow drag-and-drop rearrangement of notes, permit resizing of the applet and fonts for different monitors and different eyes, support searching and selective printouts, and provide graphical maps of the webs of perspectives and nodes.

8.3 Naturalness of the Perspectives Metaphor

Despite the fact that WebGuide has been designed to make the perspectives metaphor seem natural and simple to navigate, people express confusion as to how to use the perspectives. What perspective should I be working in, browsing for other people's ideas or entering for discussions? The metaphor of perspectives as a set of alternative (yet linked and over-lapping) textual workspaces is a new notion when made operational, as in WebGuide.

The fact that an individual note may have different edited versions and different linking structures in different perspectives, that notes may have multiple parents within the discussion threads, and that new perspectives can be added dynamically and may inherit from multiple other perspectives sets WebGuide apart from simple threaded discussion media. It also makes the computations for displaying notes extremely complex. This is a task that definitely requires computers. By relieving people of the equivalent of these display computations, computer support may allow people to collaborate more fluidly. This is the goal of WebGuide. Although the software now hides much of the complexity, it is not yet at the point where people can operate smoothly without worrying about the perspectives.

8.4 Representation of the Web of Perspectives

One problem that aggravates acceptance of the perspectives metaphor is that the web of inheritance of content from perspective to perspective is hard to represent visually within WebGuide. The WebGuide interface relies on an *outline* display, with hiding of expansion of sub-notes. This has many advantages, allowing users to navigate to and view notes of interest in an intuitive way that is already familiar. However, an outline display assumes a strictly hierarchical tree of information. Because the web of perspectives has multiple inheritance, its structure is not visible in an outline, which always shows a perspective under just one of its parents at a time. Thus, for instance, there is no visual representation of how a comparison perspective inherits from several personal perspectives.

The same is true at the level of notes. A note that has been linked to several other notes that it may summarize is always displayed as the child of just one of those notes at a time.

Two solutions suggest themselves for future exploration. One is to provide an alternative representation such as a *graphical map* in place of the outline view. As appealing as this idea sounds, it may be technically difficult to do on-the-fly. A bigger problem is that graphical maps are notoriously poor at scaling up. Already in our two trial situations—in which there are on the order of twice as many perspectives as participants—it would be hard to clearly label a graphical node for every perspective within the applet's confined display area. The second alternative is to indicate additional links with some kind of *icon* within the outline view. This would require more understanding on the part of the users in interpreting and making use of this additional symbolic information.

8.5 Structuring of Learning Situations

We have argued based on previous experience that the crucial aspect of supporting collaborative learning has to do with structuring social practices (Koschmann, Ostwald, & Stahl, 1998). *Practice*, in the sense of Bourdieu's concept of *habitus* (Bourdieu, 1972/1995), is the set of generally tacit procedures that are culturally adopted by a community. In introducing WebGuide into its two user communities, we have tried to establish certain usage practices, both by instruction and by enforcement in the software. Looking back at figure 1, one can see that Logan students are only allowed to navigate to certain perspectives—namely their personal perspective and those group perspectives that inherit from that perspective. Seminar participants were originally given *permission* to navigate throughout the system and to make changes anywhere. That was subsequently modified (as shown in figure 3) to restrict their abilities when not in their personal perspective. The governing principle was that everyone should be able to do anything they want within their personal perspective, but no one should be able to affect the display of information in someone else's personal perspective.

When the ability to enter notes everywhere was restricted, facilities for copying and linking notes from other computational perspectives into one's own computational perspective were introduced. This was intended to encourage people to integrate the ideas from other figurative perspectives into their own figurative perspective by making a conscious decision as to where the new note should go in their existing web of notes. However, this added a step to the process of communication. One could no longer simply select a note that one wanted to comment on and press the "add discussion" button.

In order to facilitate discussion of notes that one did not necessarily want to integrate into one's own perspective, the "add discussion" (annotation) button was then made active in all comparison perspectives. This led to minor problems, in that one could then not edit discussion notes that one had contributed in these perspectives. This could be fixed at the cost of additional complexity in the rules by allowing the author of a note to edit it in comparison perspectives.

More significantly, our experiments with changing permission rules pointed out that people were using WebGuide primarily as a threaded discussion medium for superficial opinions and socializing—and rarely as a *knowledge* construction space. Furthermore,

their ability to construct shared group perspectives on discussion topics was severely hampered by the lack of support for *negotiation* in the system.

8.6 Distinguishing the System's Capabilities

In iterating the design of WebGuide it became increasingly clear that what the system “wanted to be” (the design vision) was a *medium for construction of knowledge*. Yet, users were more familiar with discussion forums and tended to ignore the perspectives apparatus in favor of engaging in threaded discussion. These are very different kinds of tasks: collaborative knowledge construction generally requires a prolonged process of brainstorming alternative ideas, working out the implications of different options and negotiating conclusions; discussion can be much more spontaneous.

This suggests that more clarity is needed on the question: what is the *task*? If people are going to use WebGuide for collaborative knowledge construction then they need to have a clear sense of pursuing a shared knowledge construction task. The Logan students have such a task in articulating positions on acid mine drainage. However, much of their knowledge construction takes place in classroom discussion. They use WebGuide largely as a repository for their ideas. The seminar has been concerned with understanding a series of readings, so its participants have been more interested in exchanging isolated questions or reactions than in formulating larger integrative positions.

Our experience to date already suggests the complexity of trying to support collaborative learning. We should probably distinguish the software interface functions that support discussion from those that support knowledge construction. But this should be done in such a way that spontaneously discussed ideas can later be readily integrated into longer-term knowledge construction processes. Similarly, additional functionality—most notably support for group negotiation—must be added, differentiated and integrated. New capabilities and uses of WebGuide can increase its value, as long as confusions and conflicts are not introduced. For instance, providing facilities for people to maintain lists of annotated Web bookmarks, things-to-do, favorite references, upcoming deadlines, etc. within their personal perspectives might not only give them familiarity with using the system, but would also build toward that critical mass of usage necessary for meaningful adoption.

It has become a cliché that computer mediation has the potential to revolutionize communication just like the printing press did long ago. But the real lesson in this analogy is that widespread literacy required gradual changes in the skills and practices of the populace in order to take full advantage of the technological affordances of the printing press. In fact, the transition from oral tradition to literacy involved a radical change in how the world thinks and works (Ong, 1998). Although social as well as technical changes can be propagated much faster now, it is still necessary to evolve suitable mixes of practices and systems to support the move from predominantly individual construction of knowledge to a new level of *collaborative cognition*.

Our investigation of the above six issues will guide the next stage of our on-going exploration of the potentials and barriers of perspectives-based computer mediated collaborative learning on the Web.

9. Dialog with JIME reviewers

In fall 2000, the preceding part of this chapter was reviewed through the JIME on-line review process. I thought the reviews nicely brought out what the paper was trying to do. They added, in a generally supportive way, confirmation of one person's experiences from much broader backgrounds. The reflections on key issues significantly enriched the discussion.

Rather than disrupting the narrative flow of the report above, situated as it was in its particular phases of WebGuide development, responses to the reviewer comments and inquiries will be presented in question/response format below. This may serve as another layer of reflection, from a somewhat later vantage point.

Question:

A slight doubt, which I think it would be hard to understand without using the system for a while, would be if it could feel/be restrictive. When computer-mediated collaboration is "well used," users systematically attend to convergence (using the divergent discussion as a resource) by writing summaries and essays based on the shared material. Would WebGuide confine learner freedom to synthesize/converge because of the complexity of its linking systems... just a doubt.

Response:

While WebGuide's interface has improved considerably since its first usage, problems remain of trying to think about ideas on a computer monitor. It is still a less convivial environment to play with complexly inter-related ideas than is paper. There is also the difficult trade-off between simplicity and clarity of the interface and the desire to support complicated functionality. The mechanisms to support convergence are only partly automatic, transparent and natural. And yet, if we want to think and write collaboratively then paper will not suffice.

Question:

Does this WebGuide software provide a more straightforward discussion area to be used alongside of the work on perspectives? Somewhere here there seems to be confusion between a virtual collaborative discussion space and a tool to aid collaborative work. Another point which confused me was the idea of the software as artifact in the same way as a piece of sculpture or a narrative... even if as the author points out, software "represents a very different way of objectifying experience." Can various perspectives be represented with a single graphical image? Perhaps Cubist painting rather than sculpture makes a better analogy?

Response:

As detailed below, I have subsequently added a "discussion perspective" that provides a space for threaded discussion. Previously, threaded discussion took place directly in the comparison perspectives—leading people to ignore their personal perspectives and aggravating the conflict between discussion and construction. One of the hardest things I have had to figure out as a designer is how to integrate this into the perspectives framework, so that ideas entered one place would be available for the rest of the

knowledge-building process. I have just now implemented this and have not yet released it to my users. I have still not implemented the sorely needed negotiation procedures. Discussions with Thomas Herrmann and his colleagues in Germany have helped me to understand the issues related to these new perspectives, and why the system should include explicit discussion and negotiation perspectives.

An artifact is never a simple object. A sculpture, for instance, opens up a rich world: it not only structures physical space and offers a sensuous surface, it also evokes other objects, meanings and works. Software is yet harder to characterize: what is its form and substance, where are its resistances and affordances? A communication and collaboration artifact like WebGuide makes possible new forms of interaction and knowledge building—but how do people learn how to take advantage of this without being overloaded? The artifact here is not so much the buttons and windows of the user interface as the discussion content that gets built up through the interface. These issues have led me to another iteration of theory with a seminar in fall 2000 on how artifacts embody meaning and subsequent analysis of empirical data on how people learn to understand and use meaningful artifacts (see chapter 12 as well).

I like the cubist image. But sculptures also encourage and facilitate viewing from different visual perspectives. I have thought of replacing the mono-perspectival pictures in the chapter with video clips that could be run in the JIME publication. Perhaps I could just use animated gifs of each sculpture, that cycle through several views—creating an effect that cubism anticipated before perspectival technology was available.

Question:

This article does give us a lot of good, clear, qualitative description of the two situations in which this software is being looked at. At some point, though, there seems to be a need for firmer ground and a few numbers.

Response:

The middle school classroom had 12 students. During the several months of sporadic usage, 835 notes were entered (including revisions of old notes). This count includes guiding questions and organizing headings that the teacher and I entered.

The graduate seminar had 8 active students. During the semester, 473 notes were entered.

This semester (which is half over as I draft this response), there are 11 active participants. We have entered 497 notes already, but many of these are headings, modifications or entry of data to be shared. This probably represents an average of two entries per week per student. While I work on some technical problems that have arisen, I am not encouraging heavy use of WebGuide. Mostly entries are comments and questions on the class readings, with some follow-on discussion. If I defined some collaborative tasks, we might get much higher usage.

I try to hold class in a computer lab at the beginning of the semester so that we can learn the systems together and students can help each other. Most students can now access WebGuide from home, although this remains problematic. When we all use WebGuide at the same time in the lab, the worst technical problems come up (multi-user issues that are hard to test without class usage). Also, problems arise of how the entries are organized (how to find what one's neighbor just said she put in) and how discussion

relates to one's personal perspective. The main beneficiary of class usage of WebGuide is still the designer, who sees what problems need to be solved and what new functionality is desirable. For the students this is a glimpse into the future, but not yet a powerful cognitive and collaborative tool. In each class that uses WebGuide the students participate in reflecting on the process of designing the software artifact—and this is integral to the course curriculum as an experiment in collaborative learning.

Question:

I see an advantage of being able to see the work of other roles in progress. Figure 2 shows that joining of perspectives takes place way (too) late, namely in Gulch class comparison. It means students are not having enough time to prepare counterarguments and it also means that students miss out on constructing their perspectives along the same lines as that of other groups. In addition, I doubt whether it is desirable to have students think only about their own role or perspective since this is rather unrealistic (I may have this wrong, I am not sure how the system actually was used in practice).

Response:

I fear there is still some confusion on how perspectives work. The inheritance diagrammed in figure 2 takes place continually as notes are added, not just when perspectives are somehow complete. Every user of WebGuide can visit every perspective and read what is there at any time. The restriction is that you can only modify (edit, delete, rearrange) notes in your own perspective. Recently, I have added "private" notes that you can add in any perspective but are only viewable by you. This way, you can annotate any notes in the system privately.

I have also added "discussion" notes that you can add in any perspective; rather than staying in that perspective (and thereby modifying someone else's perspective), the discussion note and the note it is discussing are copied to a new "discussion perspective." The new discussion (and a new negotiation) perspective provides a space for inter-personal discussions to take place. Your contributions in the discussion perspective are also copied into your personal perspective so that you have a complete record of all the ideas you have entered into WebGuide and so that you can integrate these ideas with others in your working perspective.

These changes are part of a rather radical re-design—or at least extension—of the WebGuide perspectives system that has not yet been tried out by users. However, it is worth presenting here in some detail because it shows my response to the worrisome issues that have come up about conflicts between discussion and knowledge building (as discussed especially in point 5 of section 8 above). It brings the presentation up to date as of spring 2001.

Figure 6-5 shows the new interface of WebGuide. It now consists of two separate windows, a Java applet interface and an HTML window. Previous interfaces included an HTML frame within a fixed size main interface window. The user can now resize and overlap the two windows to optimize and personalize use of screen real estate. The main interface consists of (a) an expandable hierarchy of notes (either their titles or the first line of their content is displayed in the hierarchy—the full content of the currently selected note is displayed in the HTML window), (b) a bar of buttons for selecting a perspective across the top and (c) a control panel of function buttons on the right side.




Figure 6-5. The new interface to WebGuide 2000.




Figure 6-6. The new bar of perspectives buttons in WebGuide 2000.

Figure 6-5 goes approximately here

Figure 6-6 shows a close-up of the perspectives buttons, providing direct access to the most common perspectives and a pull-down list of all defined perspectives in the current database. Note that in addition to the *group* (or class) perspective, the current user's *personal* perspective and the (group or class) *comparison* perspective, there are now perspectives for *discussion*, *negotiation* and *archive*. We will see how these are inter-related in figure 6-8 below.

Figure 6-6 goes approximately here

Figure 6-7 goes approximately here

Figure 6-7 shows a close-up of the function controls, with restricted options grayed-out. The *comment* button allows a user to enter a quick comment below the selected note. The *new note* button is similar to the *comment*, but allows the user to choose a label for the kind of note and to position the new note after (i.e., at the same level of hierarchy) the selected note rather than indented below it (i.e., as a child of it). Subsequent buttons let the user *edit*, *delete*, *move*, and *copy* or *link* a selected note. *Copy to home* or *link to home* is used when one has selected a note that is not in one's personal perspective and wants to create a physical or virtual copy of it there. *Email* lets one send an email and have the content of the email and its responses inserted below the selected note. *Search* conducts a simple string text search across all notes (their author, title and content) in the database and displays the resulting notes in the HTML window (where they can be easily printed

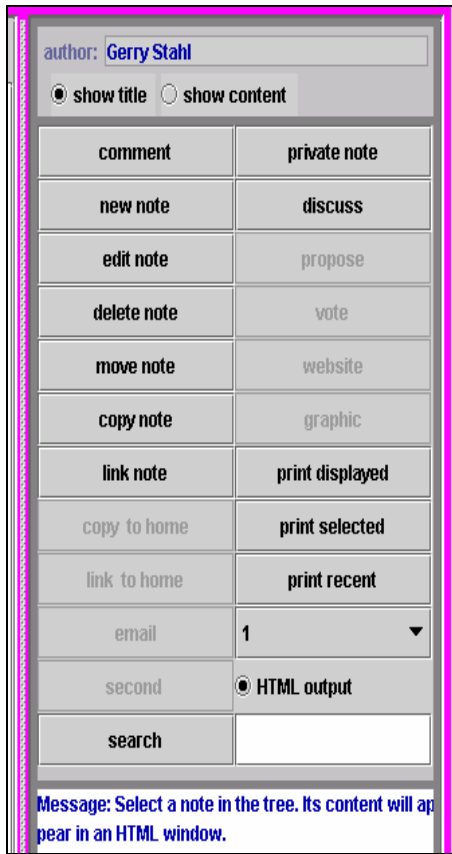


Figure 6-7. The new knowledge management control panel in WebGuide 2000.

out). *Private note* is similar to *comment*, except that one can insert it in any perspective and that it will only be displayed when the author is logged in as the current user. *Discuss* and *promote* create notes in the discussion and negotiation perspectives; they will be described in the next paragraph. The *vote*, *website* and *graphic* buttons are for adding votes on negotiation issues, live links to URLs and *graphic* (multimedia) URLs to be displayed in the HTML window—these functions are not yet implemented. The *print displayed* button causes all notes whose titles are currently displayed in the hierarchy display to have their content shown in the HTML window for printing. The *print selected* button lets a user select multiple notes that are not sequential and have their content displayed in the HTML window. Finally, the *print recent* button displays in the HTML window the content of all notes that were created in the past N days, where a value for N is selected below this button. These search and print buttons are important steps toward providing more effective knowledge management—offering convenient access to selected notes.

How should the *discuss* and *propose* buttons work? A user should be able to start a discussion based on any other user's note found in the system. The resulting discussion should be available to everyone in the group. The two perspectives available to everyone are the group and the comparison perspectives. The comparison perspective quickly becomes over-crowded and confusing, so I decided to create a new discussion perspective derived from the group perspective. Similarly, proposals for negotiations should be able to build on anyone's notes and should be generally available, so I also created a negotiation perspective linked to the group perspective. Recall that the group (or class) perspective contains notes agreed to by the group at large (or seeded by the teacher to provide a shared starting point). The group perspective therefore provides an over-all context for collaborative discussion and negotiation, as well as for individual efforts at knowledge building. So, while we do not want discussion and negotiation notes that have not yet been adopted by the whole group to show up directly in the group perspective (and therefore to be inherited into all other perspectives), we do want to have the discussion and negotiation perspectives inherit from the group perspective in order to provide some context and structure. Moreover, we want the negotiation to inherit from the discussion so that a note in a discussion thread can be proposed for negotiation and so that discussion threads can be viewed in relation




Figure 6-8. The old inheritance structures for perspectives in WebGuide (on the left) and the new structures (on the right).

to negotiation proposals. As shown in figure 6-8, individual personal perspectives should inherit from the group but not from the discussion or negotiation perspectives.

Figure 6-8 goes approximately here

The trick with putting notes in the discussion and negotiation perspectives is to situate them meaningfully in the hierarchy with at least some context. Suppose you have entered a note that I want to comment on and to present for group discussion. Your note is in your personal perspective and I may have found it in the comparison perspective. So I either select your note in the comparison perspective or go to your personal perspective and select it there. I click on the *discuss* button. The system then wants to start a thread in the discussion perspective starting with your note, which would then be followed by my note. To do this, the system finds your note's (the one that I want to comment on) parent-note—the note that your note is threaded from in the hierarchy of your personal perspective—and designates that note the anchor note. If the anchor note happens to already appear in the discussion perspective (which inherits the whole group perspective), then everything is simple and the system simply makes a copy of your note below the anchor in the discussion perspective and attaches my note below that. Alternatively, if an ancestor of your note appears within the discussion perspective in the notes hierarchy, then that closest ancestor is used as the anchor. Otherwise, the system attaches a copy of your note to a special “Discussions” heading note in the discussions perspective and then attaches my note below that. Then we have a discussion thread that anyone can add to in the discussions perspective.

In addition to setting up the new thread in the discussions perspective, the system makes a copy of your note with mine attached to it and places them below the anchor note (which I inherit from the group) in my own personal perspective. This is so that my personal perspective contains all of my contributions to discussions and negotiations.

modes:	group perspective	discussion (group)	negotiation (group)	personal perspective	other's perspective	comparison perspective	archive (deleted)
buttons:							
discuss		X	X	X	X	X	X
propose			X				
vote			X				
private note	X	X	X	X	X	X	X
new note				X			
edit				X			
delete				X			
move				X			
copy				X			
link				X			
copy home		X	X		X	X	X
link home		X	X		X	X	X
email				X			
search	X	X	X	X	X	X	X
print	X	X	X	X	X	X	X
displays:							
Statement	X	X	X	X	X	X	w. deleted
Discussion		X	X	viewer's	owner's	X	w. deleted
Proposal			X	viewer's	owner's	X	w. deleted
Decision			X	viewer's	owner's	X	w. deleted
Private	viewer's	viewer's	viewer's	viewer's	viewer's	viewer's	viewer's

Table 6-1. Table of permissions for WebGuide 2000 buttons and displays.

That way, I see all of my ideas and I can conveniently manipulate them in my own workspace. The dotted line in figure 6-8 from negotiation to viewer's perspective indicates that these entries will appear in my perspective when I am viewing it.

Similarly, figure 6-8 indicates that private notes that I created with the private note button will appear in whatever perspective I created them in when—and only when—I am viewing them. Finally, the archive perspective is simply the group comparison perspective, including notes that have been deleted. This is primarily for the convenience of researchers who want to view old versions of work. Figure 6-8 shows how the inheritance structure has changed with the recent addition of the discussion, negotiation, private and archive perspectives. The possibility of extending the perspectives metaphor and the underlying computational mechanism to include new perspectives like these confirms the power and generality of the approach.

Table 6-1 summarizes the relationships of the buttons and display modes to the different perspectives. The top of the chart ("buttons") indicates the perspectives in which each of the buttons is active (i.e., not grayed out).

Table 6-1 goes approximately here

The bottom of the chart (“displays”) indicates which notes are displayed in each type of perspective (and in some cases to whom it is displayed). “Statements” are notes created with the *comment* or *new note* button; “Discussions” are created with the *discuss* button; “Proposals” with the *propose* button; “Decisions” with the *vote* button; and “Privates” with the *private note* button. The viewer is the currently logged-in user; the owner is the person to whom the personal perspective belongs.

Yes, this is obviously the designer’s story. I think it is premature to give a user’s story. For a number of reasons that are my fault (technical problems and poor definition of tasks), WebGuide has been at best used as a threaded discussion forum. I hope that the new structures of Negotiation, Discussion and Private perspectives will help users to engage in personal and group knowledge building. Perhaps then we will see some insightful user scenarios.

Question:

Which ideas led to WebGuide? Obviously, the author has not started to work on WebGuide with zero theory. Unfortunately, the author does not make his starting notions (theory) very clear in the design narrative. I think this is an omission that should be corrected. If I presume correctly, the author has detected some real-life problem for which there was no adequate solution. This triggered WebGuide development.

Response:

As for fore-runner systems, I make no claims that WebGuide is superior to other research prototypes or even commercial systems for supporting collaborative learning. It is consciously based on Scardamalia and Bereiter’s extensive theoretical, technical and pedagogical work on knowledge-building communities and their CSILE (now Knowledge Forum) system that is used in schools around the world. WebGuide’s only attempted innovation is perspectives. The idea of perspectives grows out of my dissertation work with Ray McCall and is based on ideas cited in this chapter and chapter 4.

Where did the original impetus for WebGuide itself come from? This is another story, told sketchily in chapter 5. While introducing the software from chapter 2 in another local middle school, I observed problems of students collecting and retaining website addresses as part of their Web research projects. I thought it would be nice to let them save these addresses on the Web, rather than on disks or floppies that never seemed to be available when they needed them. So WebGuide was originally conceived of as their personal guide to the Web, with their collected website links. Then I wanted them to be able to share their links and negotiate class-adopted lists of links. Then I added the idea of annotating and eventually discussing the links, and finally categorizing and reorganizing them. Soon, the superstructure took over and I have still not made it easy to store links in WebGuide. The WebGuide interface has always included an HTML window as well as the Java applet display. The content of notes is displayed in the HTML window—specifically so that website links can be live and one can click on them and go to the site. This also means that graphics and other media can be stored in WebGuide and viewed, and that HTML markup can be used in the content. As for the philosophy

behind WebGuide, the notion of perspectives goes back to a former life when I studied Heidegger and hermeneutics (Stahl, 1975a), as well as to my more recent computer science dissertation (Stahl, 1993) that argued for this kind of software perspectives mechanism—warranted by reference to ideas of design theorists Rittel, Alexander, and Schön (see chapter 4). So persistent questioning pushes the horizon of context further and further back through forgotten cycles of practice and theory, complexly evolving trajectories of inquiry that had no clean starting point *ex nihilo*.

Question:

The problem of system under-utilization is worldwide and well known. Which functionalities that you consider to be key functions did users underutilize?

Response:

Sure, I do not care if students do not use all the features of WebGuide either—and I do not provide a lot of formatting, etc. in the first place. But I would like to see them get beyond mere threaded discussion—the superficial exchange of off-the-cuff opinions—to deeper collaborative knowledge building. Seriously taking up each other's ideas and formulations, worrying about terminological disagreements, negotiation of innovative insights that merge multiple perspectives: these would be exciting to see emerge from the more sophisticated use of WebGuide's functionality, which allows notes to be modified, copied, rearranged, etc. across perspectives.

Question:

I wondered whether the author has worked in a situation in which there were not enough computers for all students forcing the formation of groups.

Response:

Periodically, students have teamed up on computers. This is nice for collaboratively learning how to use the system. It is also useful if I want to videotape the usage and analyze the discourse within the little group for a fine-grained view of what is going on from the user perspective. The problem with doing this with WebGuide is that it is text-based and only one person can type text into the shared computer at a time.

Question:

Perspectives? How should I fit in the notion of 'role' (e.g., that of landowner) within the typology of literal, figurative and computational (section 3). And what should I think about different points of view within roles? Are these perspectives too (as seen from the designer's point of view)? And what about class perspective, team perspective?

Response:

The perspectives mechanism of automatic inheritance of content down the hierarchy is very general. In some cases I have used it to define a hierarchy of domain knowledge (my dissertation), of roles (the middle school Gulch project), of academic disciplines (the interdisciplinary seminar), or just of different people (this semester). The group or class perspective is supposed to display the state of knowledge that has been mutually agreed upon, and thus requires the still-missing negotiation support. So now it contains mostly what the teacher has defined by fiat as the shared knowledge structure in order to get the

process going in an organized way. Yes, this is all hard to explain or comprehend, especially without actually using the system.

Summary:

I found the JIME reviews very heartening. The reviewers clearly understood and appreciated what I was trying to do with my narrative approach to reporting on recent research. Furthermore, they added important critical perspectives. My concluding response to the reviews consists of a brief overview of the adventure of composing this chapter.

The chapter grew out of a submission to AERA '99 (the annual conference of the American Educational Research Association). To have a paper accepted to this conference, one simply submits an abstract. When my abstract was accepted, I felt free to write in whatever vein I chose. The freedom from having to write to traditional reviewers with narrow paradigms of scholarly publication allowed me to experiment stylistically as well as to think about what format would be most appropriate to the level of experience with WebGuide that I wanted to report.

The paper session at AERA was coordinated by Ricki Goldman-Segall, who served as the discussant as well. She shares my enthusiasm for "perspectivity software." I had just read her book (Goldman-Segall, 1998), which has a "thick description" style of interwoven themes and which precedes each chapter with one of her photographs. This gave me the impetus to tell my story by talking about the diverse themes which were important to me. I also decided to introduce a decorative element to the page like Ricki did, and to tie my ideas about sculpture loosely to my content. (In the book version, I have removed pictures of my own sculpture and added figure 6-4.)

It was clear to me that providing a traditional analysis of the software usage would have been wildly premature. While the use of WebGuide by one teacher and his dozen students over several months had made a number of technical and social issues painfully clear to me, and while the experiment had been an experience for the students, there was nothing entered into the database to illustrate the ultimate vision I had for the software approach. Similarly, in my graduate seminar with about eight students for a semester, WebGuide served more as an example of what we were thinking about than as a tool that let us think about it more deeply. What was interesting was not the empirical data about the software usage, but the process ("dance," "structural coupling") by which our understanding of what was needed developed in the classroom settings where a crude version of WebGuide was used.

The work on WebGuide continues to be the focus of my activities in Colorado. Many of the weaknesses pointed out in the reviews are being gradually addressed in new software functions, theoretical papers and funding proposals. This evolving article remains my fullest discussion of perspectives and their inheritance, a topic that is devilishly hard to explain clearly. WebGuide 2000 is now being used in my seminar on CSCS. Every month I produce a new version with additional improvements. However, while some students are starting to use it regularly to formulate and discuss ideas, its use still falls far short of the goal.

It has become clearer to me that WebGuide needs to be a *collaborative knowledge management environment*. It needs to better support the browsing, modifying and re-organizing of inter-related ideas. "Knowledge building" has become a more central

concept for me and I am trying to understand how it proceeds or could proceed: what activities are involved and what tools could support these interpersonal activities. Talking about knowledge building (a concept I attribute to Carl Bereiter) seems to be a productive way to think about learning in a social and collaborative framework. The subtle intertwining of group and personal perspectives is a central structure of collaborative knowledge building.

The notion of “artifacts” has become ever more central to my theoretical interests. My seminar this semester is on the question of how artifacts—particularly computer-based artifacts like WebGuide—affect our cognitive abilities. How do artifacts embody meaning, how do people design that meaning into them and how do others learn what that meaning is? What are the implications for designing new media to support thinking and collaborating? This week we are reading Heidegger’s discussion of how works of art like sculpture not only make explicitly visible their forms, meanings and material, but actually open up whole new worlds in which human activities can take place. What kind of world do we want to create for future WebGuide users? What kinds of intellectual worlds do we want students to collaboratively construct for themselves?

The problems of getting communities of students to adopt Web media like WebGuide are daunting. Look at our use of the JIME technology in reviewing the journal version of this essay. The idea that the online JIME discussion medium might support a back-and-forth knowledge-building discussion among the reviewers and with the author—grounded in the artifact of the submitted article, section by section—was not realized. None of the reviewers knew how to use it effectively. They probably first typed up typical reviews in their word processors and then pasted them into the top of the discussion hierarchy. Then they broke them up and stuck some pieces under different headings, but never in the places that were linked to article sections. Months later, the author had to respond in a similar way. The editor of the reviews did not even post his thoughtful contributions to the online JIME site at first, but emailed them separately. Unfortunately, this is typical not only of JIME and WebGuide, but of groupware in general (Lenell & Stahl, 2001). These are the pressing issues that need to be discussed at this stage, more than details of technology and statistical assessment methodology.