

“ITR/IM: Perspectives on Collaborative Knowledge-Building”

Project Summary

Knowledge-Building: This project proposes to develop an approach to information management grounded in an innovative theory of *learning and collaboration as knowledge-building*. The theory is oriented toward guiding development of technology that can better fulfill growing societal requirements for Web-based support of groups. It offers an analysis of collaborative knowledge-building activities and of Knowledge-Building Environments (KBEs) to support these activities. To do this, it brings together and synthesizes approaches and concepts from situated learning, activity theory, hermeneutics, distributed cognition, and related theories from philosophy, social science, education, and computer science.

Knowledge-Building Environments: The proposed work centers on a new way of using computers to support collaboration by *integrating information management support for various activities of knowledge-building*. Such KBE software prototypes go beyond superficial discussion or chat and beyond choices and decisions among fixed options to support the co-construction of deep knowledge, innovation, and shared understanding. They support learning, working, and innovation over time within groups that may be physically distributed.

Perspectives: A key innovation for KBEs in this project is the modeling of the *interplay between individuals and the group* in collaborative settings. This is done by providing personal and group computational perspectives: virtual workspaces whose contents are interrelated by automated inheritance mechanisms. Computational perspectives provide a new, dynamic, personalized form of on-line information management that supports the fundamental structure of collaboration. They help to manage a shared information space so that participants view information relevant to themselves and can process (edit, rearrange, reconceptualize) that information without affecting anyone else’s personal perspective. Then, through supported negotiation activities, information is migrated to sub-group and group perspectives, where it represents shared knowledge.

Research and Training Focus: The project will directly employ four students and will involve many more in seminars and class projects. Because KBE software is essentially a new form of learning technology, education students as well as computer science and other students will be involved in designing, developing, deploying, and assessing the software. This will help to create an *interdisciplinary educational technology research focus* that the PI and his colleagues have already begun to foster.

Software Development: The proposed work will build an infrastructure for local, national, and international collaboration on KBE software development, providing a more reliable *basis for assembling KBE prototypes* customized to particular deployment sites. An open source perspectives server will be released, allowing researchers to develop KBE interface components that simply call this server for database access and dynamic computation of perspective contents. A standard for data interchange with the server and for interoperability among KBE systems will facilitate a component architecture and the use of shared tools to assess KBE usage.

Study Sites: The project will assess KBEs with perspectives in realistic study sites: an academic research group, a collaborative learning seminar, a corporate training setting, and an industrial design group. *Quantitative analysis* of captured textual contents will be compared with results from non-integrated threaded discussion systems and other groupware. *Qualitative analysis* of surveys and field notes will investigate issues of deployment, adoption, social practice, utility, and effectiveness.

Project Impacts: The proposed project should result in progress in the development and assessment of KBEs, an emerging form of software with a potential to significantly extend human cognition by supporting collaborative knowledge-building activities and by providing persistent external memory of what took place during the collaboration. The release of a perspectives server with its associated standards will provide a concrete basis for catalyzing *local, national, and international collaboration* among KBE researchers. This will promote KBE research as an important new research focus.

“ITR/IM: Perspectives on Collaborative Knowledge-Building”

Project Description

This project is guided by a *theory of collaborative Knowledge-Building Environments* (KBEs) that we are developing. This theory proposes the following principles:

- *Collaborative knowledge-building* is a particular view of group learning that focuses on a range of activities that take place within communities, as opposed to focusing on learning as the transmission of bits of information to individual learners.
- Collaborative knowledge-building takes place largely through the interaction among people with different understandings from *multiple personal and group perspectives*.
- Such knowledge-building within groups can be helped by appropriately designed *information technology* (IT) that supports various knowledge-building activities and supports interaction among alternative perspectives.

The form of IT that we are interested in – collaborative Knowledge-Building Environments – represents a distinctive approach that overlaps related work in Computer-Supported Collaborative Learning (CSCL) and Computer-Supported Cooperative Work (CSCW). IT support for learning is traditionally oriented toward the transmission of information to individual students. Even where it is based on a view of student construction of knowledge, as with Intelligent Tutoring Systems (ITS) for algebra or physics, the goal is measured by testing the incorporation of pre-defined content or methods into the individual's understanding (Wenger, 1987). A more student-centered, constructivist approach is taken by Interactive Learning Environments (ILE), which might, for instance, allow students to create ecologies in *SimLife* to learn biology, or programs in *Turtle Logo* to explore math concepts (Papert, 1980). In contrast, a KBE primarily supports the group process and leaves matters of content up to the participants (which may include a teacher who raises particular content issues and helps maintain focus). In this way, it applies CSCW approaches to CSCL. A review of CSCW technology for groups (Kraemer & Pinsonneault, 1990) distinguishes group communication support systems (GCSSs) from decision support (GDSSs). GCSSs are specific communication media like email and video-conferencing. In providing computational tools for group decision making, GDSSs tend to support isolated, focused activities that integrate products of individual work. In contrast, a KBE aims to support a broad spectrum of knowledge-building activities – both individual and group – in a more seamless fashion. It supports the construction of areas of knowledge through group inquiry over extended periods of time. It also supports the interplay of individual and group more comprehensively, through integrated mechanisms of "computational perspectives" and negotiation that treat the group as more than just the sum of the individuals.

Assessments of CSCL and CSCW systems have defined a number of key issues for evaluating the problems and successes of such systems. For instance, in simple threaded discussion forums common problems include: short threads (a tendency for discussions to die quickly), low participation (lack of motivation to participate), few cross-references (little convergence of ideas), and superficial content (minimal depth of investigation) (dePaula, 1998; Guzdial & Turns, 2000; Hewitt & Teplovs, 1999). On the other hand, GDSSs and GCSSs attempt to decrease communication barriers within the group, while increasing task-oriented focus, depth of analysis, and decision quality (Connolly, 1997; Kraemer & Pinsonneault, 1990). Social informatics studies have raised additional issues of software deployment and adoption in addition to questions of usability and utility (Kling, 1999). These are some of the dimensions along which KBEs must be assessed within realistic learning and working social contexts.

To date, the PI and his colleagues have begun to develop KBE theory in conjunction with Web-based KBE prototypes that support many of the activities described in the theory and that have been tested informally in collaborative learning classrooms. In particular, computational support for *personal and group perspectives* has been developed and tried out. Support for computational perspectives was explored in the PI's dissertation (Stahl, 1993a) and has since been refined and adapted to the Web (Stahl, 1999a). This work has been described in relevant CSCL and CSCW conferences (see references by Stahl). The proposed project will build on existing concepts and prototypes, extending them substantially by: implementing a technical infrastructure to support data interoperability, integration of functionalities, and rapid prototyping; deploying customized KBEs in specific study sites; observing the social impacts of these IT systems in the work settings; revising the theory based on empirical findings; and fostering a community of researchers working on IT support for knowledge-building in workgroups.

1. PREPARATION FOR PROPOSED WORK UNDER PRIOR NSF SUPPORT

1.1. Organizational Memory and Organizational Learning (CSS)

“Conceptual Frameworks and Computational Support for Organizational Memories and Organizational Learning (OMOL),” PIs: Gerhard Fischer, Gerry Stahl, Jonathan Ostwald, September 1997 – August 2000, \$725,000, from NSF CSS Program #IRR-9711951.

This grant prepared much of the background for the proposed work. The OMOL project started from a model of computer support for organizations as Domain-Oriented Design Environments (DODEs) in which both domain knowledge and local knowledge are stored in the form of artifact designs and associated design rationale (Fischer, 1994). This CSCW model evolved into one of Collaborative Information Environments (CIEs), that emphasized the interactive, asynchronous, persistent discussion of concepts and issues within an organization (Stahl, 1998; Stahl, 2000a). Gradually, interest in organizational learning aspects led to involvement in CSCL and the model of collaborative Knowledge-Building Environments (KBes) (Fischer et al., 1999). A number of software prototypes were developed to explore the use of the Web as a communication and collaboration medium. Of these, the most important for the proposed work are the following:

- **DynaClass:** A discussion forum for use in college courses. It features ties to **DynaGloss** and **Sources** as well as email notification and specialized displays (Ostwald, 1999).
- **WebGuide:** Differs from **DynaClass** in providing more control over rearrangement of notes; features computational Perspectives (Stahl & dePaula, 1998; Stahl et al., 1998).
- **DynaGloss:** A system for defining technical terms and keywords and for debating the definitions and reviewing the history of debate; linked to **DynaClass** and **Sources** in that each term shows all the locations in these other systems where the term is explicitly referenced.
- **Sources:** A system for annotating bibliographical entries; uses terms from **DynaGloss** as keywords.
- **InfoMap:** An interface component for creating a graphical display of linked notes like a threaded discussion; providing convenient drag-and-drop functionality.

Work on this grant led to the focus on KBes as models of computer support for organizational memory and organizational learning. In particular, it provided a number of different systems, each with useful functionality, and brought home the need to define component standards so the functionalities can be combined more flexibly. As we tested and deployed these systems, we confronted serious issues of adoption and focused our concerns increasingly on socio-technical and social informatics (Kling, 1999) issues: motivation, media competition, critical mass, social practices, seeding, management, re-seeding, convergence of ideas, peer-to-peer collaboration, deployment strategies. These issues led to a new research agenda (Stahl, 1999b) and this proposal.

1.2. WebGuide and Environmental Perspectives (NOAA)

“Collaborative Web-Based Tools for Learning to Integrate Scientific Results into Social Policy,” PIs: Ray Habermann, Gerry Stahl, November 1998 – July 1999, \$89,338, NSF, #EAR-9870934.

This grant funded the initial implementation of **WebGuide** as an integrated Java applet KBE supporting personal and group Perspectives. It was a joint effort between the PI, a middle school teacher, and a research group at the National Oceanographic and Atmospheric Administration (NOAA) labs in Boulder. The teacher taught an environmental science class in which he wanted to spend the year having his students interview various adults and construct a set of contrasting perspectives (conservationist, regulatory, business, community) on a particular local environmental issue that the students had previously been involved in. **WebGuide** was used by the students to collect notes on their interviews and to formulate personal and team perspectives on the issue. Results of this software trial were analyzed and presented at conferences (Stahl, 1999a; 1999b; 1999c; Stahl & Herrmann, 1999). These findings led to a number of revisions of **WebGuide**, including the separation of the Perspectives mechanism from the Web interface, and recognition of the need for software architectures, standards, and components to support flexible rapid prototyping of KBes.

1.3. Collaboration in KBes (CILT)

“Interoperability Among Knowledge-Building Environments,” PI: Gerry Stahl, September 1999 – August 2000, \$9,124.21, from NSF-funded Center for Innovative Learning Technology (CILT), Subcontract #17-000359 under NSF grant #EIA-9720384.

This is a current seed grant whose purpose is to stimulate collaboration among KBE research groups. Part of the intention of the grant was to prepare a proposal for fuller funding, such as the present proposal and its currently pending complementary NSF proposals for “IT Support for Knowledge-Building in Workgroups” and “ITR/EWF: Collaborative Research on Knowledge-Building Environments: Growing a National and International Research Community for Distance Learning Information Technology.” This grant has already resulted in a semester-long student project involving three graduate and three undergraduate students (one collaborating virtually from Germany) creating an XML DTD that defines a data format for data imported from several different KBE prototypes and displayed in a Web browser using XSL. The grant supported a workshop entitled “Collaborating on the Design and Assessment of KBEs in the 2000's” at CSCL '99 at Stanford. This workshop attracted over 60 participants and was preceded by an on-line discussion of 28 submitted position papers. This grant has led to the emphasis on collaboration among KBE research groups and the need to put into place some of the technical and social conditions for such collaboration (Stahl, 1999a), as proposed here.

2. OVERVIEW OF PROPOSED WORK

2.1. Theory of Collaborative Knowledge-Building Environments

Collaborative Knowledge-Building

Information Technology (IT) is a broad field that can be conceptualized in various ways. Traditionally, the computer was thought of as a medium for storing and delivering data, that can then be used by people in their work. More recently, the computer (especially with the Web) has become a medium of communication, through which people share information and knowledge. This communication can take a variety of forms. In simple forms of e-commerce or on-line voting, people submit their decisions about a fixed list of choices. In chat and most email, people exchange greetings and opinions, generally without changing those opinions. Many systems in recent years have tried to support a particular form of communication or social interaction like brainstorming or decision-making – often with very positive results (Connolly, 1997; Vogel et al., 1987). We are interested in a distinct but broader process of communication which we term *collaborative knowledge-building*. Here, groups of people construct new knowledge through interaction of their ideas and perspectives, usually eventually preserved in documents or other artifacts.

Our theory of collaborative knowledge-building (Fischer et al., 1993/1998; Stahl, 1975; 1993a; 1993b; 1999c; 2000a; 2000b; Stahl & Herrmann, 1999) proposes a concept we call *the synergistic moment*; we intend to investigate the validity of this concept in the proposed project. The synergistic moment is the critical point during collaboration in which a group constructs meaning that transcends what any participant may have “in mind.” The *shared understanding* that is generated in this process is a subtle phenomenon: It does not mean that everyone is in complete agreement or even that each individual has the same internal cognitive representations of what is discussed. Rather, it means that a certain group view has been expressed. The unit of analysis for describing this is the group, and is manifested in the group's discourse. Individuals may agree to disagree with the group understanding, and careful investigation may reveal that individual understandings differ from the group's view (Hatano & Inagaki, 1991). The intersubjective “sharing” is not a correspondence or overlapping of individuals' mental content, but a coordination or interaction of their participation in joint socio-cultural activity (Matusov, 1996). The synergistic moment is an emergent property of the group dialog as a cacophony of voices (Bakhtin, 1986). It could easily pass unnoticed as a magical fount of creativity; to more deeply understand it likely requires “thick description” (Geertz, 1973) and detailed interaction/discourse analysis (Jordan & Henderson, 1995), and therefore presupposes that the interaction was captured in some medium. Fortunately, the literature on CSCL contains a number of incisive analyses (Roschelle, 1998) of the synergistic moment, although they do not highlight it as such.

The synergistic moment is a result of perspective-sharing (Boland & Tenkasi, 1995), but at the group rather than the individual level. It overcomes the problem pointed out by Feltovich et al. (Feltovich et al., 1996), that any one perspective may limit the ability to comprehend creatively the complexity of a topic under discussion. What typically happens is that one person makes a statement from her personal perspective; someone else interprets that statement from his own perspective and responds accordingly; others continue this process so that the discourse consists implicitly of reinterpretations from various perspectives. The drive to establish intersubjectivity and shared knowledge is powered by socio-cognitive conflict and contention among perspectives according to studies by Piaget and his followers (Perret-Clermont & Schubauer-Leoni, 1981). The dialog proceeds through sequential turn-taking

and attempts to repair “misunderstandings” as understood from particular perspectives and reinterpreted from others. Thanks to the human drive to impose coherent social meaning structures (Geertz, 1973), a synergistic group understanding emerges. This shared understanding can play a central role in the further activity of the group and can be more or less adopted by individuals into their personal perspectives. Although the synergistic moment seems to the participants to emerge spontaneously, it can be understood as the result of many identifiable knowledge-building activities, as represented in our model (below).

Perspectives in Knowledge-Building

According to hermeneutics – the philosophy of interpretation – human understanding is fundamentally perspectival. We construct knowledge from our situated perspective in the world: our historical position, cultural tools, and practical interests (Gadamer, 1960/1988; Heidegger, 1927/1996; Stahl, 1975). Computational support for knowledge-building can represent our interpretive perspectives with computational Perspectives (Boland & Tenkasi, 1995; Nygaard & Sørgaard, 1987; Winograd & Flores, 1986). (In this proposal, Perspective-with-a-capital-P will refer to the proposed computational mechanism that mirrors human interpretive perspectives-with-a-lower-case-p.) In this sense, Knowledge-Building Environments (KBEs) with computational Perspectives are designed to support the essential structure of collaboration. A key hypothesis of the proposed work is that KBEs benefit from an approach that represents the perspectival nature of collaboration. A goal of the project is to facilitate the incorporation of a computational Perspectives mechanism in KBEs – both in our own prototypes and in the work of other KBE research groups around the world.

Computational Perspectives have been explored by the PI in a number of software prototypes, in his dissertation system, and in his theoretical publications (Stahl, 1993a; 1993b; 1995; 1998; Stahl & Herrmann, 1999; Stahl et al., 1995). In a single-user system, computational Perspectives may correspond to different domains or professional viewpoints on a design problem, such as electrical, plumbing, structural, and heating concerns in architecture (Fischer et al., 1993; 1993/1998). In a KBE to support collaboration, computational Perspectives typically provide personal or group workspaces for the development of different sets of ideas. In this way, they can model the relationships among the various personal and group interpretive perspectives at work in the construction of collaborative knowledge.

We hypothesize that computational Perspectives can support the synergistic moment in collaborative knowledge-building by providing the necessary contact among different personal Perspectives, allowing them to interact, and then locating the results in a group Perspective. By situating the traditionally ephemeral synergistic moment within an explicit structure of computational Perspectives and by doing so in a persistent way, a KBE provides new opportunities for group self-reflection.

An important complement to Perspectives is *negotiation*. Negotiation is a process through which divergent personal perspectives converge on a collaborative shared understanding. When Perspectives and negotiation are effectively “intertwined” in a KBE, they compensate for each other’s potential problems: Negotiation converges ideas so that everyone can benefit from the ideas of other perspectives, while personal Perspectives allow people to work on their own views while potentially time-consuming negotiations are underway (Stahl & Herrmann, 1999).

For instance, when *WebGuide* – a KBE with computational Perspectives implemented by the PI – was used in a middle school environmental science classroom, students each had their own personal Perspective in which to develop their own responses to questions posed by the teacher. The teacher’s questions to the whole class were posed in the class’ group Perspective. From there they were automatically inherited into the team Perspectives. The content of the team Perspectives was, in turn, inherited into the personal Perspectives of team members. Gradually, students migrated their ideas to team Perspectives that represented either conservationist, governmental, corporate, or citizen perspectives on the ecological controversy – depending on which perspective team the student was part of. Then they could work with the ideas of their team-mates and negotiate their team position. In the end, the different teams negotiated to spell out agreements and disagreements (Stahl, 1999c).

The analysis of the synergistic moment suggests that negotiation need not take the explicit, rationalist forms typical of GDSSs, such as voting. Group results may emerge naturally out of the intertwining of Perspectives in group discussion. A challenge of the proposed work will be to develop software support for capturing such results and migrating them un-intrusively to group Perspectives.

The Potential of IT Support for Knowledge-Building

IT support has the potential of transforming the activities underlying the synergistic moment. For one thing, it would make those activities publicly accessible. The group could then reflect upon the emergence of its shared understanding by looking over the persistent record of its dialog. Such reflection might prove especially useful in contentious situations or for newcomers who were not part of the original dialog and are motivated to re-open the issue – as illustrated by Matusov (1996). Furthermore, computer support of perspectives could make explicit the interplay of different personal Perspectives and the migration of ideas and their interpretations between personal and group Perspectives. Ironically, perhaps, the “asynchronous” medium of the Web would allow group members to interact simultaneously – without waiting for sequential turns – thereby overcoming what Peters (1998) characterizes as “the hardest argument against democracy: the ability of only one person to speak and be heard at a time” (p. 261). Of course, as we have already discovered with the Web in general, the increased flood of ideas raises complex information management issues. We do not yet understand the full social impact of the envisioned KBEs – we will only know how they are used once they have been implemented, deployed in naturalistic settings, and observed.

Theories of human cognitive development emphasize the important role of external memories to extend short-term and long-term human memory (Donald, 1991; Norman, 1993). They also stress that individual cognition is a social product, highly mediated by social symbol systems, cultural artifacts, processes of structuration, and group collaboration (Bourdieu, 1972/1995; Geertz, 1973; Giddens, 1984; Hutchins, 1996; Vygotsky, 1930/1978). This suggests that computer support for collaboration has the potential to significantly advance the power of human cognition. In addition to maintaining a persistent external memory, IT can help people to be more reflective and creative – as has been demonstrated in computer support for brainstorming and decision-making (Connolly, 1997; Vogel et al., 1987). However, as our research to date indicates, despite the fact that the Web seems to offer a promising technological base for such a development, computer-supported collaboration is a complex process that requires a sophisticated body of knowledge that we are just beginning to assemble. Moreover, the potential is beyond the reach of any single research group.

We believe that IT support for collaborative knowledge-building has not yet been developed to near its potential. KBE research has been carried on for over a decade now, starting with the CSILE system and continuing with KIE, CoVis, etc. (Cuthbert, 1999; Pea, 1993; Scardamalia & Bereiter, 1991; 1996). Recently, commercial systems like KnowledgeForum, WebCT, and LearningSpace are catching on. However, as yet there has been no systematic attempt to support the variety of activities that are involved in knowledge-building. There is no general theory of collaborative knowledge-building as a social process. Existing research tends to target specific contexts like middle school science with specialized closed systems, rather than developing interchangeable, open source components that can be applied in a full range of contexts. While networks of KBE researchers are coming together in other countries, there is little organized effort to collaborate in the US. The proposed project aims to change this situation. However, collaboration across institutions cannot be started by just wishing for it. This project tries to put some of the necessary conditions in place by developing technical infrastructure (standards, a Perspectives Server, interface components) and initial results that can be used to stimulate discussion and collaboration among KBE researchers locally, nationally, and internationally. Interoperability and collaboration will allow isolated advances to be exchanged, new functionalities to be shared, and test data to be compared.

A Model of Collaborative Knowledge-Building

One approach to better understanding how to design computer support for collaborative knowledge-building in social settings is to conceptualize the various constituent activities involved in individual and social knowledge-building. The diagram below from (Stahl, 2000b) provides a starting point for this, combining aspects of activity theory, situated learning, hermeneutic philosophy, and distributed cognition theory (Chaiklin & Lave, 1993; Cole, 1996; Engeström et al., 1999; Gadamer, 1960/1988; Hutchins, 1996; Lave & Wenger, 1991; Nardi, 1996).

The idea of this diagram is that knowledge-building can proceed through many different activities. The sequential structure of the model is only illustrative of an ideal conceptualization. We understand that these activities complexly overlap in practice. The possible relationships among the individual activities – and particularly the interactions between the personal and social – can be complex and varied. The purpose of the diagram is to suggest a number of distinct activities that could be supported by a KBE with multiple functionality. The sequential labeling of these activities corresponds to proposed KBE components listed in Table 1 below, and it is not intended to imply a necessary order to the activities.

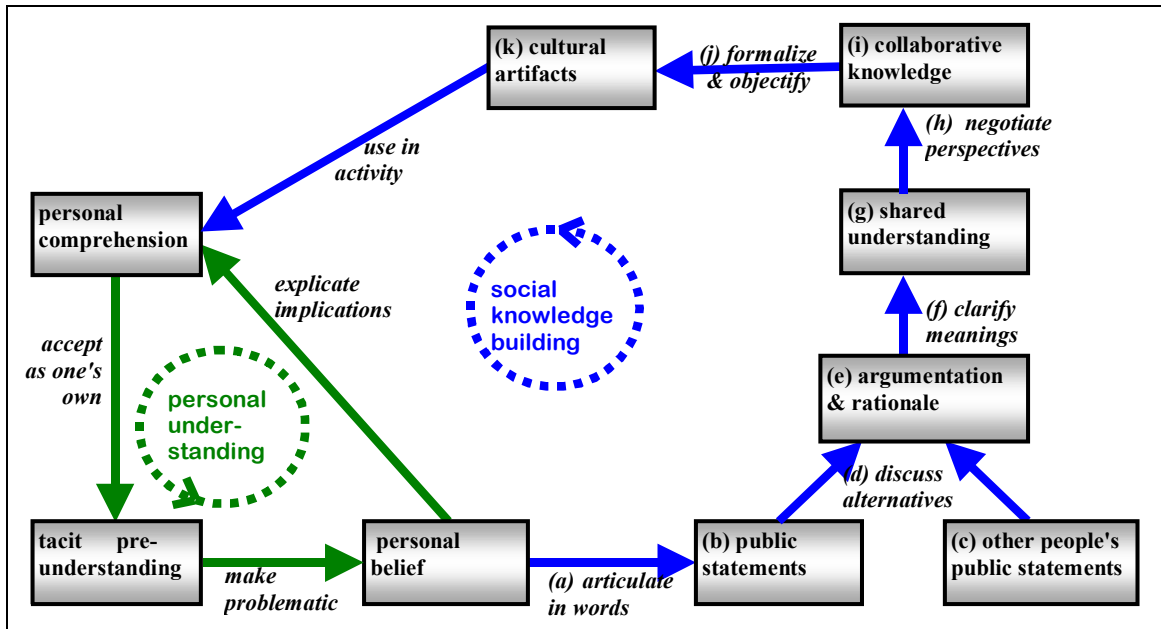


Figure 1. A model of personal understanding and social knowledge-building.

A set of seminal books and articles in Computer-Supported Collaborative Learning (CSCL) has formulated a view of *learning* as a social process of collaborative knowledge-building within communities of practice (Brown & Campione, 1994; Brown & Duguid, 1991; Lave, 1991; Lave & Wenger, 1991; Pea, 1993; Scardamalia & Bereiter, 1996; Wenger, 1998). However, these texts do not make the set of cognitive and social activities that underlie such a view explicit in the manner attempted in our KBE theory.

Starting in the lower left corner, Figure 1 shows a cycle of personal understanding. The rest of the diagram depicts how personal beliefs can be articulated in language and become part of social interaction. Note that the results of social knowledge-building eventually feed into personal understanding, providing the evolving toolkit of culturally-based individual cognitive capabilities. The depicted knowledge-building activities are discussed briefly below in the context of proposed computer support.

IT Support for Knowledge-Building Activities

Each of the activities of social knowledge-building pictured in Figure 1 can be supported computationally. Table 1 lists an illustrative form of support for each. It also lists corresponding prototypes that we have developed. Support for each activity is briefly discussed following the table.

Table 1. Forms of computer support for knowledge building activities.

	<i>Knowledge-building activities</i>	<i>Forms of computer support</i>	<i>Prototype systems</i>
a	articulate in words	articulation editor	DynaClass
b	public statements	personal Perspective	WebGuide
c	other people's public statements	comparison Perspective	WebGuide
d	discuss alternatives	discussion forum	DynaClass
e	argumentation & rationale	argumentation graph	InfoMap
f	clarify meanings	glossary discussion	DynaGloss
g	shared understanding	glossary	DynaGloss
h	negotiate perspectives	negotiation support	WebGuide

i	collaborative knowledge	group Perspective	WebGuide
j	formalize and objectify	bibliography discussion	Sources
k	cultural artifacts and representations	bibliography or other community repository	Sources

- (a) Computer support should facilitate the process of articulating ideas and preserving them in convenient forms. Most KBEs, including discussion forums like DynaClass, provide an *editor* for articulating ideas. Some KBEs have tried to introduce procedural facilitation, scaffolding, or prompting to encourage someone to articulate an appropriate expression (Slotta & Linn, 2000). Other approaches would be to provide an outline editor or a brainstorming area.
- (b) Public statements by one person confront those of other people. Computer support can represent the different *perspectives* from which these statements emerge. Perspectives are more general than representations of individuals themselves, because one person can offer statements from multiple perspectives and several people can agree on a common perspective. Perspectives can be related to one another, for instance deriving from a common perspective that they share. Computational representations of perspectives in a KBE like WebGuide (Stahl, 2000a) make explicit the important relationships among personal and group perspectives, as well as providing means for individuals and collaborative teams to articulate their own perspectives.
- (c) A KBE with support for Perspectives should provide *comparison Perspectives*, in which one can view and contrast alternative Perspectives and adopt or adapt ideas from other people's Perspectives. Comparison Perspectives in WebGuide aggregate ideas from various individual and/or group Perspectives and allow for comparison of them (Boland & Tenkasi, 1995; Stahl, 1999c). Other systems like D³E (Sumner & Buckingham Shum, 1998b) facilitate commentary on documents by other people, such as reviews of journal articles.
- (d) The most common element in current KBEs is the *discussion forum*. This is an asynchronous, interactive communication system like DynaClass that allows people to respond to notes posted by one another. Typically, there is a thread of responses to entered notes, with a tree of divergent opinions. A KBE should go beyond superficial undirected discussion to converge on shared understandings (dePaula, 1998; Guzdial & Turns, 2000; Hewitt & Teplovs, 1999).
- (e) Although every note in a discussion forum is a response to another note, the discussion may have a more complex implicit structure. One note might argue for or against another or provide evidence to back up the claim of another note, for instance. Such an argumentation structure can be made explicit and formalized in a representation of the *argumentation graph*. A component like InFoMap that displays the structure of notes graphically can contribute to participants' meta-level comprehension of their knowledge-building activity, pointing out where additional evidence is needed or where alternatives have not been explored (Buckingham Shum & Hammond, 1994; Donath et al., 1999; Suthers, 1999).
- (f) An important requirement for constructing group knowledge is the establishment of shared understanding. This can be fostered by clarifying the meaning of important terms used in various competing claims. A *glossary discussion* can make explicit how different participants understand the terms they use, as in DynaGloss or DocReview (Hendricksen, 1999).
- (g) The glossary discussion should result in a group *glossary* of the agreed upon definitions of important terms. Such a glossary already represents a form of group knowledge. The glossary is, of course, subject to future debate and emendation; it may make sense to define the glossary as a particular display of information from the glossary discussion (Stahl & Herrmann, 1999).
- (h) Perhaps the most delicate phase of knowledge-building is negotiation. Computer *support of negotiation* tends by nature to make explicit the factors entering into the negotiation process. This can be extremely harmful to the subtle processes of persuasion if not done sensitively. On the other hand, negotiation is critical to helping multiple perspectives to converge on shared knowledge. Computer support can provide a useful tool – as long as it is carefully integrated with other social activities that allow for implicit, culturally established interpersonal interactions (Stahl, 1999b). Group Decision Support Systems (GDSSs) have traditionally been independent systems, not integrated with the broader context of knowledge-building (Kraemer & Pinsonneault, 1990; Vogel et al., 1987).

- (i) The accumulation of negotiated shared knowledge results in the establishment of a *group perspective*. Like the alternative individual and team (or subgroup) perspectives, the group perspective may be represented in a KBE. In *WebGuide*, the content of the group Perspective is inherited into the individual and team Perspectives, because it has been accepted by the group. Individuals can then build on this shared knowledge within their own Perspective and even begin to critique it and start the whole cycle over (Stahl, 1999a).
- (j) Shared knowledge can be further formalized. It can be represented in another symbolic system or combined into a more comprehensive system of knowledge (Stahl, 1999c). For instance, in academic research knowledge is incorporated in new classroom lectures, conference presentations, journal articles, and books. The discussion of knowledge that has been compiled into publications can be carried out in a *bibliography discussion* component of a KBE.
- (k) Finally, representations of the new shared knowledge in publications and other cultural artifacts are themselves accepted as part of the established paradigm. Although still subject to occasional criticism, ideas in this form more generally provide part of the accepted base for building future knowledge. In academic circles, an annotated *bibliography* like *Sources* might provide a useful KBE component to support this knowledge building activity (Sumner & Buckingham Shum, 1998a).

A KBE goes beyond a single-purpose system – like a simple discussion forum – and supports more than one collaborative knowledge-building activity (Muukkonen et al., 1999). It retains a record of the knowledge that was incrementally collected – unlike common chat, newsgroup, and listserv systems that erase contributions after a short period of time. We hypothesize that it should help people to express their beliefs, to discuss them with others, to differentiate their own perspectives and adopt those of other people, clarify disagreements or misunderstandings, critique and explicate claims, negotiate shared understandings or agreements, and formulate knowledge in a lasting representation. Because KBEs are computational, they can provide facilities like searching, browsing, filtering, tailoring, and linking in order to group related ideas together automatically. KBEs can interface with other agents and software utilities – for instance sending emails to notify collaborators when important knowledge-building events have taken place (McLean, 1999). They can also dynamically format sets of notes in convenient displays for different purposes.

2.2. An Infrastructure for KBEs

Computational Perspectives

Computational Perspectives provide a new, dynamic, personalized form of on-line information management (Stahl, 1995). A Perspective defines an electronic workspace in which a person or group can develop ideas and manage information that belongs together – for instance because it represents the beliefs and viewpoint of a particular person, group, domain, or intellectual position. Perspectives structure a shared information space so that special coherent views can be built up and displayed. Although the mechanism of computational Perspectives is very general and flexible, the simplest way to use it in a small group is to define a personal Perspective for each member, one team Perspective for agreed upon ideas, and a comparison Perspective that collects the ideas from all the personal Perspectives.

The design philosophy behind computational Perspectives as implemented by the PI in *WebGuide* is that users have complete control over the content in their personal Perspectives. Thus, if my personal Perspective inherits conflicting ideas from different team Perspectives that I belong to, I can delete, edit, and rearrange those ideas at will. Other users can view the contents of my personal Perspective (except for content that I have designated as private) and they can copy items, link to them, initiate public discussions of them, and propose them for incorporation in team Perspectives – but none of this affects how the content of my Perspective is displayed to me. This allows me to build my own Perspective on the topics that are under consideration by the group. I can see what knowledge others are building, incorporate that knowledge into my Perspective, or join in with others to share, discuss, and negotiate. The same design philosophy applies of course to team Perspectives: team members jointly (through negotiation processes) have complete control over the content of their team Perspective.

Inheritance is a central defining mechanism of computational Perspectives as used in this proposal. The ability to define arbitrarily complex networks of Perspectives with multiple layers of sub-groups between the group Perspective and the individual personal Perspectives, and to have the automatic inheritance of content through the network distinguishes this approach from all other systems of “views” and “perspectives.” Inheritance in this sense

is not class inheritance, but “content inheritance.” A given Perspective can inherit content from multiple other Perspectives. This content is aggregated (logical union) in the given Perspective, where it can be over-ridden with edits, deletions, rearrangements, and additions. The inheritance mechanism is derived from efficient approaches explored in hypermedia, including “delta memory” and “transclusion” (Bobrow & Goldstein, 1980; McCall et al., 1990; Mittal et al., 1986; Nelson, 1981; Nelson, 1995). For a discussion of related work, see (Stahl & Herrmann, 1999).

Because new Perspectives can be defined (either in advance or during system use) to inherit from any (non-cyclical) other Perspectives, it is generally useful to define “comparison Perspectives” that aggregate the ideas from team members, including those ideas that have not been agreed upon and migrated to the team Perspective. This is handy for keeping an eye on what one’s fellow team members are thinking. Typically, we have set up the inheritance network of Perspectives to have a diamond-shaped profile, diverging out from the total group Perspective via teams to all the personal Perspectives, and then converging back via team comparisons to the group comparison Perspective. This models a collaborative knowledge-building process that combines divergent brainstorming and convergent negotiation.

Functionally considered, a KBE with Perspectives like WebGuide consists of two primary subsystems: a Perspectival data selection computation and a set of interface displays of the selected data. When a display is requested, the system must search the database to determine which content notes should be displayed to the particular user in the requested Perspective. For instance, if I request to view your Perspective, the system must select notes that are defined within your Perspective or within any Perspective from which yours inherits (recursively), except for notes that are private or that have been over-ridden. Various special displays can also be computed using this inheritance computation by treating discussions, negotiations, historical archives, etc. as pseudo-Perspectives that have special inheritance and exclusion rules. Once the Perspectival data computation has been returned, the content can be displayed in specialized interfaces that provide different kinds of functionality useful for further knowledge-building.

An Open Source Perspectives Server

A specific task of the proposed project is to separate out the Perspective computation from WebGuide and structure it as a self-contained module with a well-defined application programming interface (API). This will form a Perspectives Server, a Java application that runs on the Web server along with the database system. It will be separate from the WebGuide client that will still run in a Web browser on the client’s computer. This separation of functions into a server and a client will have many advantages. It will speed the functioning of WebGuide because the intensive computation of Perspective content will be done on a central server that is faster than typical student computers. Also, calls to the database system will take place locally rather than across the Internet. In terms of system development, it will mean that developers can build systems that incorporate Perspectives without having to worry about the Perspective algorithms or the database calls. They will use an API that lets them request data that should be shown to a given user in a given Perspective. They can then just focus on how best to display this data in the interface.

The Perspectives Server will be a self-contained Java application. It will be released as open source with clear documentation on how to use it to get Perspective data for display. The data will be delivered as an XML text stream that can be used by any Web technology, such as HTML, Perl, or Java. The data will be human-readable, making it easy for programmers to see what data is being passed. Although it is anticipated that the Perspectives Server will generally be used as a black box, its open source availability will allow programmers to modify it if necessary, such as to incorporate improvements to the XML DTD or in response to changes in Web technology. However, the Perspectives Server will be designed to make expansions of the database schema easy to incorporate without changes to the source code. This will allow new data structures corresponding to new multimedia data types.

The Perspectives Server will be a form of middle-ware, operating between the database and the client software. It will instantiate a three-tier, model-controller-view architecture that defines independent layers for the data schema or model, the data computation or control, and the interface display or view. The database management system can be any standard relational SQL system like MySQL or Oracle. The middle layer can be the Perspectives Server or a stripped down version that does not compute Perspectives. And the interface can be any kind of applet, Web page, or Web application that conforms to the API standard.

A Component Architecture for KBEs

The release of the open source Perspectives Server will not only facilitate the rapid prototyping of Perspectives-based systems for this project to use in its study sites, it will also allow other researchers to incorporate computational Perspectives in their KBEs. We have already had requests for this from researchers in California and in Germany.

The PI of this proposal is involved in several efforts to promote collaboration among KBE researchers. Among these, he is the PI on a seed grant from the NSF-funded Center for Innovative Learning Technologies (CILT) to foster data interoperability among KBE systems. Work on this is currently producing a draft XML DTD (eXtensible Markup Language – Document Type Definition) to provide a common data format that KBE data can be imported to and exported from. Data in this format can be displayed using CSS and XSL. Tools we are now developing will allow such data from any KBE to be analyzed with standard measures, allowing for instance the volume and characteristics of discussion threads on different systems to be quantified and compared.

This DTD will also provide the format for data transfer between the Perspectives Server and interface clients. Clearly, interface clients will have to be designed to accept and make use of data received in this XML format. An important aspect of the proposed work will be to define a set of standards: the XML DTD for data interchange, the API for the Perspectives Server, and the ability of interface clients to call the Server and make use of the data.

The PI will be communicating with other KBE researchers nationally and internationally to solicit their concerns about these standards and to work toward a consensus and adoption of a set of such standards. These standards do not have to be formally approved by international standards bodies; informal agreement within a set of collaborating research groups is all that is needed for substantial practical benefits. Use of these standards will allow for rapid prototyping and customization of systems for various study sites both within the proposed project and by other researchers. It will also move us significantly toward a future in which KBE components from different research groups can be intermixed so that new systems can take advantage of functionality developed at different sites. Certainly, every effort will be made to incorporate related international standards, such as those for XML, XLINK, and metadata.

Intelligent Hyper-Linking of Related Perspectives

System development efforts in Year I will focus on release of the Perspectives Server and in Year II on prototypes for the study sites. In Year III, we will enhance the power gained by organizing notes into personal and group Perspectives by adding the functionality to locate the notes most closely related to a given note, such as a new idea just entered into my personal Perspective or an old note proposed for inclusion in the group Perspective. Research in KBEs like CSILE has shown that it is difficult to locate related ideas within a shared database of discussion notes (Hewitt et al., 1998; Hewitt & Teplovs, 1999).

We will use Latent Semantic Analysis (LSA) (Landauer & Dumais, 1997; Landauer et al., 1998) to analyze the semantic content of notes and to measure the semantic relatedness of pairs of notes. LSA is based upon a statistical analysis (singular value decomposition) of co-occurrences of terms in a large corpus of text. It determines the relatedness of words even if they did not occur together explicitly – hence the term “latent”. LSA incorporates some refinements that make its performance closer to that of humans than similar methods (see special issues of *Discourse Analysis* 1997 and *Interactive Learning Environments* 2000 on LSA assessment studies). The PI and his graduate assistant are currently completing a four year project (sponsored by the McDonnell Foundation CSEP Program) that successfully uses LSA in a Web-based educational system tested in middle school classrooms (Kintsch et al., 2000; Stahl et al., 2000; Steinhart, 2000).

Automated linking of related notes will involve a fairly straight-forward application of LSA. It will be handled within the Perspectives Server, running on a computer with access to the necessary files for LSA. For each major application or study site, a corpus of domain-related text will be subjected to LSA analysis to define a semantic space. Periodically (e.g., each night) the site’s shared database of notes will be folded into this corpus to redefine the space and to compute the vector for each note within this space. In real time, when linking is requested for a new note, the note’s vector can be quickly computed and a list of notes with the closest vectors in the semantic space can be produced without noticeable delay.

We will experiment with different interfaces to try alternative approaches to incorporating this functionality into KBEs. For instance, it can be left to users to ask for lists of notes related to a given note. Alternatively, an agent can automatically check to see if there are notes within a given closeness to certain notes: newly entered notes, notes

proposed for negotiation, notes being read or edited, etc. The agent can then suggest that links be established from the given note to similar ones. The different interfaces can be tried out in our study sites.

2.3. Study Sites for Evaluating Knowledge-Building in Workgroups

We will use local sites under our own control as alpha sites for testing our software, not only to eliminate bugs, but to try out different functionality and to refine the interface. More formal evaluation of the social impact of KBEs in workgroups will be conducted in corporate sites, primarily in the Boulder area. We have targeted two local situations (academic research and university learning) and two corporate sites (corporate training and industrial design).

Academic Research: the L³D Center

The proposed work will take place within the Center for LifeLong Learning and Design (L³D), a research group within the Department of Computer Science at the University of Colorado. We will try out our prototypes in a variety of applications within L³D. Such self-application will give us first-hand experience with the requirements for the use of KBE software and with the practical problems of deployment and adoption. Increasingly, research at L³D involves participants from different disciplines and even virtual subgroups, like colleagues at other universities and other countries. When, e.g., a research project involves participants from Boulder, Colorado, and from Dortmund, Germany, a Web-based collaboration medium is essential, and means for defining and negotiating personal, subgroup, and whole-group Perspectives seem particularly appropriate. Because users in this group are themselves software developers and researchers, they are particularly accepting of glitches and are reflective about design issues. This will provide a convenient and forgiving initial test site.

University Learning: College Seminars

Several members of L³D, including the PI, offer undergraduate classes and graduate seminars. These courses typically emphasize student on-line discussion and group projects. They often have a content focus on the Internet. We have used WebGuide and other Web-based tools in past courses and will continue to use them in the future. We are exploring courses that not only span multiple disciplines, but also span the oceans. Again, here, Perspectives for subgroups (interest-, content-, discipline-, or location-based) make sense. Collaborative classrooms will provide a secondary test site in which KBE functionality and prototypes can be tried out and knowledge-building activities can be monitored.

Corporate Training: Athenaeum International DesignShops

We foresee our primary study site for this project being a corporate training setting, in which knowledge-building takes place under settings that may be advantageous for study. Athenaeum International (AI) is part of a distributed network of corporate training facilitators associated with MG Taylor. AI is located in Boulder and has established a good working relationship with the PI and with L³D. AI specializes in the design and manufacture of custom movable furniture for rapid deployment at corporate training events, as well as the facilitation of such events. They are interested in incorporating knowledge-building software tools in their furniture to support the training process. (See attached letter of support from AI.)

A typical training event – or DesignShop™ – involves bringing together decision-makers from throughout a company to “reengineer” their corporation or re-think their high-level mission. This might involve a series of three-day workshops, or even an on-going sequence of quarterly gatherings. AI staff have noted a number of problems that they think could be addressed by innovative computer support: there is tremendous time pressure and everyone cannot express all their ideas and arguments; too much time is spent introducing materials; it is hard to retain important points and decisions; follow-through is tricky; documentation is labor-intensive. KBE support could include pre-workshop preparatory discussions, capturing of ideas that arise in the face-to-face meetings, organized documentation of debates that took place, and follow-up discussion, analysis, negotiation, decision-making, or follow-through.

AI would provide a challenging study site for KBEs. DesignShops are high-stakes events involving people who need to make efficient use of their time. The groups here would be larger than the test groups at the university, and would involve more intense face-to-face interaction. A workshop series is of limited duration, so success could be assessed quickly and changes made prior to a subsequent trial. Project staff would have the aid of AI’s experienced group

process facilitators to guide the design of the software support and of the deployment strategy, as well as to analyze the impact the software had on the social systems.

Industrial Design: Seagate Technical Design Centers

Another targeted corporate study site provides a rather different opportunity for investigating the use of KBEs. We have begun to investigate a particular work group within Seagate, a major hard disk manufacturer located near our university. Seagate is an established high-tech corporation. Through a series of acquisitions and mergers, it now consists of large design centers in Longmont (outside of Boulder, Colorado), Minneapolis (MN), Kansas City (IW), Redwood City (CA) and Singapore. Each of these design centers houses a few hundred employees working rather independently of the other centers.

We will focus on the effort of a Seagate Vice-President who is in charge of coordinating research on the problems of disk drive head tracking. In order to meet market demands that are projected as exceeding Moore's law – requiring a doubling of storage density every year – the designers who work on head tracking must solve complex issues in physics and mechanical engineering. Unfortunately, engineers at different centers work almost independently of each other, duplicating research and designing products that overlap in functionality and specifications. As a result of their different histories, the centers have very different cultures of work, interaction, and outlook. For instance, one center prides itself in minimizing costs while another spares little cost to make what they consider a quality product. Designers from the different centers are accustomed to different engineering paradigms and find it hard to talk to each other.

The research objective here would likely be to use KBEs to structure communication and collaborative design among the distributed design groups. There is already a shared Lotus Notes database in which everyone can view the specifications and deadlines for each group's product line. However, there is currently no medium of communication among the groups (other than generic email) and no persistent textual discussion of the posted data. This makes it hard to share interpretations, work on establishing common understandings, or collaborate on building knowledge from the data. There is no support for Perspectives that would represent the conflicting cultures of the distributed groups and allow for negotiation of these differences.

Other Potential Study Sites

We will try to work with both AI and Seagate initially to explore their suitability to this project. We may end up focusing on one or the other site for practical or theoretical reasons. It is always hard to predict how field research will develop over years, particularly in today's volatile marketplace, where key contacts change jobs and companies alter their strategies. For this reason, we have also established access to several other industrial sites similar to AI and Seagate: *StorageTech* is another major manufacturer of digital storage media; our contacts there are especially concerned with making their company a "learning organization" in which knowledge-building is recognized to be an essential aspect of work. *IBM* has a major support center near Boulder, with a help-desk organization of 700 employees who must continually collaborate to build knowledge of the products they support. *L³D* has maintained a long-term relationship with two Japanese software companies, *SRA* and *PFU*, who are both interested in incorporating our ideas about lifelong learning and KBEs into their operations. These companies all have interesting settings where we would be welcome to deploy and observe our system prototypes if we have time during the proposed project.

Evaluation of Social Impact of KBEs at Study Sites

Evaluation will be conducted using converging methods to understand the complex, systemic issues around new technology deployment and use. Evaluation goals are two-fold:

- *Constructive:* We want to understand the environment with the objective of tailoring the design of the KBE technology to the study site, as well as constructively guide the deployment toward successful adoption.
- *Objective:* We want to objectively observe evolving use of the KBE technology; assess the nature of the IT impacts on coordination and collaboration; and evaluate the validity of the "collaborative knowledge building" concept and applicability of the KBE theory; and refine the KBE theory based on results from observed practice.

Data collection will take the form of face-to-face in-depth interviews; phone and email-based "interviews" when subjects are at a distance (particularly in the case of Seagate); real-time non-participant observation in the workplace as well as via the KBE technology; surveys (particularly in the case of the AI training sessions, where there will be many more subjects); KBE database data collection; and document collection where appropriate.

In addition to coding and content analysis of field and interview notes, analytical approaches will also include structural analysis of discussion thread lengths and participation levels based on the KBE database data. We will perform discourse and content analysis of argumentative exchanges, and KBE-captured collaboration episodes.

Issues we will investigate include:

- What is the nature of collaborative knowledge-building, and what activities comprise it?
- How does asynchronous support for articulation affect participation with respect to a variety of factors including time, location, and social status?
- How does computational support for discussion, argumentation, and clarification affect consensus-building as well as conflict?
- Do subjects understand and put into practice the concept of computational Perspectives? How do subjects interact with and manage multiple Perspectives (personal, subgroup, group, comparison)?
- Can "synergistic moments" be captured? What portion of these moments occur on-line in KBE environments?
- What do subjects do with persistent dialog? Does persistence foster reflection on "synergistic moments," and under what circumstances (what shades of consensus and conflict)?
- What role do facilitators (in the case of Athenaeum International) and management play in the use of KBE technologies?
- What are the particular hurdles that must be overcome for KBEs to be useful in these particular organizational environments?

Since the evaluation is a multiple-person effort, and because there will likely be different students participating at different times during the pre- and post-deployment stage, we will carefully organize observation efforts, and make an effort to systemize field note format to the best possible extent. We will also conduct regular group data analysis meetings to coordinate the results of our efforts.

3. OBJECTIVES OF PROPOSED WORK

3.1. Objectives for Theory Development

- Investigate the phenomenon of the "synergistic moment" in interactions captured in the database.
- Investigate the utility and actual usage of computational Perspectives.
- Refine the model of collaborative knowledge-building activities based on project findings.
- Publish results of this project in a monograph.

3.2. Objectives for Technology Prototyping

- Define a standard for data interoperability among KBEs.
- Release open source import/export/display/analysis tools for KBE data interoperability.
- Release an open source Perspectives Server.
- Release example KBE interface components using different technologies.
- Develop an agent for intelligent hyper-linking of KBE notes.

3.3. Objectives for Deployment

- Deploy a KBE prototype in settings of academic research (e.g., L³D research group).
- Deploy a KBE prototype in settings of university learning (e.g., seminars at the university).
- Deploy a KBE prototype in settings of corporate training (e.g., AI DesignShops).
- Deploy a KBE prototype in settings of industrial design (e.g., Seagate design groups).

3.4. Objectives for Evaluation

- Conduct initial evaluation of the way work is presently conducted among selected groups to: collect baseline data; carefully select groups to deploy to; and formulate deployment strategies.
- Assess existing challenges for communication and collaboration within and across groups through interviews and possibly email-based interview-type surveys for those participants at a distance.
- Observe deployment of KBEs and revise deployment strategies as necessary.
- Perform on-going qualitative (observation, interviews, and document collection) and quantitative (KBE database activity) data collection.
- Analyze collected data to evaluate the impact of the KBEs in the workgroups.

3.5. Objectives for Dissemination

- Build a local, interdisciplinary community of students and faculty to conduct KBE research.
- Collaborate with at least 3 US research groups and with at least 3 international KBE research networks.
- Maintain an active website with the results of this project, including the open source products.
- Present the work of this project in the Group '01 and CSCL '01 international conferences (to be held at the University of Colorado), and report the findings of this project at 3 or more other international conferences.

4. PLAN OF PROPOSED WORK

4.1. Year I

Refine the model of collaborative knowledge-building activities through presentations to researchers and analysis of the "synergistic moment" in face-to-face collaboration.

Define a standard for data interoperability among KBEs using XML and XLINK for notes and relations among notes, including threaded discussions; explore the adequacy of this standard using local prototypes; circulate the proposed standard among other KBE researchers; integrate the proposed standard with emerging data standards.

Define a standard for KBE interface components including JavaBeans, using XML for data interchange; explore the adequacy of this standard using local prototypes; circulate the proposed standard among other KBE researchers; integrate the proposed standard with other emerging data standards.

Define a standard for KBE database servers to provide XML data structures to interface components meeting the above standards.

Release an open source Perspectives Server that efficiently computes data visible in a requested Perspective in the standard XML format. Provide open source and documentation so that other researchers can use this Server for Perspectives-based KBEs and can propose improvements to the Server.

Deploy a KBE prototype in a setting of academic research such as a local research group. The KBE will be used to discuss and design standards and software. Learn from this deployment experience.

Begin initial observation of two primary test sites (AI and Seagate) to prepare for eventual KBE deployment.

Collaborate with American research groups interested in KBEs (e.g., Berkeley, SIU, Stanford, Georgia Tech, Hawaii, SRI) and international KBE research networks (e.g., existing CSCL research networks in Germany, England, Norway, Finland, Canada) to coordinate software component design around agreed upon standards.

Establish a website to publicize the results of this project, including the open source products, and to provide an information center for collaborations. Incorporate a KBE in the website to foster interactive knowledge-building.

4.2. Year II

Release open source import/export/display/analysis tools for KBE data interoperability, using XML to define a DTD and using XSL and CSS to display the data. Document the standard and the tools on a website that makes them available to other researchers.

Release example KBE interface components in different technologies, including Java, Perl, and HTML. Make these available on a website with documentation and tutorials to help other researchers develop compatible components.

Deploy a KBE prototype in a setting of university learning such as a seminar offered by the PI, and observe use. The KBE will be used for students to develop, exchange, and negotiate reflections on shared readings and on collaborative writing projects.

Continue investigating the AI study site and assess the requirements for deployment of a KBE there.

Deploy a KBE prototype in a setting of corporate training such as a series of AI corporate training sessions. Commence post-deployment observation.

Present approach and intermediate findings of this project at the Group '01 and CSCL '01 international conferences. During Year II, these conferences will both be hosted at CU, with the PI serving as chair of the local arrangements committee at Group and chair of the program committee at CSCL.

4.3. Year III

Deploy a KBE prototype in a setting of industrial working such as a technical design group at Seagate, and observe and analyze use there. The AI deployment will also continue to be evaluated in Year III.

Develop an agent for intelligent hyper-linking of KBE notes using LSA technology to locate notes in a KBE database that are semantically related to a given note.

Continue on-going analyses of KBE use and examine findings with respect to the project's central hypotheses about KBEs. Publish a monograph on the project findings.

Report on this research at 3 or more international conferences.

5. EXPECTED IMPACT OF PROPOSED WORK

The proposed work – with its development of theory, defining of architectures and standards, prototyping of systems, deployment at several study sites, evaluation of utility, and dissemination to collaborators around the world – may seem overly ambitious for a half-time (6 months per year for the PI) project. This would certainly be the case if a lot of the groundwork had not already been laid and if there was not an expectation that this project will take place within a context of supporting work. The PI has already published on KBE Theory and begun to sketch more in his seminars. Work is already well underway on separating the Perspectives Server out of WebGuide and defining an appropriate XML DTD for the data format – although these two projects have not yet been combined. Building a consensus on technical standards for KBEs has been initiated through the CILT grant and a popular workshop at CSCL '99. It will continue at planned conferences, including Group '01 and CSCL '01 where the PI will be involved in planning. WebGuide and other L³D prototypes have been used and evaluated in middle school and college courses, including the PI's recent seminars. Colleagues in L³D are developing systematic methodologies for analyzing and comparing discussions in KBEs.

The proposed project is conceived as the centerpiece for the PI's work in the coming years. A large NSF ITR pre-proposal is currently pending that would support the collaboration efforts associated with this work. Another proposal to NSF's CSS program would continue our OMOL project, orienting it particularly toward the deployment and adoption effort at the corporate study sites and at the analysis of the social impact of the KBE software in these settings, while a pending CISE Research Infrastructure proposal includes equipment and infrastructure for the local academic study sites. If funded, these other proposals would provide adequate support to fully achieve the vision proposed here. If they are not funded, the PI will seek other grants leveraged on the work proposed here.

The proposed work is designed to provide some of the basic conditions for the sorts of intensive collaboration that we believe are necessary to achieve the potential of KBE technology. In addition to providing technical conditions (e.g., Server, architecture, standards) and social conditions (local, national, international networks), the work will prepare specific examples of functionality (e.g., computational Perspectives, automated linking) and concrete analyses of application. These results will be disseminated through communication channels including specific related conferences.

Within the PI's home institutions the proposed work will be an important catalyst for building an interdisciplinary collaborative effort of undergraduates, graduate students, research staff, and faculty working on KBE theory, technology, and evaluation. The proposed work defines a coherent research agenda over several years that will crystallize a local research community.

“TTR/IM: Perspectives on Collaborative Knowledge-Building”

References Cited

- Bakhtin, M. (1986) *Speech Genres and Other Late Essays*, (V. McGee, Trans.), University of Texas Press, Austin, TX.
- Bobrow, D. & Goldstein, I. (1980) *An experimental description-based programming environment: Four reports*, Technical Report No. CSL-81-3, Xerox Palo Alto Research Center, Palo Alto, CA.
- Boland, R. J. & Tenkasi, R. V. (1995) Perspective making and perspective taking in communities of knowing, *Organization Science*, 6 (4), pp. 350-372.
- Bourdieu, P. (1972/1995) *Outline of a Theory of Practice*, (R. Nice, Trans.), Cambridge University Press, Cambridge, UK.
- Brown, A. & Campione, J. (1994) Guided discovery in a community of learners. In K. McGilly (Ed.) *Classroom Lessons: Integrating Cognitive Theory and Classroom Practice*, MIT Press, Cambridge, MA, pp. 229-270.
- Brown, J. S. & Duguid, P. (1991) Organizational learning and communities-of-practice: Toward a unified view of working, learning, and innovation, *Organization Science*, 2 (1), pp. 40-57.
- Buckingham Shum, S. & Hammond, N. (1994) Argumentation-based design rationale: What use at what cost?, *International Journal of Human-Computer Studies*, 40 (4), pp. 603-652.
- Chaiklin, S. & Lave, J. (Eds.) (1993) *Understanding Practice: Perspectives on Activity and Context*, Cambridge University Press, New York, NY.
- Cole, M. (1996) *Cultural Psychology*, Harvard University Press, Cambridge, MA.
- Connolly, T. (1997) Electronic brainstorming: Science meets technology in the group meeting room. In S. Kiesler (Ed.) *Culture of the Internet*, Lawrence Erlbaum Associates, Hillsdale, NJ, pp. 263-276.
- Cuthbert, A. (1999) Designs for collaborative learning environments: Can specialization encourage knowledge integration?, In: Proceedings of *Computer Supported Collaborative Learning (CSCL '99)*, Palo Alto, CA, pp. 117-126.
- dePaula, R. (1998) *Computer Support for Collaborative Learning: Understanding Practices and Technology Adoption*, Masters Thesis, Telecommunications Department, University of Colorado, Boulder, CO.
- Donald, M. (1991) *Origins of the Modern Mind: Three Stages in the Evolution of Culture and Cognition*, Harvard University Press, Cambridge, MA.
- Donath, J., Karahalios, K., & Viegas, F. (1999) Visualizing conversation, *Journal of Computer Mediated Communication*, 4 (4). Available at: <http://www.ascusc.org/jcmc/vol4/issue4/donath.html>.
- Engeström, Y., Miettinen, R., & Punamaki, R.-L. (Eds.) (1999) *Perspectives on Activity Theory*, Cambridge University Press, New York, NY.
- Feltovich, P., Spiro, R., Coulson, R., & Feltovich, J. (1996) Collaboration within and among minds: Mastering complexity, individually and in groups. In T. Koschmann (Ed.) *CSCL: Theory and Practice of an Emerging Paradigm*, Lawrence Erlbaum Associates, Hillsdale, NJ, pp. 25-44.
- Fischer, G. (1994) Domain-oriented design environments, including Reply to commentaries, *Automated Software Engineering*, 1 (2), pp. 177-203, 223-229.
- Fischer, G., Grudin, J., McCall, R., Ostwald, J., Redmiles, D., Reeves, B., & Shipman, F. (1999) Seeding, evolutionary growth and reseeding: The incremental development of collaborative design environments. In G. M. Olson, T. W. Malone, & J. B. Smith (Eds.), *Coordination Theory and Collaboration Technology*, Lawrence Erlbaum Associates, Mahwah, NJ, p. in press.
- Fischer, G., Lindstaedt, S., Ostwald, J., Schneider, K., & Smith, J. (1996) Informing system design through organizational learning, In: Proceedings of *International Conference on The Learning Sciences (ICLS '96)*, Evanston, IL, pp. 52-59. Available at: <http://www.cs.colorado.edu/~gerhard/papers/InformingSystemDesign.pdf>.
- Fischer, G., Nakakoji, K., Ostwald, J., Stahl, G., & Sumner, T. (1993) Embedding computer-based critics in the contexts of design, In: Proceedings of *Conference on Human Factors in Computing Systems (INTERCHI '93)*,

- Amsterdam, Holland, pp. 157-164. Available at:
<http://www.cs.colorado.edu/~gerry/publications/conferences/1990-1997/chi93/CHI93.html>.
- Fischer, G., Nakakoji, K., Ostwald, J., Stahl, G., & Sumner, T. (1993/1998) Embedding critics in design environments. Reprint from 1993 (*Knowledge Engineering Review* (4) 8, 285-307); including Reflections from 1998. In M. T. Maybury & W. Wahlster (Eds.), *Readings in Intelligent User Interfaces*, Morgan Kaufman, New York, pp. 537-561. Available at: <http://www.cs.colorado.edu/~gerry/publications/journals/ker/index.html>.
- Gadamer, H.-G. (1960/1988) *Truth and Method*, Crossroads, New York, NY.
- Geertz, C. (1973) *The Interpretation of Cultures*, Basic Books, New York, NY.
- Giddens, A. (1984) *The Constitution of Society*, U of California Press, Berkeley, CA.
- Guzdial, M. & Turns, J. (2000) Sustaining discussion through a computer-mediated anchored discussion forum, *Journal of the Learning Sciences* .
- Hatano, G. & Inagaki, K. (1991) Sharing cognition through collective comprehension activity. In L. Resnick, J. Levine, & S. Teasley (Eds.), *Perspectives on Socially-shared Cognition*, APA, Washington, DC, pp. 331-348.
- Heidegger, M. (1927/1996) *Being and Time: A Translation of Sein und Zeit*, (J. Stambaugh, Trans.), SUNY Press, Albany, NY.
- Hendricksen, C. (1999) Asynchronous collaboration by Document Review, *DocReview* . Available at: <http://students.washington.edu/veritas/papers/DRpaper/basedoc.html>.
- Hewitt, J., Scardamalia, M., & Webb, J. (1998) *Situative design issues for interactive learning environments*, at http://csile.oise.on.ca/abstracts/situ_design.
- Hewitt, J. & Teplov, C. (1999) An analysis of growth patterns in computer conferencing threads, In: *Proceedings of Computer Supported Collaborative Learning (CSCL '99)*, Palo Alto, CA, pp. 232-241.
- Hutchins, E. (1996) *Cognition in the World*, MIT Press, Cambridge, MA.
- Jordan, B. & Henderson, A. (1995) Interaction analysis: Foundations and practice, *Journal of the Learning Sciences*, 4 (1), pp. 39-103.
- Kintsch, E., Steinhart, D., Stahl, G., Matthews, C., Lamb, R., & the LSA Research Group (2000) Developing summarization skills through the use of LSA-backed feedback, *Interactive Learning Environments* .
- Kling, R. (1999) What is social informatics and why does it matter?, *D-Lib Magazine*, 5 (1).
- Kraemer, K. & Pinsonneault, A. (1990) Technology and groups: Assessment of the empirical research. In J. Galegher, R. Kraut, & C. Egido (Eds.), *Intellectual Teamwork: Social and Technological Foundations of Cooperative Work*, Lawrence Erlbaum Associates, Hillsdale, NJ, pp. 375-405.
- Landauer, T. K. & Dumais, S. T. (1997) A solution to Plato's problem: The latent semantic analysis theory of acquisition, induction and representation of knowledge, *Psychological Review*, 104 (2), pp. 211-240.
- Landauer, T. K., Foltz, P. W., & Laham, D. (1998) Introduction to latent semantic analysis, *Discourse Processes*, 25 (1), pp. 259-284.
- Lave, J. (1991) Situating learning in communities of practice. In L. Resnick, J. Levine, & S. Teasley (Eds.), *Perspectives on Socially Shared Cognition*, APA, Washington, DC, pp. 63-83.
- Lave, J. & Wenger, E. (1991) *Situated Learning: Legitimate Peripheral Participation*, Cambridge University Press, Cambridge, UK.
- Matusov, E. (1996) Intersubjectivity without agreement, *Mind, Culture, and Activity*, 3 (1), pp. 25-45.
- McCall, R., Bennett, P., d'Oronzio, P., Ostwald, J., Shipman, F., & Wallace, N. (1990) PHIDIAS: Integrating CAD graphics into dynamic hypertext. In A. Rizk, N. Streitz, & J. Andre (Eds.), *Hypertext: Concepts, Systems and Applications. Proceedings of the European Conference on Hypertext (ECHT '90)*, Cambridge University Press, Cambridge, UK, pp. 152-165.
- McLean, R. (1999) Meta-communication widgets for knowledge building in distance education, In: *Proceedings of Computer Supported Collaborative Learning (CSCL '99)*, Palo Alto, CA, pp. 383-390.
- Mittal, S., Bobrow, D., & Kahn, K. (1986) Virtual copies at the boundary between classes and instances, In: *Proceedings of Object-Oriented Programming Systems, Languages and Applications (OOPSLA '86)*, Portland, OR, pp. 159-166.

- Muukkonen, H., Hakkarainen, K., & Lakkala, M. (1999) Collaborative technology for facilitating progressive inquiry: Future Learning Environment tools, In: Proceedings of *Computer Supported Collaborative Learning (CSCL '99)*, Palo Alto, CA, pp. 406-415.
- Nardi, B. (Ed.) (1996) *Context and Consciousness: Activity Theory and Human-Computer Interaction*, MIT Press, Cambridge, MA.
- Nelson, T. (1981) *Literary Machines*, Mindful Press, New York, NY.
- Nelson, T. H. (1995) The heart of connection: Hypermedia unified by transclusion, *Communications of the ACM*, 38 (8), pp. 3-33.
- Norman, D. A. (1993) *Things That Make Us Smart*, Addison-Wesley Publishing Company, Reading, MA.
- Nygaard, K. & Sørgaard, P. (1987) The perspective concept in informatics. In G. Bjerknes, P. Ehn, & M. Kyng (Eds.), *Computers and Democracy: A Scandinavian Challenge*, Avebury, Aldershot, UK, pp. 371-393.
- Ostwald, J. (1999) *Dynasites*, at <http://seed.cs.colorado.edu/dynasites.documentation.fcgi>.
- Papert, S. (1980) *Mindstorms: Children, Computers and Powerful Ideas*, Basic Books, New York, NY.
- Pea, R. (1993) The collaborative visualization project, *Communications of the ACM*, 36 (5), pp. 60-63.
- Perret-Clermont, A.-N. & Schubauer-Leoni, M.-L. (1981) Conflict and cooperation as opportunities for learning. In W. P. Robinson (Ed.) *Communication in Development*, Academic Press, New York, NY, pp. 203-234.
- Peters, J. D. (1998) *Speaking Into the Air*, University of Chicago Press, Chicago, IL.
- Roschelle, J. (1998) Activity Theory: A Foundation for Designing Learning Technology?, *Journal of the Learning Sciences*, 7 (2), pp. 241-255.
- Scardamalia, M. & Bereiter, C. (1991) Higher levels of agency in knowledge building: A challenge for the design of new knowledge media, *Journal of the Learning Sciences*, 1 , pp. 37-68.
- Scardamalia, M. & Bereiter, C. (1996) Computer support for knowledge-building communities. In T. Koschmann (Ed.) *CSCL: Theory and Practice of an Emerging Paradigm*, Lawrence Erlbaum Associates, Hillsdale, NJ, pp. 249-268.
- Slotta, J. D. & Linn, M. C. (2000) The Knowledge Integration Environment: Helping students use the Internet effectively. In M. J. Jacobson & R. Kozma (Eds.), *Learning the Sciences of the 21st Century*, Lawrence Erlbaum Associates, Hillsdale, NJ.
- Stahl, G. (1975) *Marxian Hermeneutics and Heideggerian Social Theory: Interpreting and Transforming Our World*, Ph.D. Dissertation, Department of Philosophy, Northwestern University, Evanston, IL. Available at: <http://www.cs.colorado.edu/~gerry/publications/dissertations/thesis.htm>.
- Stahl, G. (1993a) *Interpretation in Design: The Problem of Tacit and Explicit Understanding in Computer Support of Cooperative Design*, Ph.D. Dissertation, Department of Computer Science, University of Colorado, Boulder, CO. Available at: http://www.cs.colorado.edu/~gerry/publications/dissertations/dis_intro.html.
- Stahl, G. (1993b) Supporting situated interpretation, In: Proceedings of *Annual Meeting of the Cognitive Science Society (CogSci '93)*, Boulder, CO, pp. 965-970. Available at: <http://www.cs.colorado.edu/~gerry/publications/conferences/1990-1997/cogsci93/CogSci.html>.
- Stahl, G. (1995) *Supporting Personalizable Learning*, Technical Report No. CU-CS-788-95, Department of Computer Science, University of Colorado, Boulder, CO. Available at: <http://www.cs.colorado.edu/~gerry/publications/techreports/personalize/>.
- Stahl, G. (1998) Collaborative information environments for innovative communities of practice, In: Proceedings of *German Computer-Supported Cooperative Work conference (DCSCW '98)*, Dortmund, Germany. Available at: <http://www.cs.colorado.edu/~gerry/publications/conferences/1998/dcscw98/dcscw.html>.
- Stahl, G. (1999a) POW! Perspectives on the Web, In: Proceedings of *WebNet World Conference on the WWW and Internet (WebNet '99)*, Honolulu, HA. Available at: <http://www.cs.colorado.edu/~gerry/publications/conferences/1999/webnet99/webnet99.html>.
- Stahl, G. (1999b) Reflections on WebGuide: Seven issues for the next generation of collaborative knowledge-building environments, In: Proceedings of *Computer Supported Collaborative Learning (CSCL '99)*, Palo Alto, CA, pp. 600-610. Available at: <http://www.cs.colorado.edu/~gerry/publications/conferences/1999/csl99/>.

- Stahl, G. (1999c) WebGuide: Guiding collaborative learning on the Web with perspectives, In: Proceedings of *Annual Conference of the American Educational Research Association (AERA '99)*, Montreal, Canada. Available at: <http://www.cs.colorado.edu/~gerry/publications/conferences/1999/aera99/>.
- Stahl, G. (2000a) Collaborative information environments to support knowledge construction by communities, *AI & Society*, 14 , pp. 1-27. Available at: <http://www.cs.colorado.edu/~gerry/publications/journals/ai&society/>.
- Stahl, G. (2000b) A model of collaborative knowledge-building, In: Proceedings of *International Conference of the Learning Sciences (ICLS 2000)*, Ann Arbor, MI. Available at: <http://www.cs.colorado.edu/~gerry/publications/conferences/2000/icls/>.
- Stahl, G. & dePaula, R. (1998) poster: Learning perspectives, In: Proceedings of *International Conference on the Learning Sciences (ICLS'98)*, Atlanta, GA. Available at: <http://www.cs.colorado.edu/~gerry/publications/conferences/1998/icls98/icls98.html>.
- Stahl, G., dePaula, R., & the LSA Group (2000) Evolution of an LSA-based interactive environment for learning to write summaries (in preparation), *Journal of Interactive Media in Education (JIME)* . Available at: <http://www.cs.colorado.edu/~gerry/publications/journals/ILE/ile.html>.
- Stahl, G., Hermann, T., dePaula, R., & Loser, K.-U. (1998) demo: WebGuide: Guiding cooperative work on the Web with support for perspectives and negotiation, In: Proceedings of *Computer Supported Cooperative Work (CSCW '98)*, Seattle, WA. Available at: http://www.cs.colorado.edu/~gerry/publications/conferences/1998/cscw98/cscw_demo.html.
- Stahl, G. & Herrmann, T. (1999) Intertwining perspectives and negotiation, In: Proceedings of *International Conference on Supporting Group Work (Group '99)*, Phoenix, AZ. Available at: <http://www.cs.colorado.edu/~gerry/publications/conferences/1999/group99/>.
- Stahl, G., Sumner, T., & Owen, R. (1995) Share globally, adapt locally: Software to create and distribute student-centered curriculum, *Computers and Education. Special Issue on Education and the Internet*, 24 (3), pp. 237-246. Available at: <http://www.cs.colorado.edu/~gerry/publications/journals/c&e/>.
- Steinhart, D. (2000) *Designing, Testing, and Evaluating Summary Street, an LSA-Based Summarization Tool (in preparation)*, Ph.D. Dissertation, Department of Psychology, University of Colorado, Boulder, CO.
- Sumner, T. & Buckingham Shum, S. (1998a) From documents to discourse: Shifting conceptions of scholarly publishing, In: Proceedings of *Computer-Human Interaction (CHI 98)*, Los Angeles, CA, pp. 95-102.
- Sumner, T. & Buckingham Shum, S. (1998b) A toolkit for publishing Web discussion documents: Design principles and case studies, In: Proceedings of *Asia Pacific Computer Human Interaction (APCHI 98)*, Shonan Village Center, Japan, pp. 218-223.
- Suthers, D. (1999) Representational support for collaborative inquiry, In: Proceedings of *32nd Hawai'i International Conference on the System Sciences (HICSS-32)*, Maui, HA, pp. CD-ROM. Available at: <http://lilt.ics.hawaii.edu/lilt/papers/hicss99.pdf>.
- Vogel, D., Nunamaker, J., Applegate, L., & Konsynski, B. (1987) Group decision support systems: Determinants of success, In: Proceedings of *Decision Support Systems (DSS '87)*, pp. 118-128.
- Vygotsky, L. (1930/1978) *Mind in Society*, Harvard University Press, Cambridge, MA.
- Wenger, E. (1987) *Artificial Intelligence and Tutoring Systems*, Morgan Kaufman, Los Altos, CA.
- Wenger, E. (1998) *Communities of Practice*, Cambridge University Press, Boston, MA.
- Winograd, T. & Flores, F. (1986) *Understanding Computers and Cognition: A New Foundation of Design*, Addison-Wesley, Reading, MA.