CLARE: A Computer-Supported Collaborative Learning Environment Based on the Thematic Structure of Research and Learning Artifacts

A Ph.D. Dissertation Proposal

by

Dadong Wan

January, 1993

Department of Information and Computer Sciences
University of Hawaii at Manoa

Contents

List of Figures

List of Tables

Abstract

This research concerns the representation issue in collaborative learning environments. Our basic claim is that knowledge representation is not only fundamental to machine learning, as shown by AI researchers, but also essential to human learning, in particular, human metalearning; few existing learning support systems, however, provide representations which are intended to help the learner make sense of and organize the subject content of learning, integrate a wide range of classroom activities, (e.g., reading, reviewing, writing, discussion), and compare and contrast various viewpoints from individual learners. It is our primary purpose to construct an example instance of such a representation, and to show that useful computational manipulations can be performed on it, and that the combination of the representation and related computational services can actually lead to the improved learner's performance on selected collaborative learning tasks.

First, we propose a representational scheme, called RESRA, which characterizes the thematic structure of learning and research materials or artifacts, and which might be used to organize learning activities at the levels of summarization, evaluation, argumentation, integration, and construction. We develop a computer-based tool, called CLARE, that supports the use of RESRA and two of its own aggregates (i.e., threads and perspectives) for various learning tasks, e.g., conducting reviews, engaging in collaborative writing, participating in structured online discussions, And finally, we design two experiments that will empirically evaluate the effectiveness of CLARE and provide evidence for our research claims.

1 Introduction

While collaborative learning is common and widely encouraged in college classrooms, the level of computer-based support has been quite limited. Many existing learning support environments, such as computer-assisted instruction (CAI) and intelligent tutoring systems (ITS), are designed to facilitate individual learning. Computer-mediated communication (CMC), namely, electronic mail and bulletin-board systems (e.g. EIES), is perhaps the most commonly used collaborative learning tool [Hiltz, 1988]. They have been found quite effective in creating "virtual classrooms" by overcoming the "same-time, same-place" requirement of the face-to-face learning. However, such systems do not provide services specialized for specific learning tasks and processes. Other more sophisticated tools, e.g. Intermedia [Yankelovich *et al.*, 1988], attempt to integrate multiple applications (e.g., word processing, drawing, mail) into a single environment. In addition, they also provide hypermedia capabilities which allow easy linking and integration of various types of objects (e.g., text, graphics, video) into any application. Despite their functional versatility, these tools are still fairly general-purpose, and their usefulness is confined to mere online authoring, browsing, annotation, or information sharing, all of which have long been available as separate systems. What these systems fail to provide is task-specific support, such as helping the learner structure a research paper, prepare a presentation, or engage in a focused discussion. What is also missing is the ability to represent and highlight individual differences and similarities and the ability to leverage on them.

This research concerns the representation, or lack of same, that underlines existing collaborative learning systems. We claim that the lack of an expressive and usable representation for organizing the subject content of learning, integrating various classroom activities, (e.g., reading, writing, discussion, presentation) and comparing and contrasting various viewpoints from individual learners accounts for most of the above mentioned problems. First, we propose a representational scheme, called RESRA, which characterizes the thematic structure of learning and research materials or artifacts. We develop a computer-based tool, called CLARE, that supports the use of RESRA for various learning tasks, e.g., writing research reviews, engaging in online discussions, drafting research proposals. We also design two experiments which will allow us to empirically evaluate the effectiveness of CLARE and test our research claims.

This proposal is organized as following: Section 2 describes the problem, context, and our research claims. Section 3 elaborates the representational framework, i.e., RESRA. Section 4 depicts the main design features and the architecture of CLARE. Section 5 outlines the two experiments we plan to conduct to evaluate CLARE. Section 6 relates the current research to a broader context of existing work. And finally, Section 7 enlists the plan of actions for the current project.

2 The Problem: Context and Claims

2.1 Theoretical Background

An essential consideration of any successful computer-based support environment is the melding of appropriate theory with innovative approaches to tool development. Our work is driven by three distinct and yet related learning theories: *metalearning*, *meaningful learning* and *constructivist pedagogy*. Together, these theories underscore the inadequacies of existing tools and provide directions for new systems.

Metalearning The notion of *metalearning*, as defined in [Novak and Gowin, 1984], concerns the understanding of the nature and structure of knowledge and learning itself. The other side of the coin is *content learning*, or understanding of the content of a specific topic, such as how human brains work. In a typical classroom setting, the two types of learning are often interwined. Participants in research seminars, for instance, are expected to both understand the particular subject matter addressed in the seminar, e.g., artificial intelligence, and to learn how to collaborate, how to research literature, how to present and evaluate research artifacts, how to identify interesting problems and develop novel solutions, and so forth. The significance of this distinction are twofold: first, as described in Section 2.2, although content learning tools are improving, the support for metalearning is not, not to mention the support for the integration of the two. Second, metalearning has become increasingly important in today's world in part because of the accelerating rate at which knowledge is produced and disseminated. Students will find the subject content they learn in school to quickly become obsolete. In contrast, the metaknowledge they acquire will enable them in the long run to adapt and cope with the changing state of human knowledge.

Meaningful Learning The fundamental assumption of meaningful learning theory, also known as assimilation theory of cognitive learning, is that the single most important factor influencing learning is what the learner already knows, and that learning is evidenced by a change in the meaning of experience rather than a change in behavior, as held by behavioral psychologists [Ausubel, 1963; Aususel *et al.*, 1978]. The key question is how to help students to reflect on their experience and to construct new meanings. Novak and Gowin [Novak and Gowin, 1984] propose two metacognitive strategies: concept maps and Vee diagrams, both of which are tools intended to represent changes in the knowledge structure of students over time and help them "learn how to learn". Concept mapping is widely accepted in the educational community. Numerous studies have shown its effectiveness in facilitating meaningful learning [Cliburn, Jr., 1990; Novak, 1990; Roth and Roychoudhury, 1992]. The Vee diagram, however, is less widely known and used.

Constructivist Pedagogy Constructivism holds that knowledge in general and scientific knowledge in particular is socially constructed [Berger and Luckman, 1966; Knorr-Centina, 1981]. Such knowledge, instead of being the same for individuals, is "taken-to-be-shared" [Roth and Roychoudhury, 1992] with communities of knowers. To become a member of such a community, students need to engage in collaborative interactions and undergo learning situations which allow them to be enculturated into the discourse practice of a field. In order to form classroom communities which function like those of scientists and researchers, for example, students need to be given the opportunity to engage in authentic practices of scientists and researchers.

2.2 Problem Characterization

In an earlier phase of this research, we identified four problems in collaborative learning that could be addressed through computerized support [Wan and Johnson, 1992]:

- *Face-to-face barriers.* Face-to-face barriers can be either interpersonal or intercultural/interlingual. In a seminar setting, for example, discussions can be dominated by a few "strong personalities" or by the seminar leader(s). Individual contributions to the group discussion might be inhibited because some participants do not feel comfortable speaking openly in a group, or expressing verbal disagreement with other participants, especially the leader. In seminars composed of people from different cultural and linguistic backgrounds, intercultural gaps and differences in language fluency might prevent the minorities from full participation.

- *Same-place, same-time constraints.* In a conventional seminar, again, physical copresence in classroom is a prerequisite for participation. Interactions among participants often take place only when they meet face-to-face. Collaborative possibilities rarely go beyond the boundaries of the classroom. It is also often difficult to moderate a class session in such a fashion that all participants are heard from and that a broad range of potential collaborative activities are accomplished within a relatively short period of time.

- *Discontinuity across different seminar sessions.* Face-to-face interaction can be very effective, but its effectiveness is bounded by its transitory nature. Continuity between seminar sessions, for instance, is difficult to maintain without a disciplined long-term *memory*, i.e., explicit tracking of both processes and activity contents and physical means of connecting them together. This is particularly true when meetings take place on a relatively infrequent basis and participants are physically distributed, as often the case with most seminars.

- *Lack of organic links between reading, writing, presentation and discussion.* Seminar activities are inherently both integrative and exploratory. To support either type of activity requires explicit representation and manipulation of links between various activity structures, contents, and processes. In traditional seminars, however, the awareness, storage and retrieval of such connections are largely left to individual participants. As a group, there exists no external *pool* of mutual artifacts which group members can contribute to and benefit from.

Perhaps in part due to their significance, most of the above problems have been addressed in existing learning tools. For example, computer-mediated communication is used to increase student participation outside physical classrooms [Hiltz, 1988]. Similarly, hypermedia systems, such as Intermedia [Yankelovich *et al.*, 1988] and NoteCards [Halasz *et al.*, 1987], are found quite effective for both authoring, browsing and presenting shared information. What is missing in these tools, however, is mechanisms necessary to facilitate metalearning, to help the student extract meanings from research papers, books, presentations, and discussions, and to enable the student to do both in a collaborative fashion. Intermedia, for instance, allows multiple users to concurrently create and follow links in the same web, but provides little hint on what those links and webs mean to the user.

Concept mapping is perhaps among the few attempts to provide explicit support for meta- and meaningful learning [Novak and Gowin, 1984] in classrooms. Systems such as SemNet [Fisher, 1990] are based on semantic network theory, a model of human memory and knowledge representation first proposed by [Quilian, 1967]. In fact, the term *semantic* *network* and *concept map* are used interexchangeably in SemNet. What differentiates the two seems that semantic networks are constructed by trained knowledge engineers for machine reasoning, while concept maps are built and used by human learners. This difference implies that concept maps are much simpler than semantic networks. It also reveals two potential pitfalls of concept mapping as a metalearning tool:

- *Atomic structure:* In concept maps, all knowledge must ultimately be reduced to concepts and links among them. This, though easily achievable for introductory textbooks, is far from adequate in advanced learning (e.g., graduate seminars), which often requires analysis and synthesis using high-level constructs, e.g., claims, problems.
- *Free form of expression:* Concept map, like the designer's sketch pad, gives the learner maximum freedom in deciding what to draw and how to draw it. The representation does not dictate nor provides any heuristic guidelines on how it should be used. This flexibility, which makes concept maps extremely expressive, also adds little structure useful as the basis of computation, and which human learners may rely on to help them make sense of the map. The latter is especially significant in a collaborative sense, for this arbitrary nature means that it is difficult to compare, contrast, and integrate concept maps generated by different individual leaners.

Few existing systems support collaborative construction of concept maps. In their study, for example, [Roth and Roychoudhury, 1992] have to rely on movable paper clips instead of a computerized system. Hence, the basic problem still remains:

> *Learning theories suggest that metaknowledge, metalearning, and collaboration are essential to meaningful learning. However, existing computer-based environments are largely for supporting content learning, communication, and information sharing. Even metacognitive tools such as concept maps fail to provide adequate structural heuristics for both computation and augmenting human learning.*

The above problem is what this research is intended to address. The following sections will provide detailed description of our approach.

2.3 Thesis

The fundamental premise of our approach is that a well conceived semi-structured representation, coupled with a set of properly chosen computational services, can provide a sound basis for facilitating collaborative meta- and meaningful learning. The key phrases here are "well-conceived" and "properly chosen". By that we do not mean that we hold the "holy grail" or "silver bullet", if such a thing indeed exists. Our notion of a "well-conceived" representation is one that provides the learner with maximum heuristic values in the following three areas:

- content learning: i.e., helping the learner organize and make sense of specific learning materials;
- exploration of metaknowledge: i.e., providing a basis on which new representational primitives might be identified; and

- ongoing and incremental interplay between the above two.

Our notion of "properly chosen" services are the ones which support this interplay between representational/meta-level exploration and content learning. Because of this structural indeterminacy, we believe that examples are essential in directing proper use of a representation. In our system, we strive to provide examples at various levels of abstraction and allow them to be upgraded on an ongoing basis.

The central thesis of this research is that CLARE represents a novel and useful approach to collaborative learning: CLARE supported learning will be more effective than that of the traditional mode of learning, i.e., face-to-face and pencil-and-paper based. It will also show significant improvement over other existing computer-based tools, such as concept mapping, or hypermedia. The focus of this research is on the former. Specifically, the two experiments described in Section 5 are designed to empirically test this claim. In the near future, we plan to conduct experiments which compare CLARE with other tools of the same class, e.g., concept mapping.

2.4 Contributions

The research contributions of this work can be viewed at three levels:

1. Conceptually, RESRA represents a new approach to collaborative learning that is based on constructivist pedagogy and the Ausubel-Novak-Gowin theory of meaningful learning. It overcomes the structural weakness of concept maps. On the other hand, RESRA constructs are not intended to restrict the user's expressiveness and how they learn but rather, serve as a heuristic basis for computation and for evolving a representation appropriate for meaning extraction and knowledge construction in specific group and domain settings.

2. Technically, CLARE implements a set of services for facilitating content and metalearning based on RESRA. The five levels of abstraction it embodies, i.e., RESRA primitives, domain, template, and example libraries, and specific instances, creates room for both computation and learner exploration. The explicit treatment of perspectives and the provision of a multiway comparator illustrate collaborative services beyond simple access control, information sharing, and online annotation.

3. Empirically, our evaluation experiments will provide some primary data on how students react, use, and think about our approach. They will also allow us to assess the effect of CLARE on the student's performance of selected learning tasks. The outcome of these experiments should shed important light on proper mechanisms for supporting collaborative metalearning, and provide a basis for further investigation.

2.5 Limitations

Collaborative learning is a complex activity to study and support. In this research, we do not attempt to address every aspect of the subject, nor to create a universal system that has all neat features of existing learning support tools. Instead, we have chosen to focus on the role of representation in collaborative learning and on developing pertinent computational mechanisms to facilitate the use of such a representation in various learning activities. Below are four major limitations of the current research:

1. CLARE is not an AI system. It does not provide automatic inference, user modeling, or natural language understanding capabilities. RESRA, the underlying knowledge representation scheme, is intended for helping human learners evaluate and construct knowledge rather than for improving machine reasoning, despite that many useful computations in CLARE are centered on RESRA.

2. CLARE's artifact-based approach renders its emphasis on the structural characteristics and relationships of collaborative learning. It presumes the presence of certain learning activities, e.g., evaluation, argumentation. However, it does not prescribe any process model, i.e., how those activities should be combined. Instead, it insists that choosing a proper learning model is the responsibility of the course or experiment designer rather than the system designer.

3. CLARE helps ameliorate certain problems related to face-to-face collaboration (e.g., "dominant personality") to the extent that it is a computer-mediated environment. However, it cannot entirely overcome the interpersonal and intercultural conflicts inherent in collaboration. When used in a face-to-face setting, it is possible that the effectiveness of CLARE as an augmented learning tool be overshadowed by some interpersonal or intercultural factors. CLARE provides no means to separate one from the other.

4. At the system level, CLARE does not have certain advanced functionalities found in other existing learning support environments. Examples include multimedia, version control, and fancy graphical interface. This is in part due to the pilot nature of our system: we plan to incrementally incorporate above functions as we have opportunity to empirically validate the usefulness of the CLARE's core functionalities.

3 RESRA: the Thematic Representation

3.1 The Structure of Learning Artifacts

The structure of learning materials or, in our terminology, *artifacts*, such as journal articles and research reports, may be viewed at various levels. At the top level, the student sees an article as consisting of a title and a number of sections, subsections, figures, tables, et al. This structural distinction is based on the visible attribute, and is *presentational* in nature. Often, it has little bearing on the content of the underlying artifact. Looking down one level, however, one may note that the same artifacts also use such standard headings as "abstract", "introduction", "experimental design", "outcome", "related work", "conclusions", and so forth. These labels tell the learner the type of content immediately followed. They are, however, still primarily organizational, and therefore belong to the *presentational structure*, or PS.

Argumentation plays a critical role in both science and learning processes [Cross, 1990]. *Rhetorical structures* or RS, such as the one proposed by [Toulmin *et al.*, 1984], provide a useful means of understanding research artifacts, especially, relationships between artifacts.

For instance, the contents of a series of articles published in the "technical communication" columns of a professional journal are capturable using rhetorical models. Unlike PS, which focuses on the *surface structure* of learning artifacts, RS represents the *deep structure*, which cannot be derived without first understanding the artifact content. Some research is devoted solely to the RS-based learning (e.g., [Cavalli-Sforza *et al.*, 1992]).

The third type of structure is the *thematic structure*, or TS. As the name implies, TS characterizes the theme or essential elements of learning materials and relationships between them. Similar to RS, TS is content-oriented. It, however, goes beyond RS in that it models both the discursive and the domain structure, both intra- and inter-artifact relationships. For example, TS can include such primitives as "concept", "claim", which, when instantiated into the field of software engineering, might include "software complexity" (concept) and "Object-oriented design offers an effective solution to software complexity" (claim). These features of TS make it a valuable basis of knowledge representation, which explains why our representational scheme, i.e., RESRA, is based on it.

3.2 Knowledge Representation and Metalearning

Knowledge representation (KR) is often associated with automated reasoning and machine-based intelligence. In this context, however, we are interested in using KR to facilitate and augment human learning. Since, according to [Swaminathan, 1990], KR models are either epistemological or ontological, KR-centered learning is metalearning, that is, it helps the student deciding what kind of knowledge to use as well as how to structure it.

Our thematic representation (i.e., RESRA, see Section 3.4) falls to the category of content theories. As observed by Swaminathan, content theories are usually incomplete, vague and, as a result, do not usually lead to unique mappings from the text under study into the primitives proposed by the theory [Swaminathan, 1990]. Though these problems have been a constant source of criticism against certain KR schemes, such as Schank's theory of conceptual dependency (CD) [Schank, 1975], they are desirable features in the current context, because the representation is no longer primarily for machine reasoning, but rather a heuristic basis for human learning. The meta and heuristic values of the representation are manifested in the following four ways:

- a mapping tool that highlights essential elements and relationships within as well as across learning artifacts;
- an organizational tool that allows the learner to dynamically and incrementally integrate various types of learning artifacts at a fine-grain level;
- a communication tool, i.e., a shared "frame of reference" in group collaboration: contrasting different representations of the same artifact by different group members can highlight the differences among group members, while integrating them can lead to a fuller understanding of the subject domain; and
- a learning tool for the student on the conventions governing the written presentation of learning and research results.

3.3 Five Levels of Collaborative Learning

Learning in a group setting may take many different forms, including joint projects, writing, reading, discussion, and so forth. Our observation of these activities has led to these five common themes: summarization, evaluation, integration, argumentation, and construction, as described in Table 1). The actual division between these types, however, may never be as clear-cut as what the table suggests. For example, an argument is often filled with constructive ideas and value-laden judgement. In addition, the order in which the activity types appear in the table also does not dictate the actual process steps in which they occur, though the sequence is typical. One implication of this categorization is that a representation needs to be able to accurately express the semantics of artifacts generated from all these five types of activities, and to relate and integrate them in some useful fashion. RESRA, which will be described next, is such a language.

Activity Type	**Description**	**Example Artifacts**
Summarization	Extracting, condensing, and relating important elements from an artifact.	List of hypotheses and findings from a research paper.
Evaluation	Subjective appraisal of a given piece of work.	Criticisms on the flaws of an experimental design.
Integration	Relating, aggregating or abstracting previously scattered themes.	A state-of-the-art survey on a given topic.
Argumentation	Interactions among polarized points of view with regards to a given topic.	Recorded script of a panel discussion on information privacy.
Construction	New proposals, formulations, or interpretations of new or existing problems or solutions.	RESRA as a new way of supporting metalearning.

Table 1: Five Level of Collaborative Learning

3.4 RESRA

RESRA, i.e., REpresentational Schema of Research Artifacts, is a specialized language for representing the thematic structure of research and learning artifacts generated from both in and outside of classrooms. In essence, RESRA is composed of two primitives: *entity* and *relation*. The former describes the property and the structure of artifacts and the latter, relationships between entities. Table 2 lists the set of entities currently defined in RESRA. Figure 1 shows graphically the relationships between those entities.

RESRA is designed to support the five levels of learning described in Table 1, even though Figure 1 does not show such correspondence. Below are some illustrations.

Summarization Summarative primitives are the basis of RESRA representation. Normally, for a well-defined artifact type such as surveys, conceptual papers, empirical reports, a RESRA "template", which consists of a set of entities and relations, may be defined. For

Entity Type	Description	Example
Source (SO)	Identifiable written object, either object itself or a pointer it, i.e., reference.	An article by Ashton; the notes from Kyle's talk.
Problem (PR)	A phenomenon, event, or process whose understanding requires further inquiry;	Metalearning is not adequately supported by existing computer-based tools.
Claim (CL)	A position or statement about a given problem situation.	CLARE can help the student learn how to learn.
Evidence (EV)	Data gathered for the purpose of supporting or objecting to a given claim.	The result of our experiments has shown that CLARE users generate better quality research reviews than that of non-users.
Method (ME)	Procedures, models, or actions used for generating evidence for a particular claim.	Three-week experiment involving six groups (group size = 3), three of use CLARE, and the other three do not.
Concept (CO)	Primitive construct used as building blocks for problem statements, theories, claims, and methods.	metalearning; knowledge representation.
Theory (TH)	A systemic formulation about a particular problem domain, derivable through deductive or inductive procedures.	Ausubel's theory of meaningful learning.
Thing (TI)	A natural or man-made entity that is under study.	Atom, NoteCards.
Critique (CR)	Comments on a given claim, evidence, method, source, et al.	CLARE claim will much be strengthened by including example usage experience of the system.
Question (QU)	Aspects of a claim, theory, concept, etc., about which the learner is still in doubt.	How does CLARE differ from such systems as NoteCards?
Suggestion (SU)	Ideas, recommendations, or feedbacks on how to improve an existing problem statement, claim, method, et al.	

Table 2: Primitive Types of RESRA

example, an empirical study is expected to contain instances of such primitives as *problem*, *claim*, *method*, and *evidence*, while in a conceptual paper, *concept*, *theory*, and *claim*, and in a survey, *source*, and *claim*. These templates provide useful thematic heuristics which orient the learner's attention and lead to unusual discoveries such as uncovering implicit *problem*(s) or *claim*(s) in a research paper.

Evaluation One thrust of using RESRA as a evaluation tool is the fine-granularity it entails: instead of merely listing major strengths and weaknesses of a given work at the artifact level, *critique*, *question*, and *suggestion* are directed at microscopic structures like *claim*, *method*, *evidence*, and so forth, and relationships among them. This "deep-level" evaluation requires the learner to have a good grasp of the artifact under concern, and keeps him/her in a constructive mode by asking questions and offering suggestions.

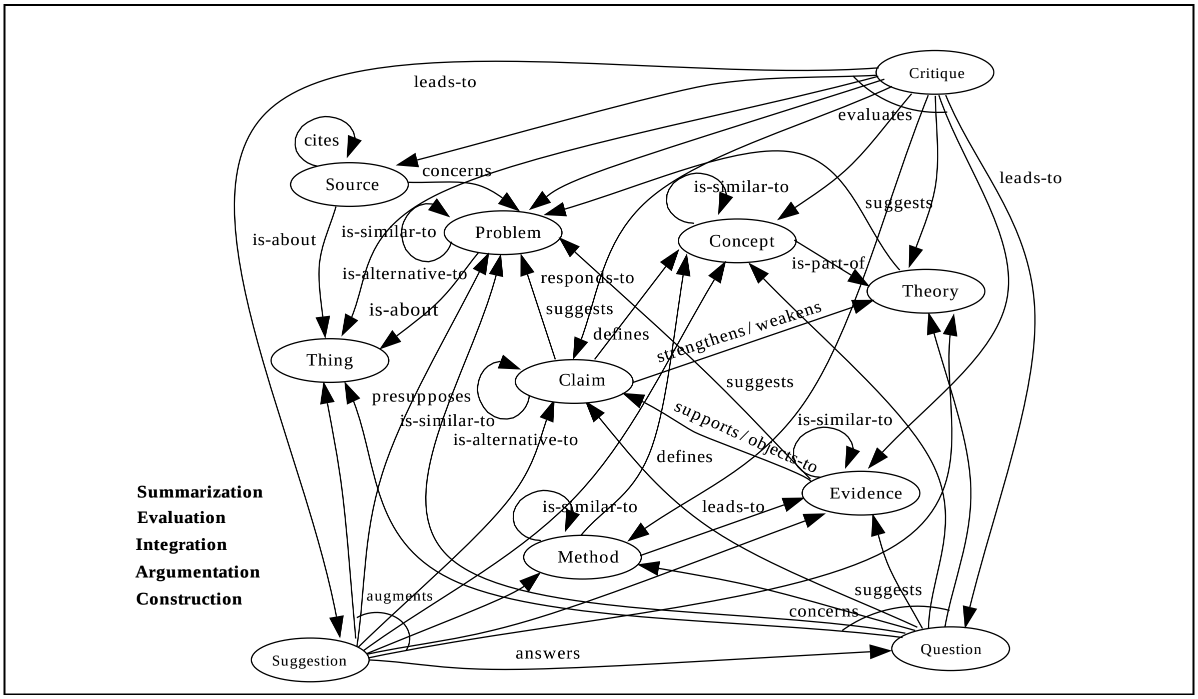

Figure 1: A Graphical Representation of RESRA

Integration As a key component of collaborative and meaningful learning, integration involves relating, linking, and consolidating RESRA instances created by different individuals and about different artifacts. As illustrated in Section 3.5, it normally takes place after individual learners complete their summarization and evaluation. In other words, a repertoire of RESRA instances must pre-exist prior to integration.

Integration often implies generalization and abstraction. For example, a good survey paper is not merely a rehashing of existing artifacts; instead, it needs a coherent framework to help bring together and make sense of related artifacts. Although RESRA does not define any aggregate, CLARE is equipped with its own aggregation mechanisms (see 4.2).

Argumentation RESRA subsumes two commonly used rhetoric models, i.e., IBIS [Conklin and Begeman, 1988] and [Toulmin *et al.*, 1984]. Argumentation in RESRA involves making alternative claims, defending existing claims using evidence, or posting questions on both. It is commonly employed as a means to resolve representational differences among individual learners so that integration might be achieved. However, RESRA can also be used to capture the rhetorical process exhibited in research literature by linking together through appropriate relation types related artifacts, e.g., a chain of articles that centered on a provocative primary work.

Construction RESRA is fundamentally a knowledge construction tool. Because of its thematic nature, RESRA can help expose gaps in existing knowledge by juxtaposing contending/related claims and different learners' perspectives, and by highlighting essential elements and relationships among various artifacts. It also encourages the learner to ask questions which in turn forms a basis for further inquiry. RESRA may also be considered as a sophisticated idea structuring tool for individuals and groups alike. It allows one to "index" ideas as they appear, and eventually leads to a systematic formulation, and perhaps, a new artifact.

In sum, RESRA offers a sound representational basis for supporting all five levels of learning. Though the description thus far is largely confined to the predefined RESRA primitives, it is important to realize that one principal feature of RESRA is its open-endedness: the users are not only allowed but encouraged to extend and adapt the initial set of primitives, and engage discussions about proper representational structures for the given domain and group settings. The value of RESRA resides in two qualities: heuristic and definitional, but more in the former than in the latter. Therefore, a computational environment that claims to support RESRA, e.g., CLARE, ought to provide services that exploit this heuristic nature.

3.5 An example use of RESRA

Figure 2 and Figure 3 are two example RESRA representations (condensed for the sake of space) from two individuals about the same research report [Kaplan and et al, 1992], both of which include summarative and evaluative instances. Though **R1** and **R2** look alike in some ways, they differ significantly in others. Those differences, as shown in Table 3, reflect different viewpoints, focuses, and levels of understanding of the artifact by the two students.

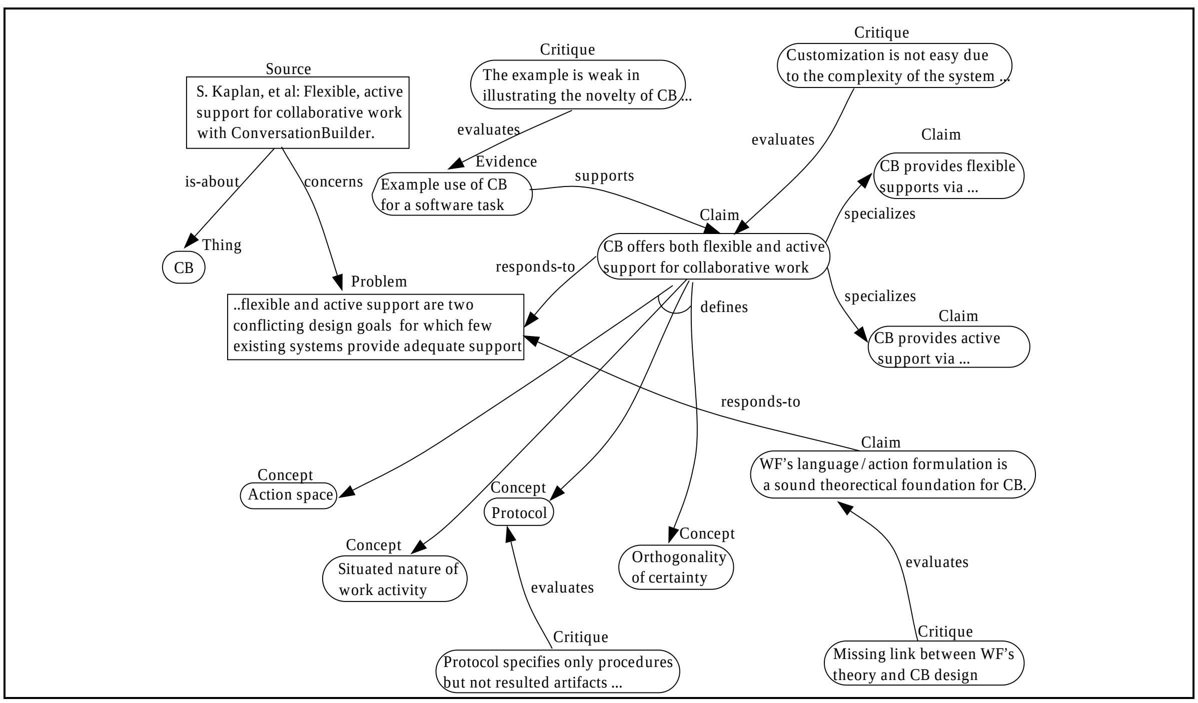

Figure 2: RESRA Representation of [Kaplan92] by Student A

Suppose that the two students who created the above representations are given opportunity to see each other's work. In light of obvious differences, they are expected to ask each other many questions. For example, student **A** may ask student **B** why the example CB usage scenario in the paper is treated as *thing* instead of *evidence* supporting the authors' claims. Similarly, student **B** may ask student **A** how he has come up with four *claims* instead

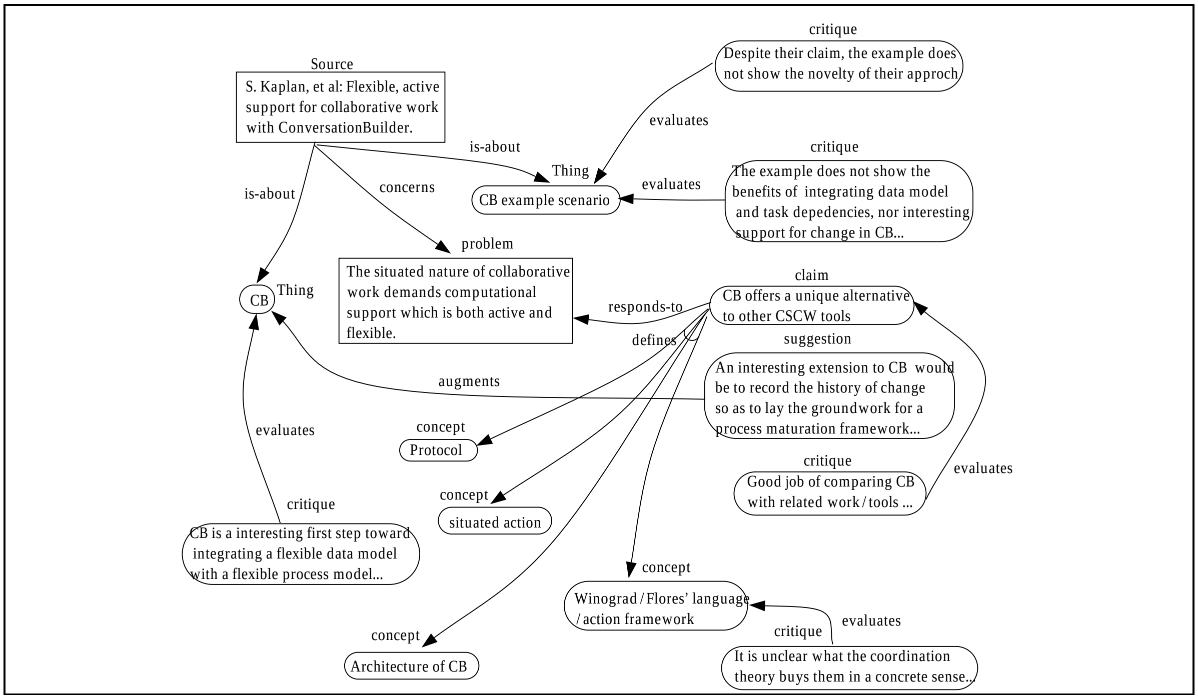

Figure 3: RESRA Representation of [Kaplan92] by Student B

of just one. Such an exchange helps the students better understand each other's perspectives with regard to the artifact and allows them to reach a consensual view of the artifact itself. It also leads to the creation of additional RESRA instances which expose the reasoning behind the initial representation. Figure 4 shows merely one of many possible reconciliated representations. Note that such a figure may be the evolutionary result of several intermediate RESRA representations.

4 CLARE: the Computational Environment

4.1 Functional Overview

CLARE is a computer-based collaborative learning environment based on the RESRA representational framework. Though the latter forms the conceptual core of the former, CLARE contains several novelties of its own, including aggregates (i.e., thread and perspective), a comparator, a hierarchical illustrator, and a number of exploratory support facilities. At the user level, CLARE is a hypertext-based system that provides the following functions:

- Supporting distributed collaboration, both synchronous and asynchronous;
- Ability to create, update, navigate, and summarize RESRA instances;
- Ability to extend, adapt, and explore RESRA primitives;
- Hierarchy of examples, templates, and reusable domain-specific instances for facilitating example-based learning;

R1	R2
4 claims, including 2 sub-claims.	1 claim.
Winograd/Flores' language/action theory is a claim.	W/F theory is a concept.
The example is evidence.	The example is a thing.
0 suggestion.	1 suggestion for future direction.
2 critiques on claims, 1 on example, and 1 on concept.	1 critique on claim, 1 on concept, and 3 on things.
4 concepts, two of which are not in R2.	4 concepts, two of which are not in R1.

Table 3: Key Differences Between Two Representations

- Aggregation in terms of user, artifact (i.e., thread), and perspective.
- Ability to compare representations from different individuals about the same artifact;
- Incrementally organized knowledge base capturing group learning experience; and
- Graphical navigator for overviewing the structure of collaborative knowledge base.

4.2 Main Features

The design of CLARE was driven by the following assumptions:

- RESRA is the structural basis of the system; in other words, CLARE should reflect the five-level model of learning and allow manipulations of RESRA objects at both representational and instance levels;
- Representational exploration is an integral part of any metalearning tool, including CLARE;
- The ability to create higher-level entities, i.e., the ability to aggregate, generalize, and abstract, is essential to meaningful learning, and therefore must be supported by CLARE.
- Collaborative learning requires explicit mechanisms for representing individual viewpoints and means of comparing, contrasting, and integrating them; and
- Examples and templates are important in reducing structural uncertainty and are a basis for generating group consensus.

Based on these assumptions, CLARE incorporates the following features: a mode manager, aggregation, an illustrator, a comparator, an explorer, and a navigator.

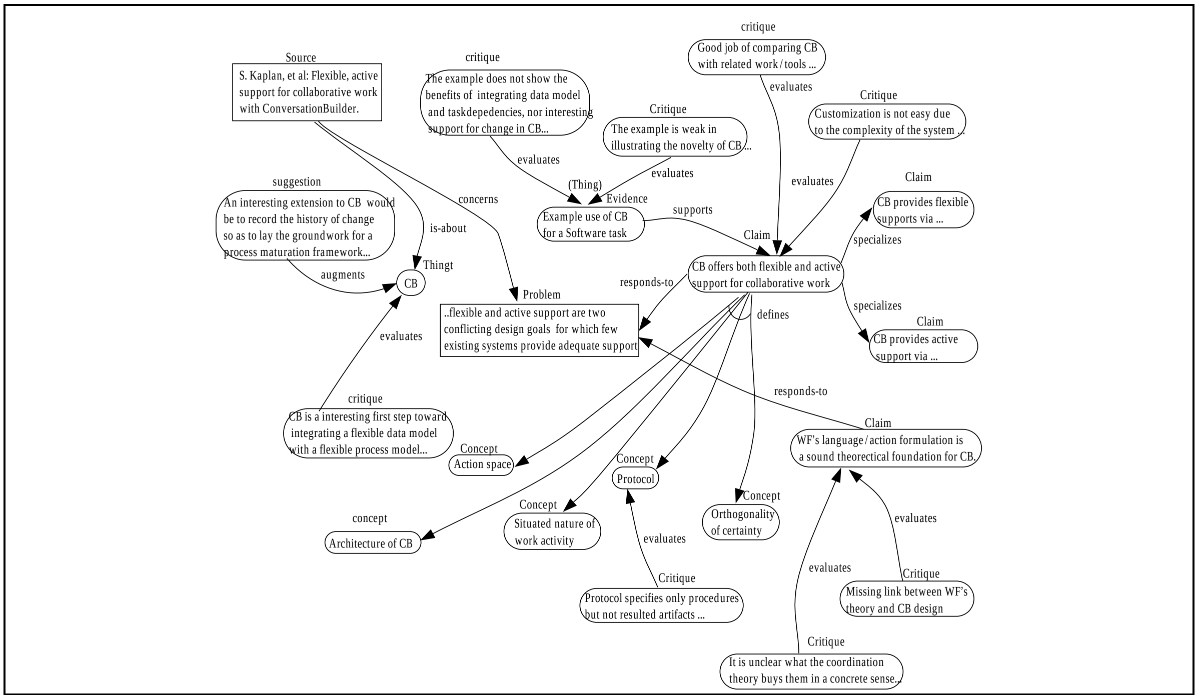

Figure 4: An integrated RESRA Representation of [Kaplan92]

4.2.1 Mode Manager

At the base level, CLARE provides five interaction modes, which correspond to the five levels of learning identified in Section 3.3, i.e., *summarization*, *evaluation*, *integration*, *argumentation*, and *construction*. In each of the five modes, the user is allowed to create, update and browse RESRA instances appropriate to the mode. One may toggle between those modes but, at a given time, only one mode is active and visible to the user (i.e., via pulldown menu). The intent of this design was to give the learner a framework to organize their activities. If the learner needs to, for example, create an *evaluation* instance while in the *summarization* mode, they can do so either by temporarily switching to the *evaluation* mode, or using instance creation function from the generic mode, in which case they are responsible for explicitly specifying the type of instances to be created. Alternatively, if such combinations are often called for, the user should consider using the explorer to add the function to the current mode (see Section 4.2.5 for more details).

4.2.2 Aggregation

As mentioned in Section 3.4, RESRA in itself does not come with any aggregate. To overcome this deficiency, CLARE defines two of its own aggregates: *thread* and *perspective*. A *thread* is a set of RESRA instances that are related, directly or indirectly, to a *base* entity, while a *base* entity is a RESRA instance that serves as a *center* of discussions. By default, all *source* and *problem* instances are base entities. The user, however, is free to add or remove instances from the default set by specifying either a primitive identifier, such as "concept" (i.e., meaning all "concept" instances) or an instance identifier, such as "metalearning". In

CLARE, threads are indexed by their base entities. They are used not only for querying and browsing but also as the basic grain for examples. See Section 4.2.3 for more details.

A *perspective* defines a set of RESRA instances that share a consistent pattern of viewing a given artifact, problem, and so forth. It is generally not as fluid as *thread*. Unlike threads, which are dynamic and automatically defined once the *base* entities are identified, a perspective must be explicitly described by the learner, and a name is normally required at the time when it is defined. The perspective provides a structural basis for facilitating group consensus building. Typically, a learner holds one perspective with regard to a given *thread*, e.g., the implementor or designer perspective. Different learners, however, may share the same perspective.

4.2.3 Illustrator

The illustrator is designed to provide CLARE users with examples and templates to facilitate the understanding and consistent use of RESRA primitives. It consists of three libraries: domain, templates, and examples.

- Domain-specific instance library. The library contains a set of standard RESRA instances, such as concepts, theories, claims, et al, and relationships between them. Those instances may not be related to any particular artifacts. Rather, they are probably created by the expert on the subject (e.g., the course instructor), or extracted from RESRA instances generated by previous learners, or a combination of both. Such a library may be viewed as the RESRA version of the core knowledge in a given domain. It serves as a point of reference for exploring less understood aspects of the subject.

- Template library. As described in Section 3.4, for each well-defined artifact type such as survey, conceptual paper, empirical report, a RESRA "template" can be defined. Such templates consist of a set of primitives. For example, an empirical study contains instances of such primitives as *problem*, *claim*, *method*, *evidence*, while a survey may contain only *artifact* and *claim*. Such templates, of course, do not always exactly match with the artifacts under concern. Nevertheless, they are "ideal types" which provide useful thematic heuristics in orienting the learner's attention, guiding the proper use of RESRA, and leading to unusual discoveries, such as uncovering implicit *problem*(s) or *claim*(s) in a research paper.

- Example library. Compared to the preceding two, the example library is less formal; instances from this library are typically created by the learners themselves. They differ from other instances in their typicality, good or bad. Examples are important in enabling the learner to learn from other other people's experiences. They also illuminate how abstract constructs (e.g., RESRA primitives) are used. Examples are typically identified by threads and indexed by respective base entities.

What is common to all three libraries is their dynamic nature: CLARE users are allowed to add, delete, and modify instances from those libraries; of course, they can always lookup, view, and navigate them as well.

4.2.4 Comparator

To better understand individual differences and similarities and to facilitate consensus building among group members, CLARE provides a comparator function which computes a similarity metric, called *similarity score*, between any two individual learners, and between that learner and the group as a whole. The similarity score, whose value range between 1 and 100, with 100 indicating the most similar, is based on a number of factors, including the number of instances created, the type and the size (for entities only) of instances, and the number of references to the domain, template, and example libraries. The similarity function, i.e., the number of factors and the weight of each factor, is customizable. Typically, similarity scores are computed on per thread basis, though it is trivial to aggregate multi-thread scores.

In addition to the similarity score, which serves as a high-level quantitative index to the group view of an artifact, the comparator also reports the number of RESRA instances by category. From that point, the learner may "zoom in" to individual instances that are of interest. See Section 4.5 for an example use of the comparator.

4.2.5 Explorer

CLARE is primarily a metalearning tool for using RESRA metaknowledge primitives to structure specific knowledge under concern. Since RESRA is inherently heuristic (see Section 3.4), its structure and semantics, which depend heavily on the characteristics of the user group and the learning task on hand, will inevitably undergo change. CLARE provides full support for evolving not only RESRA but also for CLARE's own extensions, such as aggregates, similarity scores. Its exploratory functions include:

- Extending, i.e., adding, deleting, and modifying, predefined RESRA primitives;
- Adding, deleting, and modifying the field structure of RESRA entities;
- Changing relationships between the five interaction modes and RESRA primitives;
- Defining new aggregates;
- Customizing the similarity score function, i.e., adjusting the number of factors or their relative weights, and
- Allowing online conversations on above structures.

Most of the above functions are directly implemented using the exploratory type system of EGRET, the platform on which CLARE is built (see [Johnson, 1992b] for details).

4.2.6 Navigator

CLARE is equipped with a graphical browser, called *navigator*, which allows the learner to overview RESRA network structures and "zoom in" to individual objects. This capability is available to both ordinary RESRA instances, (domain, template, example) library instances, aggregates, and RESRA primitives. Though providing an advanced graphical interface is not our primary design goal, we consider the ease of use as an important part of the system functionality. The navigator, in particular, is important in easing the "lost in hyperspace" problem.

4.3 Architectural Components

Architecturally, CLARE is composed of six components: hyperbase server, agent, EGRET, libraries, interface, and navigator. The relationships between these components are depicted in Figure 5. The major functions provided by each are briefly summarized below:

- Hyperbase: A database engine which provides persistent stores for node and link data. It ensures the data consistency through the field-level locking/unlocking, and the state consistency between multiple clients through the event mechanism.
- Agent: A distributed, specialized background process(es) that maintains necessary global data for realizing high overhead client functions. The agent is driven by events generated by the Hyperbase server based on client actions.
- EGRET: A generic platform for supporting distributed collaboration. It provides infrastructure support (i.e., gtables) and interfacing machinery with the hyperbase. EGRET's exploratory type system is the basis on which RESRA objects and CLARE's explorer are implemented.
- Libraries: It includes RESRA primitives, domain instance, template, and example libraries. They are implemented using the gtable mechanism provided by EGRET. An agent is responsible for maintaining the consistency of these libraries.
- Interface: It consists of a large set of functions that are organized into these groups: the mode manager, aggregation, the illustrator, the comparator, and the explorer. It interacts directly with the navigator.
- Navigator: A graphical browser that shows the structure of RESRA networks and allows "zoom-in" capabilities. It works with all levels of RESRA objects, from primitives, templates, examples, to ordinary instances.

4.4 Implementation Environment

The prototype of CLARE is implemented on top of Lucid Emacs, an X Window based version of the popular GNU Emacs editing environment [Stallman, 1985]. EGRET, which is runnable on vanilla Emacs, provides CLARE necessary low-level support. The database server is HyperBase, developed at University of Aalborg [Wiil and Osterbye, 1990]. The graphical navigator is implemented using the XView library.

4.5 A Sample Session with CLARE

John, Doug, and Lynn were taking a seminar in computer-supported cooperative work (CSCW) together. Since I had mentioned to them one time about the greatness of CLARE, they later came back to me and asked whether they could give the system a try; they were curious about what CLARE could give them which they couldn't get from other systems. After a 15-minute demo and showing them a few examples, I asked them to (1) select a paper from their required reading list; (2) privately (i.e., without looking over each other's shoulders) summarize and evaluate the paper using CLARE; and (3) come together to look

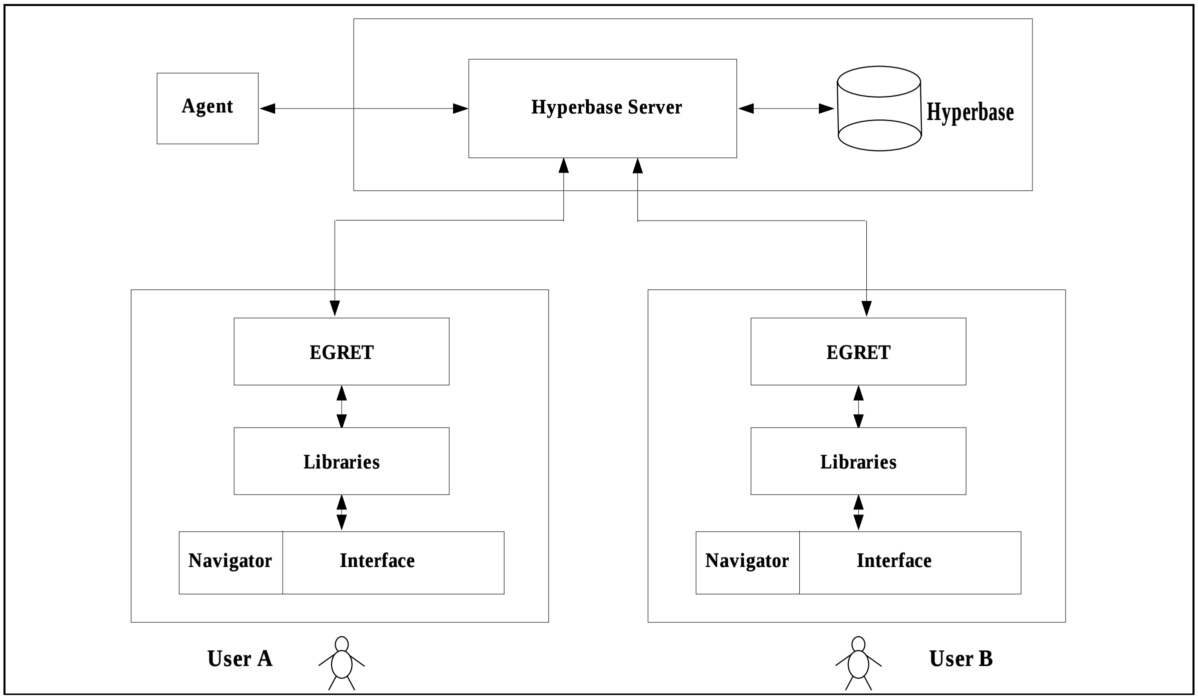

Figure 5: CLARE's Architecture Components

through what they have created individually, ask each other questions if necessary, and see whether they can use CLARE to link together their individual representations to form a consensual view of the paper.

Two days later, my three friends and I gathered in front of a workstation. We brought up CLARE, and invoked the comparator on [Kaplan and et al, 1992] – the artifact under scrutiny. Figure 6 provides a high-level index to the difference between what the three had done. The group similarity score, shown near the bottom of the screen, is 28.4%, implying that the three had little in common. The comparator table reports the number of summarative and evaluative RESRA instances created by each user; each cell contains two numbers: one is the number of entities and the other, of relations (inside parentheses). For example, Doug has created 8 summarative entities, 3 summarative links, 3 evaluative entities, and only 2 evaluative links. Note that the similarity scores listed in the table are two-way scores: since Doug is the current user, he has a similarity score of 32.1 with Lynn, and 24.7 with John. To illustrate why the scores are so low, Figure 7 and 8 provide two example *problem* instances, which are "problem" representations of the paper by John and Doug, respectively. At the first glance of the two instances, I can hardly tell they were talking about the same artifact!

Given their differences, my friends spent almost two hours navigating through each other's work, asking themselves questions, and explaining to each other why they came up with what did. At the end, they were able to link their individual views and ideas together and generate something close to a consensual representation (See Figure 4 for a skeleton of the final representation). By going through this process, John came to know why Doug had treated the example in the paper as an object of study, i.e., *thing* , instead of as *evidence* to support the authors' claims. Similarly, Doug finally realized how Lynn had come up with

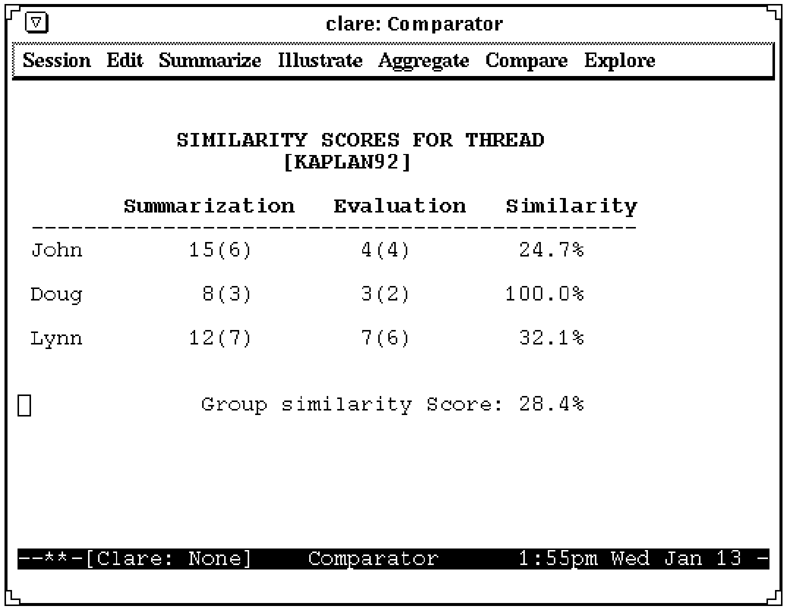

Figure 6: An Example Similarity Report

four *claims* instead of just one, as he did.

4.6 Potential Usages of CLARE

CLARE was originally conceived as a collaborative research review system to be used in advanced learning settings such as graduate seminars [Wan and Johnson, 1992]. During the course of our research, however, it has evolved into a much more generic environment that can support a wide range of learning tasks, both collaborative and individual-based ones. In this section, we focus on the collaborative aspect and describe three key areas in which CLARE might be applied: reviewing, writing, and discussion. One unique feature of CLARE is that it offers fully integrated support for all the three activities.

4.6.1 A Collaborative Review Environment

Reviews not only provide a way to discover defects in products, as used in engineering design [Freedman and Weinberg, 1990], but also is a good learning tool which allows the student to develop critical skills, learn to differentiate good work from bad ones, learn to relate existing work and from other people's experience. As such, reviews are commonly employed in classrooms, for example, book reviews, literature reviews, project reviews, and so forth.

In the CLARE environment, reviews are conducted at two levels: *summarization* and *evaluation* (see Section 3.3), with the former typically preceding the latter. To some learners, this distinction may seem arbitrary and difficult to put into practice. The assumption behind the CLARE's approach is that whether or not such a separation lead to more better learning is a question requiring empirical investigation. CLARE provides a support environment for conducting such studies. For example, one may hypothesize that the separation of evaluation from summarization reduce the level of "free riding" in collaborative learning settings. To

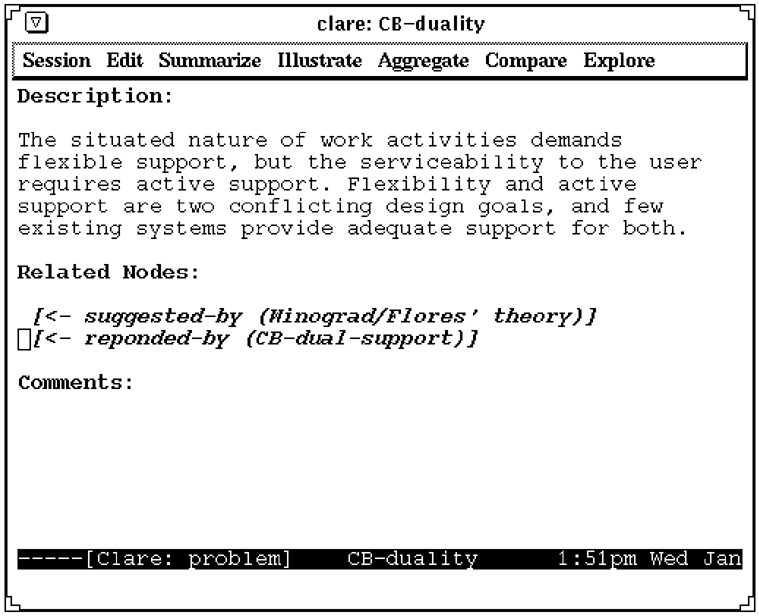

Figure 7: Doug's "Problem" Representation of [Kaplan92]

test this hypothesis, one can design a CLARE-based experiment which requires that, in one group, learners be allowed to access and make reference to each other's ongoing work, while in the other, individual work is kept private until all members finish their summarization and then, they proceed to the evaluation phase. One might expect that the second group generate higher-level quality evaluations because they are barred from taking "shortcuts" via merely linking to or "signing off" other leaners' work.

4.6.2 A Collaborative Writing Tool

As a learning activity, collaborative writing (CW) consist of not a singular task but a number of interrelated tasks, which include brainstorming, idea organizing, planning, composing, editing, revising, et al [Baecker and et al, 1993]. Through its *construction* mode (see Section 3.3), CLARE supports primarily the first three, or the early phase of CW. In a sense, CLARE is an advanced group outliner, which differs from ordinary outline editors in the following three ways:

- It is hypertextual instead of linear;
- It is primarily for collaborative idea exploration, organizing, and planning; and most important of all,
- It uses RESRA primitives and its own aggregates as a knowledge representation scheme.

By using a linearization utility which we plan to incorporate into CLARE, the learner will be able to obtain a linear representation of "idea webs" from the system, which serves as a basis for later stages of CW, namely, composing, editing, and revising.

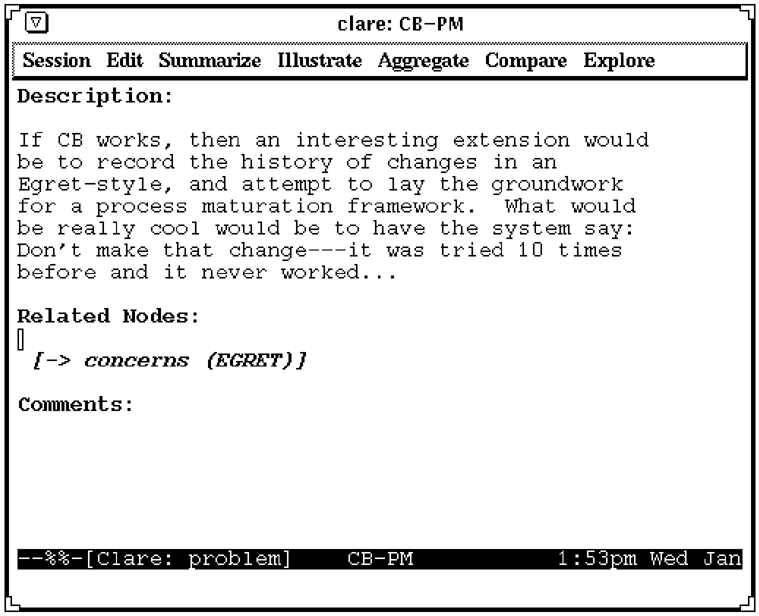

Figure 8: John's "Problem" Representation of [Kaplan92]

4.6.3 A Structured Bulletin-board System

The concept of "virtual classroom" is supported mainly through electronic mail and bulletin-board systems [Hiltz, 1988]. CLARE augments those environments by providing a number of structuring mechanisms, including RESRA, which allows fine-grained representation of conversations and discussions, and aggregates (i.e., threads and perspectives), which allow individual differences and similarities to be highlighted, compared, contrasted, and integrated. CLARE's argumentation mode, in particular, is designed to facilitates focused discussions and deliberations within tightly coupled learning groups.

5 Evaluation

The effectiveness of CLARE as a collaborative learning tool will be evaluated through two experiments. The basic claim motivating both is that CLARE can help the learner improve their performance in selected learning tasks. The experiment subjects are advanced learners, i.e., upper-level undergraduate and graduate students. The experiments involve two types of tasks: generation of study questions and joint writing of research review papers, both of which are done based on a set of research papers.

5.1 Experiment I

5.1.1 Tasks

This experiment involves a joint generation of study questions from a selected set of papers in software verification and validation (V & V). The subjects are graduate students enrolled in ICS 613 (Advanced software engineering) at the University of Hawaii. The subjects will be given 2-3 papers on a specific topic in software V & V. They are expected to read them and, based on the reading, generate and turn in a set of important questions. The purpose of this experiment is to assess the effect of CLARE on the the quality of resulted questions and on the process through which those questions are generated.

5.1.2 Procedure

The experiment involves 12 subjects, which are randomly assigned into two groups: the treatment group which uses CLARE, and the control group which relies solely on the face-to-face and pencil-and-paper based learning. Each group is randomly divided into two (2) study groups, with three (3) students in each. The experiment will start at about one-month after the semester begins, and will last for two weeks. During these two weeks, two experiment sessions will be conducted each week. Each experiment session is divided into stages: private question generation and group consolidation. For both groups, the first stage is done outside classrooms: the treatment group creates questions online, while the control group records their questions on the paper.

The second phase is conducted in the face-to-face setting. During this phase, the treatment group will be gathered in front of a workstation, reviewing, discussing, and integrating questions generated by the individual members from the previous phase. A print function will be invoked at the end of this session to print a hardcopy of all the questions generated by both individual and groups, which will be turned in for evaluation. Meanwhile, the control group will consolidate their questions through face-to-face discussions. One member in the group will be assigned to record questions agreed upon by the group. Like the treatment group, at the end of the session the control group must turn in both individual and group questions. Although both groups may use as much time as they need for the first phase, the time for the second phase is fixed, i.e., a fifty-minute class period.

Since there will be four (4) experiment sessions during the 2-week experiment period, each study group will have opportunity to alternate between the treatment and the control group. By the end of the experiment, each subject will have two sessions with CLARE, and two sessions without. A written questionnaire will be administered at the end of each experiment session.

5.1.3 Measures

This experiment will be evaluated on the basis of three grounds: *outcome*, *process*, and *learner satisfaction*. The outcome measure is the number and the quality of questions raised by the individual study group. A printed copy of all questions from each group will be collected and distributed to three experts in the field, including the course instructor. They will be graded on the scale of 1-100 based on the clarity of questions, the level of understanding of the materials shown, and the level of integration with related literature. The group identity of those questions will be hidden from the evaluators.

The face-to-face sessions for both treatment and control groups will be videotaped. The transcript will be analyzed at the end of the experiment. For the CLARE group, additional process data will be captured by the system through its instrumentation facilities. Examples include the type and the number of instances generated, the time they were created and modified, the time used in creating those instances, the usage frequency of selected functions, e.g., example and domain instances, the density of links, and the similarity scores.

The learner satisfaction measure for both treatment and control groups will be evaluated through written feedback. At the end of each experiment session, a questionnaire, which contains scaled and open-ended questions, will be administered to each group. The data will be analyzed with reference to the printed copy of questions and the process data gathered.

5.2 Experiment II

5.2.1 Tasks

This experiment involves a joint writing of a research review papers on software requirement analysis. The subjects are upper-level undergraduates, i.e., juniors and seniors, enrolled in ICS 413 (Software engineering) at the University of Hawaii in the Spring Semester, 1993. The subjects will be assigned 3-5 papers on software requirement analysis from current journals and conference proceedings. They are expected to read those papers, discuss them, and write a critical review that integrates issues/problems addressed, compare and contrast approaches proposed, and generate a list of further questions. The purpose of this experiment is to assess the effect of CLARE on the quality of resulted review paper and on the process by which the review paper is produced.

5.2.2 Procedure

This experiment involves 18 subjects, which are randomly assigned to two groups: the treatment group which uses CLARE, and the control group which does not. These two groups are further divided into three (3) study groups, with three (3) students in each. The experiment will start at around the half way into the semester and will last for two weeks. It is composed of four (4) phases: private outlining, private consolidation, group consolidation, and writing-up. Before the experiment starts, the subjects will be instructed about how the final output should be like.

- Private outlining: During this phase, both groups are required to read the assigned papers. As they do so, they write down key problems, ideas, comments, et al. The difference between what the two groups are doing is that the treatment group uses CLARE for this activity, while the control group does not.

- Private consolidation: At the beginning of this phase, the subjects within each study group will distribute to other group members what they have generated from the previous phase. Next, they will compare and contrast the ideas/questions from their fellow members with those of their own, and revise their individual review outlines accordingly. Again, the key difference between the two groups is that one does it through CLARE, and the other doesn't.

- Group consolidation: At this phase, subjects in each study group come together face-to-face to discuss major problems they have encountered while reading each other's work. In the control group, one member in each study group is assigned to record the discussion result. At the end, the discussion log, along with the revised reviews are given to a designated member, who is responsible for the final document. For the treatment group, this phase is conducted in front of a workstation; the discussion log will be recorded online. This session is fixed time for both groups, i.e., one class period.

- Writing-up: For both groups, the designated individual integrates the discussion results with individual reports into a coherent printed document, which, along with all intermediate artifacts from all group members, will be turned in to the researcher for evaluation.

5.2.3 Measures

Like the previous experiment, three types of measures will be collected on this experiment: outcome, process, and user satisfaction. The quality of the review paper will be graded by the instructor and two other experts on the subject. The evaluation criteria will include completeness, consistency, integration, level of support evidence for critiques, and quality of research questions raised. The support documents will also be taken into consideration during evaluation.

5.3 Comparison of the Two Experiments

The experiments described above illustrate two example learning tasks CLARE can support (See Section 4.6 for other potential CLARE usages). Despite their similarities, e.g., evaluation measures, these experiments differ in two important ways: subject and task. First, the two subject groups consist of two different levels of learners, i.e., graduate students versus undergraduates. Second, writing a review paper (task II) is a much more elaborate activity than generating a set of study questions (task I). The former, for example, requires not only longer time but also a higher level of integration of the learning materials on hand and what the learners already know.

The above two differences allow us to answer different empirical questions regarding the usefulness of CLARE. First, since CLARE was designed to support advanced learners, we expect that the graduate CLARE users have significantly more positive experience with the system than the undergraduate users do, whether or not their task performance reflects this difference. Second, one of CLARE's main features is its integrated support for the five levels of learning (see Section 3.3). The more sophisticated a learning task is, the more integration it requires of the five levels, and the more likely it benefits from using CLARE. Since writing a review paper is a much more sophisticated task than generating a set of study questions, we expect that the experiment involving the former be more likely to have a significant difference in the task performance.

6 Other Related Work

CLARE represents a confluence of several streams of research. Pedagogically, our work is based on the theory of meaningful learning [Aususel *et al.*, 1978; Novak and Gowin, 1984]. The illustrator function of CLARE is built on such theoretical principles as prior knowledge, subsumption, progressive differentiation, and integrative reconciliation. CLARE, like CSILE [Scardamalia and Bereiter, 1991], follows the constructivist tradition, which asserts the primacy of the social nature of knowledge, and therefore, of collaborative learning [Slavin, 1990]. However, it deviates from CSILE, Intermedia [Yankelovich *et al.*, 1988], and other similar tools in its emphasis and explicit representation of metaknowledge, and using it to help the learner organize and make sense of specific knowledge.

Representation is fundamental to design and computer science, in particular, artificial intelligence (AI) [Winograd and Flores, 1987]. The AI view of knowledge representation, however, is extremely varied, covering a broad range of ontological and epistemological schemes [Swaminathan, 1990], most of which have little to do human learning in classroom settings. The representational issues CLARE attempts to tackle bear more resemblance with what has come to be called "semi-structured" representations [Lee, 1990], for instance, IBIS [Kunz and Rittel, 1977; Conklin and Begeman, 1988], Toulmin's rhetorical model [Toulmin *et al.*, 1984]. In recent years, quite a few such schemes have been developed in the area of design rationale, e.g., [Lee and Lai, 1991; MacLean *et al.*, 1989; Conklin and Yakemovic, 1991]. More recently, we have begun to see similar approaches being used learning support environments (e.g., [Cavalli-Sforza *et al.*, 1992]). RESRA, in essence, is an extension to these approaches; its most important novelty is its treatment of such a scheme not merely as a medium for representing "other" things but itself as a subject of learning and exploration (i.e., metaknowledge).

Since CLARE is a hypertext-based collaborative learning environment, its design and implementation is shaped by previous work and experience in both of these areas, in particular, works such as [Conklin and Begeman, 1988; Halasz, 1988], and systems like NoteCards [Halasz *et al.*, 1987], gIBIS [Conklin and Begeman, 1988], and more recently, Aquanet [Marshall *et al.*, 1991]. CLARE's exploratory functions are a direct instantiation of the collaborative model provided by EGRET [Johnson *et al.*, 1992; Johnson, 1992b; Johnson, 1992a].

Empirically, our work was inspired by consistent encouraging findings from studies on the use of concept maps in science learning [Cliburn, Jr., 1990], and computer-mediated communication in augmenting traditional classroom-based learning [Hiltz, 1988]. Instead of providing the learner merely with an information access and sharing mechanism, as many CMC and hypertext systems do, we believe that the introduction of domain and/or meta structures and facilities to manipulate them can significantly enhance the usability of those learning support environments. CLARE represents our first attempt to testing this claim.

7 Research Plan

This research is divided into two phases: system development and evaluation. The first phase involves the design, implementation, and testing of CLARE prototype. The second phase requires conducting two experiments planned for evaluating the effectiveness of the CLARE system. Below are important milestones:

Dates	**Milestones**
January 25	Proposal Defense.
February 14.	Completing implementation of new functions on top of current CLARE prototype.
February 26.	Completing inhouse testing and pilot experiments.
March 20.	Completion of Phase I experiments.
April 16.	Completion of Phase II experiments.
June 30.	Completion of draft dissertation.
July 5.	Tentative date for dissertation defense.

Table 4: Important Research Milestones

References

[Ausubel, 1963] David P. Ausubel. *The psychology of meaningful verbal learning.* Grune & Stratton, 1963.

[Aususel *et al.*, 1978] David P. Aususel, J.D. Novak, and H. Hansesian. *Educational psychology: a cognitive view.* Hold, Rinchart and Winston, New York, 2nd edition, 1978.

[Baecker and et al, 1993] Ronald M. Baecker and et al. The user-centered iterative design of collaborative writing software. In *Proceedings of INTERCHI'93*, 1993. Submitted for review.

[Berger and Luckman, 1966] Peter Berger and Thomas Luckman. *The Social Construction of Reality: a Treatise in the Sociology of Knowledge.* Doubleday, 1966.

[Cavalli-Sforza *et al.*, 1992] Violetta Cavalli-Sforza, Gareth Gabrys, Alan M.Lesgold, and Arlene Weiner. Engaging students in scientific activity and scientific controversy. In Kishore Swaminathan, editor, *AAAI92-workshop*, pages 99–106, San Jose, 1992.

[Cliburn, Jr., 1990] Joseph W. Cliburn, Jr. Concept maps to promote meaningful learning. *Journal of College Science Teaching*, pages 212–217, feb 1990.

[Conklin and Begeman, 1988] Jeff Conklin and Michael L. Begeman. gIBIS: a hypertext tool for exploratory policy discussion. In *CSCW88*, pages 140–152, 1988.

[Conklin and Yakemovic, 1991] E. Jeffrey Conklin and K.C. Burgess Yakemovic. A process-oriented approach to design rationale. *Human-Computer Interaction*, 6(3,4):357–391, 1991.

[Cross, 1990] Alan G. Cross. *The Rhetoric of science.* Harvard University Press, 1990.

[Fisher, 1990] Kathleen M. Fisher. Semantic networking: the new kid on the block. *JRST*, 27(10):1001–1018, dec 1990.

[Freedman and Weinberg, 1990] D. P. Freedman and G. M. Weinberg. *Handbook of Walkthroughs, Inspections and Technical Reviews.* Little, Brown, 1990.

[Halasz *et al.*, 1987] Frank G. Halasz, Thomas P. Moran, and Randall H. Trigg. Notecards in a nutshell. In *Proceedings of ACM CHI+GI'87 Conference on Human Factors in Computing Systems and Graphics Interface*, pages 45–52, 1987.

[Halasz, 1988] F.G. Halasz. Reflections on Notecards: Seven issues for the next generation of hypermedia systems. *Communication of the ACM*, 31(7):836–852, 1988.

[Hiltz, 1988] Starr Roxanne Hiltz. Collaborative learning in a virtual classroom: Highlights of findings. In *Proceedings of ACM CSCW'88 Conference on Computer-Supported Cooperative Work*, pages 282–290, 1988.

[Johnson *et al.*, 1992] Philip Johnson, Dadong Wan, and Danu Tjahjono. EGRET design specification: Version 2.0. CSDL Technical Report 91-02, University of Hawaii Department of Information and Computer Sciences, jul 1992.

[Johnson, 1992a] Philip Johnson. An architectural perspective on EGRET. In *Proceedings of the 1992 CSCW Workshop on Tools and Technologies*, Nov 1992.

[Johnson, 1992b] Philip Johnson. Supporting exploratory CSCW with the EGRET framework. In *Proceedings of the 1992 Conference on Computer Supported Cooperative Work*, Nov 1992. Also published as Technical Report CSDL-92-01, University of Hawaii Department of Information and Computer Sciences.

[Kaplan and et al, 1992] Simon Kaplan and et al. Supporting collaborative software development with ConversationBuilder. In *CSCW92*, 1992.

[Knorr-Centina, 1981] K.D. Knorr-Centina. *The manufacture of knowledge: an essay on the constructivist and contextural nature of science.* Pergamon Press, Oxford, 1981.

[Kunz and Rittel, 1977] W. Kunz and H. Rittel. Issues as elements of information systems. Technical report, Institute of Urban and Regional Development, University of California at Berkeley, 1977.

[Lee and Lai, 1991] Jintae Lee and Kum-Yew Lai. What's in design rationale? *Human-Computer Interaction*, 6(3,4):251–280, 1991.

[Lee, 1990] Jintae Lee. SIBYL: a tool for managing group decision rational. In *CSCW90*, pages 79–92, 1990.

[MacLean *et al.*, 1989] Allan MacLean, Richard M. Young, and Thomas P. Moran. Design rationale: the argument behind the artifact. In *Proceedings of ACM CHI'89 Conference on Human Factors in Computing Systems*, pages 247–252, 1989.

[Marshall *et al.*, 1991] Catherine C. Marshall, Frank G. Halasz, Russell A. Rogers, and Jr. William C. Janssen. Aquanet: a hypertext tool to hold your knowledge in place. In *Proceedings of ACM Hypertext'91*, pages 261–275, 1991.

[Novak and Gowin, 1984] Joseph D. Novak and D. Bob Gowin. *Learning how to learn.* Cambridge University Press, 1984.

[Novak, 1990] J.D. Novak. Concept mapping: a useful tool for science education. *JRST*, 27(10):937–49, dec 1990.

[Quilian, 1967] Ross M. Quilian. Word concepts: a theory and simulation fo some basic sematic capabilities. *Behavioral Science*, 12:410–30, 1967.

[Roth and Roychoudhury, 1992] Wolff-Michael Roth and Anita Roychoudhury. The social construction of scientific concepts or the concept map as conscription device and tool for social thinking in high school science. *Science Education*, 76(5):531–557, 1992.

[Scardamalia and Bereiter, 1991] Marlene Scardamalia and Carl Bereiter. Higher levels of agency for children in knowledge-building: a challenge for the design of new knowledge media. *Journal of Learning*, 1(1):37–68, 1991.

[Schank, 1975] R. C. Schank. Conceptual dependency theory. In R. C. Schank, editor, *Conceptual Information Processing*, pages 22–82. North-Holland and Elsevier, Amsterdam and New York, 1975.

[Slavin, 1990] R. Slavin. *Cooperative learning theory and research*. Prentice-Hall, Englewood-Cliffs, NJ, 1990.

[Stallman, 1985] Richard Stallman. *The Emacs reference manual*. Free Softaware Foundation, 1985.

[Swaminathan, 1990] Kishore Swaminathan. *RA: a memory organization to model the evolution of scientific knowledge*. PhD thesis, University of Massachusetts at Amherst, 1990. also COINS Technical Report 90-80.

[Toulmin *et al.*, 1984] Steven Toulmin, Richard Rieke, and Allan Janik. *An introduction to reasoning*. Macmillan, New York, 2nd edition, 1984.

[Wan and Johnson, 1992] Dadong Wan and Philip Johnson. Supporting scientific learning and research review using COREVIEW. In *Proceedings of the 1992 AAAI Workshop on Communicating Scientific and Technical Knowledge*, jul 1992. Also published as Technical Report CSDL-92-03, University of Hawaii Department of Information and Computer Sciences.

[Wiil and Osterbye, 1990] U. Wiil and K. Osterbye. Experiences with hyperbase–a multi-user back-end for hypertext applications with emphasis on collaboration support. Technical Report 90-38, Department of Mathematics and Computer Science, University of Aalborg, Denmark., 1990.

[Winograd and Flores, 1987] Terry Winograd and Fernando Flores. *Understanding computer and cognition*. Addison-Wesley, Reading, Mass, 1987.

[Yankelovich *et al.*, 1988] Nicole Yankelovich, Bernard J. Haan, Norman K. Y. Meyrowitz, and Steven M. Drucker. Intermedia: the concept and the construction of a seamless information environment. *Computer*, jan 1988.