

Representational Effects in Asynchronous Collaboration: A Research Paradigm and Initial Analysis

Daniel Suthers, Ravikiran Vatrappu, Samuel Joseph, Nathan Dwyer and Richard Medina
Department of Information and Computer Sciences
University of Hawai'i at Manoa
suthers@hawaii.edu

Abstract

Researchers have argued that tools for online learning should provide representational support for the conceptual structure of a problem area in order to address issues of coherence and convergence and more effectively support collaborative knowledge construction. The study described in this paper sets out to investigate the merits of knowledge representations and of two alternative ways they may be related to discussion tools: embedded or linked. Analyses conducted to date suggest intriguing process and outcome differences to be investigated in future analyses. The paper also offers a methodological contribution: a paradigm for practical experimental study of asynchronous collaboration. Prior research has focused on face-to-face and synchronous collaboration due to the pragmatic problems of conducting asynchronous studies. It is crucial to understand how to support collaborative knowledge construction in asynchronous settings prevalent in online learning.

1. Introduction

The use of electronic media for online learning has expanded greatly in the past decade [1], yet too often implementations use pre-existing Internet technology to “deliver” conventional but ineffective pedagogical approaches, rather than adopting or inventing new technologies specifically designed to support effective approaches to learning. Decades of research on learning and instruction have shown the importance of learners’ active participation in expressing, testing, and revising their own knowledge (e.g., [2-4]). Therefore, electronic media should support such engagement, leveraging the computational medium’s strengths for education: its representational and analytic capabilities, its interactivity and networking support for collaboration.

Two sets of findings motivated the present research: (1) the impact of representational aids, such as

dynamic notations, knowledge maps, simulations, etc., on individual problem solving (e.g., [5-8]) and learning (e.g., [9-12]); and (2) the importance of social processes such as collaboration and mentoring to learning (e.g., [13-17]). Until recently, there has been a lack of research on how these techniques—representational tools and collaborative learning—may be constructively combined. Exceptions include [18], [19], and [20]. The limited comparative research available suggests that the form of representations used by learners during collaborative inquiry can lead to different forms of learning discourse. This effect has been shown both with representations that are constructed by learners during collaboration [20] and with representations used as a medium of discourse [21, 22].

A separate but related line of research on computer-mediated communication (CMC) has identified several problems related to typical discourse representations through which people communicate online (e.g. threaded discussion and chat). These problems include incoherence due to the violation of discourse conventions for topic maintenance [23] and lack of convergence, due to the intrinsically divergent representations used in threaded discussion [24]. The shared agreement or knowledge being constructed through the discourse is not made explicit by typical CMC tools, and hence it is difficult to find relevant contributions, place one’s own contribution in the relevant context, or quickly assess the outcome of the discourse [25, 26].

The fundamental problems are a lack of integration of discourse representations with other representations and a lack of explicit construction of the desired outcome of the collaboration, leading to weak support for online knowledge-building discourse. In response to these problems, Suthers [25] proposed better online support for artifact-centered discourse (discourse that makes reference to and is tightly integrated with visual or textual artifacts), and suggested that synergistic benefits may be obtained if these artifacts are also knowledge representations. That is, the evolving

knowledge representations become the artifacts under discussion in the CMC environment. The claim is that if each contribution to the discourse can be referenced to a component of the knowledge representation, coherence improves because comments are localized where they belong, and convergence improves because multiple contributions referencing a given topic are collected together. The knowledge representation will also serve as a summary of the status of the collaboration, available to learners and mentors to support reflection and assessment.

The present paper reports on the design of and initial results from a major experimental study of the benefits of such an approach, in which participants are enabled to construct an explicit representation of the topics and conclusions of the discourse itself as they engage in the discourse. Two forms of artifact-centered discourse are also compared in this study. Since our interest is online collaborative learning, which commonly includes a strong asynchronous component or "asynchronous learning networks" [1], we confronted the problem of experimentally studying asynchronous collaboration. A pragmatically viable methodological approach is also presented as a contribution in this paper.

2. Hypotheses and software designs

The engineering objective of this work is to improve online knowledge building environments. However, in exploring how this may best be accomplished we also address scientific objectives of understanding the role of representational tools in knowledge building processes.

2.1. Hypotheses

The following hypotheses capture the relationship between the engineering and scientific objectives. Knowledge building seeks systematicity, coherence, and convergence as participants engage in meaning making to extend their collective understanding [27]. The first hypothesis concerns the utility of explicit knowledge representations in this process.

H1: Knowledge construction is more effectively supported by environments that make conceptual relations explicit, because it is a reflective process that requires awareness of one's own conceptual understanding. The argument behind this hypothesis begins with the observation that communication media that are structured by discourse relations such as reply structure capture the historical development of discussion rather than its conceptual content, making it difficult to make contributions that move it forward [25, 26]. Explicit representations of conceptual

structure have the advantages that they encourage participants to clarify their thinking sufficiently to build these representations, make this thinking visible to others, provide resources for subsequent conversation, and can function as a "convergence artifact" that expresses the group's emerging consensus [20, 24, 25].

Even if this hypothesis were definitively accepted there remains the question of the relationship between the knowledge representations and the discourse that accompanies the creation of those representations. The next two hypotheses are alternative elaborations of H1, arguing for either maintaining the distinction between discourse and knowledge representations or combining the two.

H2: There should not be a rigid distinction between discourse and conceptual representations because the two are so tightly related. There are two versions of this hypothesis. The less radical version states that *discourse representations should be embedded in the conceptual representations because this will contextualize the discussion, facilitating ease of reference* (e.g., by simple attachment of notes to the objects to which they refer). Suthers [25] called this "embedded artifact-centered discourse" because the discourse is embedded in the artifact under discussion.

A more radical version of hypothesis H2 states that *knowledge lives in interaction: it is not possible to separate them; therefore tools for collaboration should not attempt to do so.* According to this view, it is not possible to dichotomize our interactions by saying "that is discussion" and "that is the knowledge that is the product of the discussion." For example, contributions in the discussion might be reinterpreted, elaborated, and brought to bear on other situations in a manner that elevates them to part of "what we know." When is the line from discourse to knowledge crossed? The argument for H2 states that since the two cannot be distinguished, the representational medium should not force this distinction, but should instead provide a collection of representational resources with and through which participants can interact in a discursive mutual construction of knowledge.

One could argue that this literal translation of the nature of knowledge to a recommendation for the design of tools for collaboration is a category mistake, confusing *knowledge* with *conceptual representations*. An argument about the nature of knowledge need not necessarily be literally mirrored in the representational resources we provide. Also, even if designers provide separate "knowledge" and "discourse" representations, users may not respect this distinction. Collaborators will distribute their interaction across all mutable media [28]. Knowledge may yet live in interaction regardless of how this interaction is distributed across

representational media. This point leads us to the alternative hypothesis.

H3: The distinction between discourse and conceptual representations should be reflected in the tools provided because each has a different structure, and separate tools enable optimization of the representations for each. In [25], arguments were made for “linked artifact-centered discourse” in which discourse media such as threaded discussions would be maintained separately from knowledge representations or other disciplinary representations being discussed, but referential links would be made to the relevant parts of the latter artifacts. A linked approach attempts to maintain one major advantage of the embedded approach, the contextualization of contributions, while addressing deficiencies and adding other advantages. When given its own representation, the chronological reply structure of the discourse may be maintained, and discussions that rise above particular objects in the representations are more natural. Yet, explicit “linking” or reference of discourse contributions to conceptual objects resolves some of the incoherence resulting from the violation of contiguity of related discourse contributions that is so common in electronic media. Recently these ideas have been explored in software implementations by others (e.g., [29]) as well as our own.

2.2. Software environments

These hypotheses led us to construct three software environments (Figures 1-3). All three of the environments have an “information viewer” on the left in which materials relevant to the problem are displayed. This information viewer functions as a simple web browser, but presentation of materials is constrained as discussed in the next section.

All three environments have a shared workspace or “information organizer” on the right hand side in which participants can share and organize information they gather from the problem materials as well as their own interpretations and other ideas. The three environments differ on the nature of the “information organizer,” as described below. Changes made to the workspace by each participant are propagated to other participant’s displays of the same workspace under a protocol to be discussed in the next section. In all three environments, mutual awareness of participants’ activity is also supported as follows: yellow circles are used to mark information posted by the user of the environment but not yet “read” by his or her partner, while red triangles are used to mark new information from the user’s partner that he or she has not yet read.

2.2.1. Text Condition. The shared workspace in the “Text-only environment,” or “Text” condition for

short, is a conventional threaded discussion tool (Figure 1). This environment functions as the control condition for testing the above hypotheses, since the workspace only provides explicit support for representation of discussion structure (subject headings and reply relations).

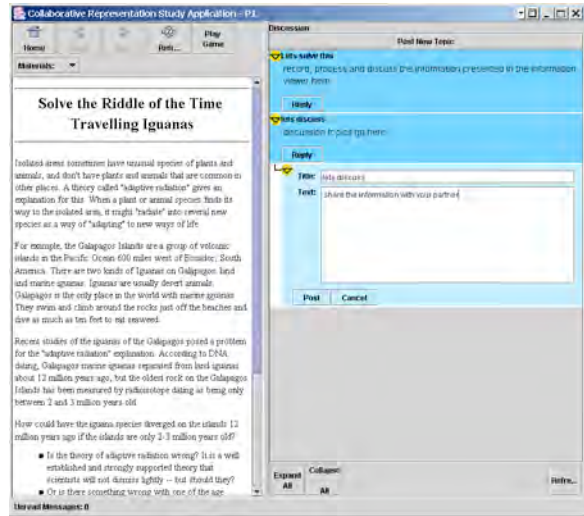


Figure 1. “Text” environment

2.2.2. Graph Condition. The shared workspace in the “Graph-only environment,” or “Graph” condition for short, consists of an integrated node-and-link graphing tool in which one can express both conceptual structure (relations of evidence between data and hypothesis objects) and commentary (notes that can be free-floating or attached to specific objects).

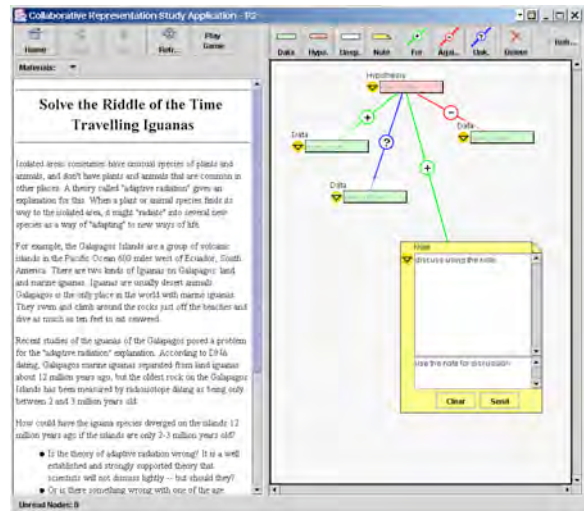


Figure 2. “Graph” environment (embedded discussion)

Like the Mixed environment (described next), the Graph workspace includes tools for constructing conceptual objects under a simple typology relevant to the task of identifying the cause of a phenomenon, including data (green rectangles, for empirical information) and hypotheses (pink rectangles, for postulated causes or other ideas). There are also linking tools for constructing “for” and “against” relations between other objects, visualized as green links labeled “+” and red links labeled “-” respectively. Two other types of objects, “unspecified” and “note,” and an “unknown” link are also provided for flexibility. The note object supports chronologically sequential accumulation of comments contributed by both participants (i.e., unthreaded discussion).

Our graph workspace reflects the weak form of H2, which claims that discussion is best supported in a contextualized manner, embedded in the conceptual representation for ease of reference. This configuration also has the advantage of simplicity in the sense that there is one workspace. The stronger version of H2, in which one cannot separate knowledge and discourse, is insufficiently reflected in the Graph software because the presence of notes that support discussions among conceptual representations still dichotomizes the elements of interaction. Another line of work (unpublished, but informed by [30]) is addressing the problem of how to provide more flexible media for interaction through representations.

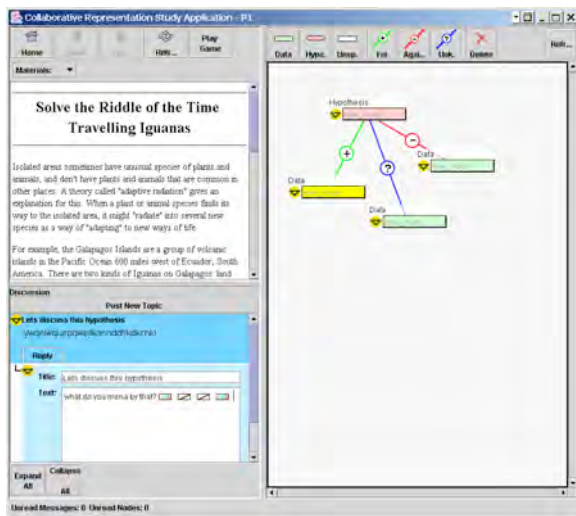


Figure 3. “Mixed” environment (linked discussion)

2.2.3. Mixed Condition. The shared workspace of the “Mixed” condition includes both a threaded discussion tool and a graphical tool for representing conceptual structure (again, relations of evidence between data and hypotheses). There are no embedded notes in the

graph. However, one can embed references to graph objects in the discussion messages simply by clicking on the relevant graph object while composing the message. The references show up as small icons in the message (Figure 3). When the reader selects the icon, the corresponding object in the graph will be highlighted, indicating the intended referent. This environment is motivated by H3, which claims that separate representations are needed to optimize support for discussion and knowledge, but that they should be logically “linked” for referential purposes.

2.3. Experimental Design

H1 is tested by comparing performance of users of the Text environment to performance of users of the Graph and Mixed environments. H2 and H3 are tested by comparisons of performance with the Graph and Mixed conditions to each other. The present paper reports only initial analyses of the rich data we collected (described later), and does not definitively resolve these issues.

3. A protocol for experimental study of quasi-asynchronous collaboration

The majority of experimental studies of computer-mediated communication have been undertaken in synchronous collaboration settings, while a significant portion of applications of computer-mediated communication to online learning are primarily asynchronous. Based on personal communication with other researchers, the first author concluded that a major reason for the lack of studies of asynchronous collaboration is logistical: it is easier to conduct a study in which participants come to the laboratory for one session rather than a study in which participants must return to the controlled setting at different times repeatedly over a period of time. In the latter situation, the experimenter must be concerned with a potentially higher attrition rate (a significant amount of work can become useless if a participant fails to show up for the final session), and with whether participants would engage in other activities between sessions that invalidate the assumptions of a controlled design. Faced with these challenges and concerns, we designed a study protocol that simulates many of the properties of asynchronous communication while still enabling us to conduct sessions with participants in the laboratory at the same time.

The fundamental criterion was that there be no particular timing constraint between the actions of participants (e.g., waiting for the participant’s action before being able to continue one’s own work), nor temporal affordances to be exploited in a synchronous

manner (e.g., sending a message and expecting an immediate reply). A second aspect of asynchronous work that we sought to simulate (albeit necessarily less faithfully) is that one might stop working on a problem for a while, do something else, and then return to the work. We achieved these desiderata through a study protocol in which (1) participants took occasional “breaks” from their work to play a computer game, and (2) the work of the other participant became available only after these breaks. We discuss this protocol in further detail below.

3.1. Chunking of materials and “breaks”

The study materials were divided into six sets of materials. In each set, a participant was presented with four pages, each containing a short article. The contents of these pages will be described later. Participants were expected to work with the material presented in the four pages (updating the shared workspace as they deemed appropriate). When done, they could not obtain the next set of materials until they had “taken a break” by playing a computer game, a Java version of Tetris™, for 1-5 minutes (pilot studies showed that longer breaks made the session excessively long). Tetris™ was chosen for its familiarity and because it presents a perceptual motor activity quite different from the cognitive task of the study, in this sense presenting a “break” from the primary task. Paper-based activities were also considered, but rejected because we wanted to automate the timing and logging of breaks.

3.2. Protocol for workspace updates

The actions of each participant in the shared workspace were not displayed immediately in the other participant’s workspace. Instead, these actions were queued and displayed in the receiving participant’s workspace when she returned from the “break.” That is, as a person worked, the actions of that person were sent to the other participant’s client application, but queued rather than displayed. When a participant “resumed work,” all of the currently queued actions on that client were displayed. Conflicts that might arise when both participants edited the same object were resolved through *operational transformations* [31]. Operations O_a enacted by client A and O_b enacted by client B are transformed into O_a' and O_b' such that $O_a'(O_b) = O_b'(O_a)$. As a result, clients A and B are guaranteed to converge on the same state. The delayed updating protocol simulates one aspect of the experience of asynchronous collaboration: a participant sees what one’s partner has done upon returning to a workspace after a period of time. It excludes the possibility of synchronous “conversation” in which one

participant posts a message in the workspace and receives an immediate reply.

Pilot studies suggested that this protocol as just stated would be a little too strict. One participant would sometimes fall far behind another, who was wondering whether any work was being done in the workspace. Also, we recognized that in an asynchronous environment sometimes two people are working at the same time, and it is possible to get updates by refreshing the workspace with respect to a server. To address these concerns we cautiously introduced a “refresh” feature that enables one to get all updates to that point in time. We were concerned that, upon discovering this feature, participants might use it to engage in synchronous interaction by alternating between posting messages and refreshing the workspace while waiting for a reply. However, in our pilot studies and in the actual study itself, this did not happen very often. Participants used the refresh feature primarily at the end of the session when one person finished first and was waiting for the other person to finish their work (as the instructions required that they come to a final conclusion based on material they had shared with each other).

In order to assess the extent to which this protocol simulates asynchronous interaction, we compared “true” asynchronous interaction to our protocol on Clark & Brennan’s “grounding constraints” [32], well known dimensions for analyzing properties of communication media. The comparison shows that asynchronous interaction as enacted in asynchronous learning networks (ALN) and our quasi-asynchronous protocol (QAP) provide or fail to provide exactly the same grounding constraints, and are therefore equivalent according to these constraints (Table 1).

Table 1. Grounding constraints in ALN and QAP interactions

| | ALN | QAP |
|---|-------|-------|
| Copresence: <i>A and B share the same physical environment.</i> | False | False |
| Visibility: <i>A and B are visible to one another.</i> | False | False |
| Audibility: <i>A and B communicate by speaking.</i> | False | False |
| Contemporality: <i>B receives at roughly the same time as A produces.</i> | False | False |
| Simultaneity: <i>A and B can send and receive simultaneously.</i> | False | False |
| Sequentiality: <i>A's and B's turns cannot get out of sequence.</i> | False | False |
| Reviewability: <i>B can re-view A's messages.</i> | True | True |
| Revisability: <i>A can revise message for B (edit before sending)</i> | True | True |

Copresence is nonexistent in either protocol, although in both cases participants do share a workspace. Although we can imagine situations under which B receives the moment after A produces, *Contemporality* is the exception under both ALN and QAP. One can argue that A and B can both type messages at the same time, but the *Simultaneity* constraint is concerned with whether one can receive and process a communication *while* producing one. The important point is that any argument applies equally to both ALN and QAP. *Sequentiality* can be met in some ways, but not others; again, the arguments apply equally to both ALN and QAP. *Reviewability* and *Revisability* are not supported by face-to-face (spoken) interaction.

We are not naïve enough to claim that the quasi-asynchronous protocol produces a situation literally identical to ALN. It is interesting that Clark's constraints do not capture the ways in which our study protocol differs from "real" online collaboration. These differences include the time-span of interaction (possibly spread over days in ALN rather than a few hours, providing time to think about a problem between sessions) and the knowledge in QAP that one's partner is present in the same building working on the same problem at the same time (which may influence participants even though they cannot take advantage of this communicatively). We leave extension of Clark & Brennan's model to capture these aspects for future work.

4. Methods

Most of the substantial experimental design issues have already been discussed. In this section we summarize the remaining aspects of experimental method.

4.1. Participants

Pairs of participants were recruited from introductory courses in the College of Natural Sciences at the University of Hawai'i. Participants were paid US\$50 each for participating in the study. We recruited participants in pairs of acquaintances so as to eliminate the social awkwardness of interaction between persons who do not know each other (found to be problematic in our previous work).

Excluding pilot studies and disqualified sessions, we conducted a total of 30 experimental sessions involving 30 pairs or 60 participants. There were 10 pairs of participants (20 participants) for each of three treatment groups: Text, Graph and Mixed.

Female-female, female-male and male-male pairs were assigned to treatment groups in a gender-balanced

manner, because previous studies showed that gender pairing substantially influenced the style of interaction. We verified that the groups were randomly balanced as to age and grade point average, and that none of the participants had prior experience with the study problem.

4.2. Materials

4.2.1. Topics. The study presented participants with "science challenge" problems, consisting of relatively recent or ongoing issues in science and public health. The "Riddle of the Time Traveling Iguanas" problem (resolving a discrepancy in the dating of speciation of Galapagos iguanas) was used as a "warm-up" exercise with which participants could become familiar with the software and collaborating with each other through that software. The "Protect the Islanders from the Muscle- and Mind-killers" problem challenged participants to identify the cause of a disease on the island of Guam known as ALS-PD. In part because it shares symptoms with Alzheimer's and Parkinson's diseases, ALS-PD has been under investigation for 50 years. Only recently have investigators converged on both a plausible disease agent (a neurotoxic amino acid in the seed of the Cycad tree) and the vector for introduction of that agent into people (native Guamanians' consumption of fruit bats that eat the seed). Over the years numerous diverse hypotheses have been proposed and an even greater diversity of evidence of varying types and quality explored. These facts along with the relative obscurity of the problem make it a good problem to use when one wants participants to grapple with interpretation of multiple explanations and ambiguous data.

4.2.2. Organization. Source materials were provided in the form of short articles or information pages, typically consisting of one to two brief paragraphs and an image. Each article was designed to provide one key item of information relevant to the generation or evaluation of a hypothesis. The remaining information in a given article elaborated on this item or provided tangentially related "distractor" information. We prepared the material to provide evidence both for and against six major hypotheses. In some cases, the information needed to draw a conclusion was distributed across several articles.

As noted before, the articles presented to a given participant were clustered into six groups of four articles. Each participant received a different sequence of articles, although there was some overlap between both the articles given to participants and the information in non-identical articles. We used a classic paradigm in studies of group problem solving: information was distributed across participants such

that a participant relying only on information he or she directly received would come to a suboptimal conclusion. For example, one participant initially received evidence for aluminum as a disease agent and later received evidence against genetic causes, while the other participant received evidence for genetic causes and later received evidence against aluminum. Information sharing between participants was required in order for either participant to reject these hypotheses and identify the most complex explanation concerning bats as a vector for the toxin.

4.3. Procedure

After signing of consent forms, participants filled out a demographic survey. They were then introduced to the software and format of the sessions through a standardized set of instructions and demonstrations designed to be as equivalent as possible across all conditions.

Participants were then led to their respective stations in different rooms from each other, and began work on a "warm-up" problem, the Galapagos iguanas, to familiarize themselves with the software. After a maximum of 30 minutes of work on the warm-up problem, participants were instructed to halt work and begin work on the main problem, Guam ALS-PD. Participants were given up to 120 minutes to work through all of the information available for this problem. The update protocol described earlier was applied during these sessions.

At the conclusion of their problem solving session, each participant working alone was given up to 30 minutes to write an essay on the hypotheses that were considered, the evidence for and against these hypotheses, and the conclusion reached. The online environment remained available to each participant during the essay writing, but there was no further communication between participants.

Debriefing included administration of a usability questionnaire, followed by informal discussion with the experimenter of software usability and strategies used during the session. One week after the experimental session, each participant was required to complete the online posttest before payment was sent.

4.4. Data collection

Demographic information was collected through a survey and by obtaining SAT scores and Grade Point Averages from the University (with participants' permission).

Process data was collected through two primary means. First, the Morae™ video recording system was used to capture both the computer screen and a webcam sized image of each participant as digital

video. Second, our software was designed to generate complete logs of all the events at each client workstation. These events included message and graph object creation, edits, moves, and read events, whether generated by the local or remote participant.

Post-session data included the essay and usability questionnaire elicited immediately after the session, and the posttest elicited one week later, as previously discussed.

5. Results and discussion

An ambitious program of analysis is planned for this data, ranging from analyses of process and outcome data based on the "coding and counting" techniques of experimental psychology to qualitative analyses including a micro-analysis of how knowledge construction and intersubjective meaning-making is accomplished in interaction through and via appropriation of affordances of the software, as discussed in [33, 34].

Our initial analyses, reported in this paper, included a diversity of methods intended to obtain an overview of our data. These included analysis of the usability questionnaire, an exploratory examination of the video data to identify recurring issues in the sessions themselves, quantitative analyses conducted to pursue hypotheses raised by the video analysis as well as to examine some basic parameters of the sessions, scoring of the posttest, and a preliminary examination of the hypotheses mentioned in the essays. The in-depth analyses required to fully evaluate the hypotheses motivating this work have not yet been conducted. We describe each of the completed analyses in turn below.

5.1. Usability results

Quantitative analysis of the usability instrument verified that there was no significant difference across groups in participants' satisfaction with the instructions and software demonstration given by the experimenter. Analysis of questions pertaining to the software itself yielded a significant difference in satisfaction: Graph received the lowest subjective satisfaction scores and Text the highest. Questions dealing with management of layout of the graphical representation contributed strongly to this result. Examination of comments confirmed that Graph and Mixed received more negative comments, particularly with respect to screen clutter. Undo was the most requested feature. Participants also wanted their contributions to be distinguished from those of their partner.

5.4. Exploratory analysis of session data

Several sessions from each condition were skimmed with the video tool, and log data examined where needed for more precise determination of events. Anticipating a future focus on hypothesis formation as part of our analysis of knowledge construction, our exploratory analysis focused on the creation, discussion, modification, and referencing of hypotheses.

Our most salient observation concerned the timing and handling of hypotheses. In Graph and Mixed conditions, participants considered the first hypothesis much earlier than in the text condition. There seemed to be little discussion in the Text condition compared to the other two. For example, many messages were created simply by copying and pasting articles to be shared with the partner. Exploratory analysis also suggested that there was little subsequent referencing to hypotheses in the Text condition. In general, substantial discussion of hypotheses in the Text condition took place late in the session. These observations prompted us to conduct quantitative analysis of the time to create the first hypothesis.

5.5. Quantitative activity analyses

A test of the time to consider the first hypothesis was motivated by the exploratory analysis. Do the representational tools used differ in how early they encourage participants to state a hypothesis? The analysis measured the time in seconds for each individual participant to introduce the first hypothesis in any medium. Significant results ($p < .0002$) were obtained favoring early creation of hypotheses in Graph (618 sec.) and Mixed (1162 sec.) as compared to Text. (2433 sec.) This result is consistent with the representational guidance effect demonstrated by [20]. Its significance is that early introduction of a hypothesis can lead to evaluation of subsequent data in terms of this hypothesis.

5.2. Posttest results

In the ALS-PD problem, relevant information about the different possible causes of the medical condition was embedded within articles that included much corollary information not directly related to the condition. The posttest contained two classes of questions. *Memory* questions were based purely on corollary information, while *integrative* questions derived from information that would be used in reasoning about aspects of the medical condition and was distributed across clusters of articles and participants. We reasoned that information more intimately tied to the complex knowledge structure that the participants formed while solving the problem would be slower to fade from memory and be easier to

recall. In addition, items within that structure would be more likely to be the subject of collaborative discussion, and thus provide a social association that would increase long term retention. Thus, if performance improved on integrative questions in one condition, that would indicate better collaboration in that condition. Distractor responses to the questions were designed to discriminate different kinds of errors, e.g. of recall versus reasoning.

No significant differences were found in total scores across conditions, nor when considering memory or integrative questions. The posttest may have suffered from insufficient power, as it is difficult to construct many questions of the nature just described. However, we did find significant differences in the types of errors made on integrative questions: the Mixed condition made more logical reasoning errors than the Graph condition. No compelling explanation has occurred to us for this result. (Recall that GPA and SAT scores are equivalent across groups.)

5.3. Essay hypotheses

We compared the hypotheses mentioned in the essays across treatment conditions to assess differences in (1) convergence, as measured by whether pairs come to mutual agreement on the cause for the disease, and (2) quality of solution, as measured by whether individuals identified the optimal "bats as vector for toxin from cycads" hypothesis. Two analysts conducted this analysis, obtaining similar results and agreeing to select a final analysis by consensus. The results from this analysis are shown in Table 1.

Table 1. Conclusions selected in essays

| | Pair agreement | Bat hypothesis |
|----------|----------------|----------------|
| Text | 4/10 | 5/20 |
| Graph | 8/10 | 2/20 |
| Mixed | 2/10 | 2/20 |
| χ^2 | $p \leq 0.025$ | $p \leq 1.0$ |

From the standpoint of (1) convergence, Graph seems to be advantageous. We speculate that having a single visually oriented workspace (which was available during the essay writing) makes it easier for participants to see and be reminded of their work together, leading to convergence in the contents of the essay. The dual workspaces of Mixed provide more variation in strategies for using the workspaces while writing the essays, increasing the possibility that members of a pair will look at different material. The additional cognitive load of using two representations may have also been a factor in Mixed.

From the standpoint of (2) quality of solution (under an admittedly crude measure), the difference is not significant under χ^2 . The slight difference might reflect the tendency of the Text participants to simply cut and paste entire articles into their text messages and leave discussion for the end, when evidence for the bat hypothesis was salient in the final set of messages available in the sequential representation.

6. Summary and conclusions

Along with others, we have argued that tools for online learning should provide representational support for conceptual structure in order to address issues of coherence and convergence and more effectively support collaborative knowledge construction. The study described in this paper set out to investigate the claimed merits of conceptually oriented representations and of two approaches to the relationship between conceptual and discourse representations: embedded or linked. Analyses conducted to date did not yield differences on memory recall, but do suggest other intriguing process and outcome differences to be investigated in future analyses. The initial process analysis focused on the creation and discussion of hypotheses. A representational effect was identified: users of a knowledge representation tool that includes primitives for hypotheses are more likely to state hypotheses early in their sessions, and therefore have more opportunity to discuss these hypotheses than users of the threaded discussion tool. These latter participants tended to simply record the literal text of the information articles, and not discuss hypotheses until later in the session. Examination of the final conclusions stated in the essays shows that pairs of users of the graphical representation were more likely to converge on the same hypothesis. A great deal of further analysis is planned, especially focusing on the ways in which participants appropriate the affordances of the media to engage in collaborative knowledge construction [33].

The paper also offers a methodological contribution: a paradigm for practical experimental study of asynchronous collaboration. Prior research on the effects that representational tools have on collaborative learning has focused on face-to-face and synchronous collaboration. Little research has studied representational effects in a controlled manner due to the pragmatic problems of conducting asynchronous studies. The study described in this paper extends this line of research to asynchronous settings. It is crucial to understand how to support collaborative knowledge construction in such settings due to the prevalence of asynchronous approaches to online learning.

7. Acknowledgments

David Burger and Niels Pinkwart have contributed to the design of this study and the implementation of the software on which it is based. This work was supported by the National Science Foundation under award 0093505. Any opinions, findings, and conclusions or recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of the National Science Foundation.

8. References

- [1] F. Mayadas, "Asynchronous Learning Networks: A Sloan Foundation Perspective.," *Journal of Asynchronous Learning Networks*, vol. 1, pp. www.aln.org/alnweb/journal/jaln_issue1.htm#mayadas, 1997.
- [2] M. T. H. Chi, M. Bassok, M. Lewis, P. Reimann, and R. Glaser, "Self-explanations: How students study and use examples in learning to solve problems," *Cognitive Science*, vol. 13, pp. 145-182, 1989.
- [3] D. N. Perkins, D. Crismond, R. Simmons, and C. Unger, "Inside understanding," in *Software Goes to School: Teaching for Understanding with New Technologies*, D. N. Perkins, J. L. Schwartz, M. M. West, and M. S. Wiske, Eds. New York: Oxford University Press, 1997, pp. 70-87.
- [4] M. Scardamalia and C. Bereiter, "Computer support for knowledge building communities," *Journal of the Learning Sciences*, vol. 3, pp. 265-283, 1994.
- [5] K. Kotovsky and H. A. Simon, "What makes some problems really hard: Explorations in the problem space of difficulty," *Cognitive Psychology*, vol. 22, pp. 143-183, 1990.
- [6] J. H. Larkin and H. A. Simon, "Why a diagram is (sometimes) worth ten thousand words," *Cognitive Science*, vol. 11, pp. 65-99, 1987.
- [7] L. R. Novick and C. E. Hmelo, "Transferring symbolic representations across nonisomorphic problems," *Journal of Experimental Psychology: Learning, Memory, and Cognition*, vol. 20, pp. 1296-1321, 1994.
- [8] J. Zhang, "The nature of external representations in problem solving," *Cognitive Science*, vol. 21, pp. 179-217, 1997.
- [9] K. Koedinger, *On the design of novel notations and actions to facilitate thinking and learning*, Proceedings of International Conference on the Learning Sciences ed. Charlottesville, VA: Association for the Advancement of Computing in Education, 1991.
- [10] J. Novak, "Concept mapping: A useful tool for science education," *Journal of Research in Science Teaching*, vol. 27, pp. 937-49, 1990.
- [11] M. Scardamalia, C. Bereiter, C. Brett, P. J. Burtis, C. Calhoun, and N. Smith Lea, "Educational applications of a networked communal database," *Interactive Learning Environments*, vol. 2, pp. 45-71, 1992.
- [12] J. Snir, C. Smith, and L. Grosslight, "Conceptually enhanced simulations: A computer tool for science teaching," in *Software Goes to School: Teaching for Understanding with New Technologies*, D. N. Perkins, J. L. Schwartz, M. M.

- West, and M. S. Wiske, Eds. New York: Oxford University Press, 1995, pp. 106-129.
- [13] A. L. Brown and J. C. Campione, "Guided discovery in a community of learners," in *Classroom Lessons: Integrating Cognitive Theory and Practice*, K. McGilly, Ed. Cambridge: MIT Press, 1994, pp. 229-270.
- [14] J. Lave and E. Wenger, *Situated learning: Legitimate peripheral participation*. Cambridge: Cambridge University Press, 1991.
- [15] M. Scardamalia and C. Bereiter, "Higher levels of agency for children in knowledge building: A challenge for the design of new knowledge media," *The Journal of the Learning Sciences*, vol. 1, pp. 37-68, 1991.
- [16] G. Stahl, *Collaborating with Technology: Mediation of Group Cognition*. Cambridge, MA: MIT Press, in press.
- [17] N. Webb and A. Palincsar, "Group processes in the classroom," in *Handbook of Educational Psychology*, D. Berliner and R. Calfee, Eds. New York: Simon & Schuster Macmillian, 1996.
- [18] P. Dillenbourg, "Designing biases that augment socio-cognitive interactions.," in *Barriers and Biases in Computer-Mediated Knowledge Communication-and How They May Be Overcome*, R. Bromme, F. W. Hesse, and H. Spada, Eds. Dordrecht: Kluwer, 2005.
- [19] J. T. Roschelle, "Designing for cognitive communication: Epistemic fidelity or mediating collaborative inquiry?," *the Arachnet Electronic Journal of Virtual Culture*, vol. 2, 1994.
- [20] D. Suthers and C. Hundhausen, "An Empirical Study of the Effects of Representational Guidance on Collaborative Learning.," *Journal of the Learning Sciences*, vol. 12, pp. 183-219, 2003.
- [21] M. J. Baker and K. Lund, "Flexibly Structuring the Interaction in a CSCL environment," presented at Proceedings of the EuroAIED Conference, Edições Colibri, Lisbon, 1996.
- [22] M. Guzdial and C. Hmelo, "Integrating and Guiding Collaboration: Lessons Learned in Computer-supported Collaboration Learning Research at Georgia Tech," presented at CSCL'97, Toronto, Ontario, 1997.
- [23] S. C. Herring, "Interactive coherence in CMC," presented at Proceedings of the 32nd Hawai'i International Conference on the System Sciences (HICSS 32), Maui, Hawai'i, 1999.
- [24] J. Hewitt, "Beyond Threaded Discourse," *International Journal of Educational Telecommunications*, vol. 7, pp. 207-221, 2001.
- [25] D. Suthers, "Collaborative Representations: Supporting Face to Face and Online Knowledge-building Discourse," presented at Proceedings of the 34th Hawai'i International Conference on the System Sciences (HICSS-34), Maui, Hawai'i, 2001.
- [26] M. Turoff, S. R. Hiltz, M. Bieber, J. Fjermestad, and A. Rana, "Collaborative discourse structures in computer mediated group communications," *Journal of Computer Mediated Communication*, vol. 4, pp. <http://jcmc.huji.ac.il/>, 1999.
- [27] G. Wells, *Dialogic inquiry: Toward a sociocultural practice and theory of education*. New York: Cambridge University Press, 1999.
- [28] D. Suthers, L. Girardeau, and C. Hundhausen, "Deictic Roles of External Representations in Face-to-face and Online Collaboration," presented at Proceedings of the International Conference on Computer Support for Collaborative Learning 2003, 2003.
- [29] M. Mühlfordt and M. Wessner, "Explicit Referencing in Chat Supports Collaborative Learning," presented at Proceedings of the Computer Supported Collaborative Learning 2005 (CSCL2005), Taipei, Taiwan, 2005.
- [30] N. Dwyer and D. Suthers, "A Study of the Foundations of Artifact-Mediated Collaboration," presented at Proceedings of the Computer Supported Collaborative Learning 2005 (CSCL2005), Taipei, Taiwan, 2005.
- [31] C. Sun, X. Jia, Y. Zhang, and Y. Yang, "A generic operation transformation scheme for consistency maintenance in real-time cooperative editing systems," presented at Proc. of ACM Conference on Supporting Group Work, 1997.
- [32] H. H. Clark and S. E. Brennan, "Grounding in communication," in *Perspectives on Socially Shared Cognition*, L. B. Resnick, J. M. Levine, and S. D. Teasley, Eds.: American Psychological Association, 1991, pp. 127-149.
- [33] D. Suthers, "Collaborative Knowledge Building through Shared Representations," presented at Proceedings of the 38th Hawai'i International Conference on the System Sciences (HICSS-37), Wakoloa, Hawai'i, 2005.
- [34] D. Suthers, "Technology Affordances for Intersubjective Learning: A Thematic Agenda for CSCL," presented at Proceedings of the Computer Supported Collaborative Learning 2005 (CSCL2005), Taipei, Taiwan, 2005.