

A QUALITATIVE ANALYSIS OF COLLABORATIVE KNOWLEDGE CONSTRUCTION THROUGH SHARED REPRESENTATIONS

DANIEL D. SUTHERS

*Department of Information and Computer Sciences
University of Hawai'i at Manoa
1680 East West Road, POST 317
Honolulu, HI 96822, USA
suthers@hawaii.edu*

Abstract. This paper takes one step towards addressing the question of how activity mediated by shared representations—notations that are manipulated by more than one person during a collaborative task—might constitute knowledge construction activity, and how the shared representations are appropriated for this purpose. The primary contribution of this paper is a methodology for qualitative analysis of activity in a workspace built on the concept of “uptake”: how participants take up and build on prior contributions. By examining patterns of information uptake we can see ways in which participants’ activities constitute an intersubjective cognitive activity distributed across persons and the representations they are manipulating. The analysis is conducted in three phases: identification of acts of media manipulation, identification of information uptake relations between these acts, and application of appropriate theoretical perspectives to interpret the uptake graph. The uptake graph is intended to minimize theoretical commitments in order to support eclectic analysis. Several theoretical perspectives on how representations might mediate collaborative knowledge construction are surveyed to identify the kinds of events we would look for as evidence of knowledge construction via a representational medium. The methodology is applied to data from a prior study in which participants collaborated via a graphical representation as well as a verbal “chat” tool. These case examples illustrate how the methodology uncovered argumentation and knowledge construction conducted solely through the graph workspace.

Keywords: computer supported collaborative learning, knowledge construction, representational affordances, qualitative analysis, online synchronous interaction

1. Introduction

This paper is part of a larger research agenda that seeks to identify how collaborating learners use software-based knowledge representations, and consequently how to design such tools to more effectively support collaboration. The value and role of external representations (inscriptions) in mediating collaborative inquiry was demonstrated by Roschelle (1994) in his study of face-to-face collaboration, now considered to be one of the seminal papers of computer supported collaborative learning (Stahl, Koschmann & Suthers, 2006). Through an analysis that

looked carefully at students' conversation and use of gestures, Roschelle showed that the various representations available to collaborating students served more as resources for their meaning-making conversations ("mediating collaborative inquiry") than as a means of communicating expert knowledge ("epistemic fidelity").

Given the importance of co-present cues for grounding communication (Clark & Brennan, 1991) and other issues with interaction at a distance such as reduced shared context and the difficulty of reference (Olson & Olson, 2000), we might expect that shared representations have reduced utility online. However, the possibility has also been raised that computer media can go "beyond being there" (Hollan & Stornetta, 1992) in enabling new forms of interaction: perhaps the shared representations will be used in unexpected ways not limited to imitating face-to-face interaction. Our research agenda seeks to examine how participants appropriate the affordances of computer media for communication and collaborative knowledge construction¹ in ways that we can leverage for improved designs (Suthers, 2006).

The work reported in this paper was motivated by questions arising out of our prior work. In an earlier study (Suthers & Hundhausen, 2003), we found that the form of representation provided to learners can influence their face-to-face collaborations in educationally relevant ways, and termed this phenomenon "representational guidance." In a follow-up study, we decided to investigate whether the role of representations changes in online synchronous collaboration. A study partially reported in Suthers, Girardeau & Hundhausen (2003) documented differences in how face-to-face and online collaborators made use of a shared evidence map (a graphical representation of evidence relations). Face-to-face collaborators were more likely to negotiate changes to the evidence map by discussing the map and potential changes before undertaking those changes. The evidence map was either itself the topic of discussion or used as a deictic resource² during discussion of the problem domain. In contrast, online collaborators were more likely to propose new ideas by first expressing them in the evidence map, even though they had a textual chat tool within which they could have discussed these changes first. Discussion of the changes in chat was usually limited to brief confirmation dialogues. They treated the evidence map as a medium *through* which collaboration took place as well as its object. Furthermore, in contrast to face-to-face users of shared representations, gesture was almost never used online. Verbal deixis³ almost always referenced temporally recent information, while direct manipulation of the evidence map was used to reintroduce old

1 See Suthers (2006) for a discussion of "knowledge construction" and related terms, and the learning epistemologies that they imply. Briefly, I use "knowledge construction" for a Piagetian epistemology in which learners actively construct their knowledge in making sense of their world, and "collaborative knowledge construction" to acknowledge that at least part of this activity takes place in the social realm, perhaps even prior to the psychic realm as Vygotsky (1978) tells us. Unlike Stahl (2006), I do not use the similar phrase "knowledge building" because I understand that Carl Bereiter and Marlene Scardamalia intend it to refer to the intentional effort of a community to expand its collective knowledge base, which is a larger enterprise than I can claim I am studying here. Although "argumentation" can be a part of collaborative knowledge construction, I sometimes mention it separately because it is an active area of study (Andriessen, Baker & Suthers, 2003) and is addressed explicitly in one of the examples. At the end of the paper I will introduce "intersubjective meaning making" as the concept that I now believe is most useful, but throughout this paper I stay with the conception of "collaborative knowledge construction" under which the work was done.

² Deixis, derived from the Greek 'deiktikos' or 'demonstrative indicating or pointing,' is used in linguistics to refer to the function of a variety of grammatical and lexical features that relate an utterance to the spatiotemporal context of the act of utterance (Lyons, 1977, p. 636-637). In studies of multimedia communication, the term is generalized to include similar functions of nonlinguistic acts of communication, such as pointing.

³ Grammatical and lexical devices for deixis: see prior footnote.

information. In general, actions in the evidence map appeared to be an important part of participants' conversations with each other, and in fact was at times the sole means of interaction.

These observations led to questions concerning the nature and quality of interactions through the graphical evidence map. To what extent are participants engaging in argumentation and knowledge construction through the graphical representations as well as in the chat tool, and what are the methods by which they accomplish such collaborations? How are these methods distributed across the media we made available, and which affordances of that media do they rely on? Do the answers to these questions tell us how to design new media to provide better affordances? It was clear that the qualitative methodologies of our prior studies could not adequately answer such questions.

To begin to answer these questions, we identified four sessions (of 10) from our corpus that appeared to exhibit considerable interaction through the evidence map and made varying uses of the chat; and undertook qualitative analysis of these examples. Initially, progress was difficult. The video records of each user's screen are slow and tedious to watch: relevant events are spread out over time and are difficult to keep track of, and tools for video analysis are generally limited to annotating points of the video rather than supporting structural analysis. Computer logs provide a complete and compact representation of the interaction, but do not make relevant patterns of interaction salient. To address these analytic difficulties, we invented a graphical representation of the transcripts that make interaction (or its absence) more salient. This analysis bears similarities to the analysis of reasoning in conversation in Resnick, et al. (1993), but is designed to be applied to communication via a computer medium rather than spoken conversation and takes into account manipulations of visual representations as well as linguistic communication. The analysis was layered in a bottom-up manner similar to Mühlenbrock & Hoppe (1999), but conducted by a human analyst rather than automated. The layers of analysis included (1) a media-level of analysis that identifies individuals' literal actions in the media, (2) a referential level of analysis that identifies how these acts refer to, manipulate, or otherwise take up the products of previous acts, and (3) an intentional level of analysis that identifies phenomena of interest such as the accomplishment of argumentation and collaborative knowledge construction. This layering is intended to work from (1) a level of description that is relatively objective through (2) an intermediate level of description that requires some analytic inference but attempts to remain as theory neutral as possible to (3) theoretical analyses of the collaboration. As part of the work undertaken in Suthers et al. (2003), we had already coded (1) the literal actions taken by participants in the shared workspace, for example, adding, editing and linking objects in the evidence map. For the present work, my student Ravikiran Vatrappu and I (2) identified ways in which information "flows" between participants through the evidence map as well as chat, as evidenced by their references to information, whether verbally, by restatement or by direct manipulation. We then layered on top of this "uptake graph" (3) our own interpretations of the intentions behind these references, and looked for patterns of interaction that could be taken as evidence of knowledge construction activity.

The purpose of the present paper is primarily to describe the analysis method, particularly the uptake graph, so that it may be applied by others, and to secondarily to illustrate its value by showing how it helped answer some of the questions that motivated development of the method in the first place. Most of the paper will develop the analysis method, preparing us to examine case examples of how collaboration can take place through graphical media such as the evidence map.

Limitations of the paper should be mentioned to avoid misunderstanding. The paper does not directly address design: it does not analyze the design of the software used nor compare it to other designs. Investigation of design questions using the methodology described in this paper is on our future agenda. See also Dwyer & Suthers (2005) for a study deriving design criteria from users, and Suthers, et al. (2006) for an experimental study of some design choices. The present work is primarily concerned with uncovering how participants can make use of the available semiotic resources, including graphical resources, to collaborate, rather than with the effect of design variables. Therefore the methodology is not experimental or even comparative. Instead, it is descriptive, seeking adequate ways to capture online collaboration. However, this descriptive method could be applied within experimental designs in the future (see Suthers, 2006 for a discussion of this proposal). Finally, the paper is not proposing a theory. At one level, the objective is just the opposite: the second level of analysis is an attempt to offer a theory-neutral characterization of the interaction to which multiple theories may be applied, whether in a comparative or eclectic approach. It is intended as a “boundary object” (Star, 1990) supporting discourse between theoretical disciplines. The application of several theories is discussed in this paper in order to illustrate how the uptake graph may be used in the third level of analysis.

The remainder of this paper is organized as follows. First, in order to provide some context and an example used throughout the paper, the situation in which our data was gathered is summarized and a portion of the data at the first level of analysis (identification of observable events in the transcript) is presented. The obscurity of this transcript motivates the development of a notation to visualize interaction in the next two sections. The next section develops the second level of analysis, proposing the concept of “uptake” as the basic unit of collaboration, presenting a notation for uptake graphs, and applying this notation to our sample data. This brings us to the third level of analysis, at which we consult with several relevant theories to identify what might count as evidence of knowledge construction in the uptake graph. Then, having described all levels of analysis, this method is applied to two case examples to illustrate how it led to findings of argumentation and knowledge construction through manipulations of the evidence map. The paper concludes with a discussion of further research.

2. Study from which the data was derived

The data that is the object of the current analysis was taken from a previous study, reported in Suthers et al. (2003). The specifics of the previous study are not relevant to our present concern: it merely provides a convenient sample of interaction through computer media in order to gain insight into how people might accomplish collaborative knowledge construction through such media. If there is “order at all points” (Sacks, 1984), then this data will have something relevant to tell us, our concern being with an existence proof rather than universally quantified claims. However, we summarize how the sessions were conducted so that the case examples to be analyzed and nature of the data will make sense to the reader.

2.1. Participants and Procedure

Participants consisted of self-selected, same-gender pairs from introductory natural science courses at the University of Hawai'i. The participants' task was to propose and evaluate hypotheses concerning the cause of ALS-PD, a neurological disease with an unusually high occurrence on Guam that has been studied by the medical community for over 50 years. (A recent hypothesis proposes that human consumption of fruit-eating bats provides a vector for concentration and transmission of a neurotoxin to humans: see Lieberman, 2004).

Participants worked in a software environment shown in Figure 1. An information viewer enabled participants to advance individually through a series of textual pages presenting information on ALS-PD. One could not go back to previous pages, and the sequence of pages was designed such that later pages sometimes affected the interpretation of information seen several pages earlier, making the use of a representations important for recording and later reinterpreting information.

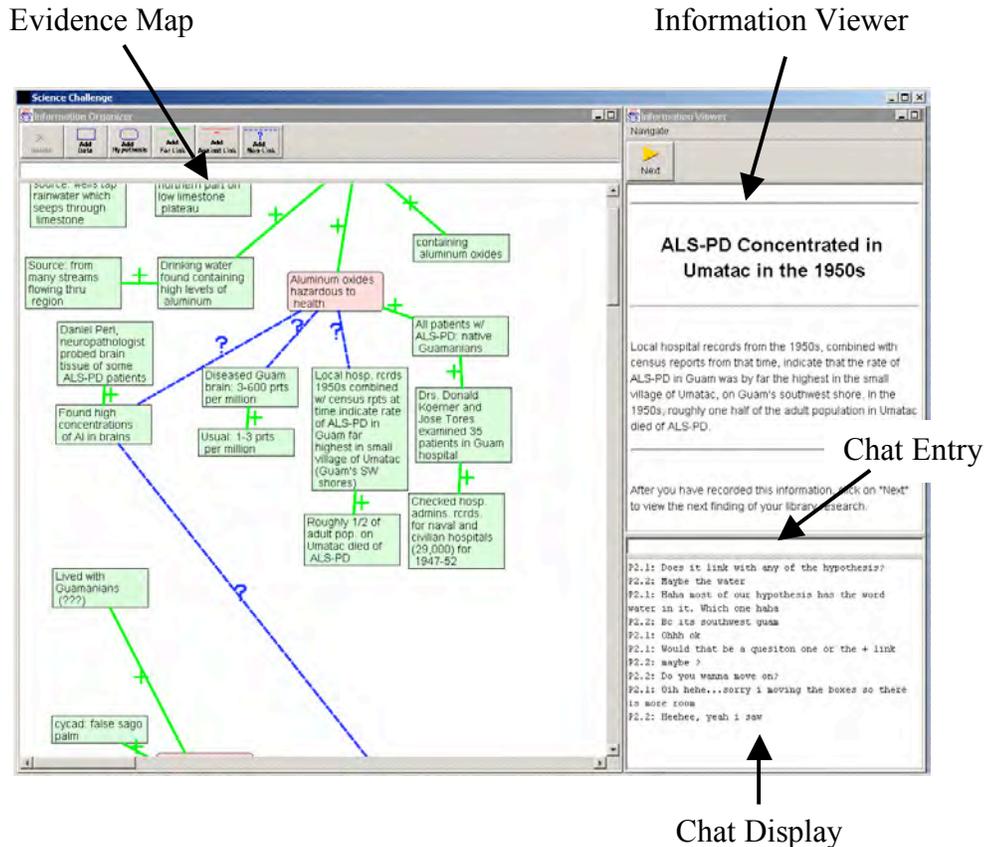


Figure 1. Collaborative workspace

This external memory took the form of a graphical evidence map (Figure 1) for constructing representations of the data, hypotheses, and evidential relations that participants gleaned from the information pages. The evidence map was based on version 3 of Belvedere (Suthers et al., 2001). Belvedere was used for historical reasons: its availability to the investigator. The design of Belvedere is not the topic of the present paper. It provided a workable representational tool for conducting the present study. The software enabled participants to build a graph of nodes (data items and hypotheses) and links (evidential relations) representing an evidence model. Links can be created to represent consistency (+), inconsistency (-) or unspecified (?) relations. The evidence map workspace was shared: manipulations enacted by one participant were immediately propagated to the other participant's display.

The software supported synchronous online communication with a chat tool. Messages typed into a chat entry box were sent to both participants' shared chat displays once the Return

key was pressed. Also, the software supported gestural deixis⁴ in two ways; one being automatic and the other requiring more deliberate action on the part of the user. If the user passed the cursor over an object in the evidence map, the fill-color of the object changed to blue. This was intended to enhance the deictic value of the cursor by making its location more visible. If the user selected an object with the cursor, the object was highlighted in yellow. The online version of the software replicated both of these color changes to the remote display. To maximize the potential for online participants to use this option for gestural deixis, we demonstrated this highlighting to them.

The pairs were given a 10-minute introduction to the problem, the task and the software. One participant was then led to a separate computer in a different room. They then engaged in a 12-minute warm-up exercise on an unrelated problem (the causes of mass extinctions). The main problem consisted of 15 informational pages on the ALS-PD disease, and participants were allowed to continue their interaction until they felt they had reached a conclusion. At the conclusion of the session, participants were given additional tasks not relevant to the present paper (see Suthers & Hundhausen, 2003; Suthers et al., 2003).

Table 1. Portion of log, session 3 (edited for readability)

<i>Time</i>	<i>Who</i>	<i>Action</i>	<i>Object(s)</i>	<i>Content</i>
14:33:10	P1	Edited	H02	(fadang) cycad seeds in medicine cause guam diseases
14:33:18	P2	Added	D11	cycads grow in thier soil/ soil tests are high for aluminium content
14:33:24	P1	Display	Page 12	Aluminum Abundant Throughout the World
14:33:28	P2	Display	Page 12	Aluminum Abundant Throughout the World
14:33:52	P2	Linked	D08+D11	"for link" (+)
14:34:04	P1	Chat		so its not aluminum then?
14:34:11	P2	Chat		maybe it is?
14:34:15	P2	Chat		aluminium poison?
14:34:17	P2	Chat		ing
14:34:24	P2	Chat		i dunno?////////////////////
14:34:25	P1	Chat		maybe but never heard of that
14:34:31	P2	Chat		ya
14:34:33	P2	Display	Page 13	BMAA-fed Monkeys Exhibit Signs of ALS-PD
14:34:34	P1	Display	Page 13	BMAA-fed Monkeys Exhibit Signs of ALS-PD
14:34:43	P1	Added	D12	aluminum poinsoning????
14:34:51	P1	Linked	D12?D08	"unspecified link" (?)
14:35:40	P2	Added	D13	animals tested for BMAA an amino acid didn't have the same ...symptomts as some one w/ als
14:36:10	P1	Linked	D13?H02	"unspecified link" (?)
14:36:20	P2	Display	Page 14	Natural Toxin in Food Causes Paralytic Disease
14:36:21	P1	Selected	D13?H02	"unspecified link" (?)
14:36:23	P1	Deleted	D13?H02	"unspecified link" (?)
14:36:28	P2	Linked	H02-D13	"against link" (-)
14:36:36	P1	Added	D14	BMAa in cycad seeds
14:36:51	P2	Linked	D14+D13	"for link" (+)
14:36:51	P1	Linked	D14+H02	"for link" (+)
14:36:55	P1	Selected	H02-D13	"against link" (-)
14:36:56	P1	Deleted	H02-D13	"against link" (-)

⁴ Gestural deixis: referencing accomplished by pointing. See previous footnote for definition of "deixis."

Transcripts of chat messages and user actions in the graph were automatically logged in the online sessions. An example of a log fragment, edited for readability, is shown in Table 1. P1 and P2 refer to the two participants. H02 is the 2nd hypothesis object created in the evidence map, while D11 is the 11th data object created, etc. The content column shows the contents of these evidence map objects or the title of the article displayed. The notation “D08+D11” refers to a “+” (consistency) link between D08 and D11, while “H02-D13” names a “-“ (inconsistency) link, etc. The items linked are listed in the order in which the participant selected them (e.g., D08+D11 means the participant selected the “+” linking tool, selected D08, and then selected D11).

The reader may find it difficult to see collaborative knowledge construction in this log. This is the same problem that we faced as analysts. Even though we were more familiar with the log data than the reader and also had the video at our disposal, we also found it to be difficult to uncover the occurrence and nature of collaboration. This problem was part of the motivation for the invention of the visualizations presented in this paper. Therefore the reader is requested to persevere, as the rest of the paper will give Table 1 more meaning.

3. Method of Analysis

The present study seeks to uncover and characterize how collaborative knowledge construction is accomplished via manipulations of representations. Given a transcript such as that in Table 1, how would one go about identifying episodes of collaborative knowledge construction?

As mentioned previously, we found the quantitative methods used in our previous analyses to be unsuitable for addressing the question of how collaborative knowledge construction is accomplished. We call these methods “coding and counting,” as the fundamental approach is to segment a record of the interaction into units such as actions and utterances, annotate these units under some coding system that describes attributes of the units in isolation, count up the number of units within a given session, and then compare these counts across treatment conditions. Such quantitative methods are suitable for testing predictions of certain types of differences between treatment groups (e.g., that users of one version of the software will talk more about evidence), but cannot capture the accomplishment of collaborative acts as they unfold in time. Therefore we turn to family of analytic methods known as sequential analysis (Sanderson & Fisher, 1994). A fundamental assumption is that the meaning of an act or utterance is a function of its context of the prior sequence of acts and utterances: this is why coding isolated segments as in quantitative analysis is inadequate. For example, the analysis later in this paper will expose the addition of a “for link” between D14 and H02 at 14:36:51 in the transcript of Table 1 as not merely an assertion of this relationship, but as part of an argument in which it is posed as an alternative to another relationship.

In order to analyze collaborative knowledge construction, we need more than to simply consider an act within a context. Any collaboration requires some form of interaction between participants. Therefore, we need some means of identifying events in the media transcripts that constitute interaction, as distinguished from parallel yet unrelated activity that happens to be in the same medium.

Methods for analysis of spoken conversation often focus on *turns* or *adjacency pairs*: utterances by different speakers that are temporally adjacent to each other (Goodwin & Heritage, 1990, Sacks, Schegloff & Jefferson, 1974), because each utterance is normally relevant to that which immediately preceded it. When this assumption is violated, the utterance is considered to be a change of topic (Grosz & Sidner, 1986). However, analysis based on turn-taking or

adjacency pairs is not suitable for most online data because of two ways in which the synchronicity of computer-mediated communication (CMC) may differ from face-to-face conversation: parallel production and persistence.

First, adjacency can be disrupted in CMC, potentially creating incoherencies in the dialogue. CMC (other than voice media) differs from spoken conversation in not suffering from *production blocking*: a property of media where only one person can produce a contribution (e.g., speak) at a time if all contributions are to be accessible to others. Most online media are “synchronous” in a different way: multiple participants can produce their contributions simultaneously. Since these contributions are sent as complete units, they may become available to other participants in unpredictable orders. For example, if C_{A1} is a contribution and C_{A2} is a reply to it, while C_{B1} is an unrelated contribution and C_{B2} is a reply to it, it is possible to receive contributions in order $C_{A1} C_{B1} C_{A2} C_{B2}$ (or other permutations that place C_{A1} before C_{A2} and C_{B1} before C_{B2}), even though C_{B1} is not a reply to C_{A1} . People do find ways to manage these disruptions by inferring and even explicitly marking the intended reply structure (Herring, 1999). However, the implication for analysis remains: we cannot simply focus our analysis on the relationships between adjacent events. Like the participants, analysts need to identify relevance relations between contributions. (See Stahl, this volume, for a related discussion of logical adjacency.)

A second reason that analysis cannot assume that relevance follows temporal adjacency is persistence. In CMC, “communication becomes substance” as Dillenbourg (2005) so aptly puts it: the ephemerality of audio is replaced with the persistence of linguistic and graphical inscriptions in a computational medium that enables replication of these inscriptions over space and review of them over time. If one can review what was said five minutes, hours or days ago as easily as what was said five seconds ago, it is far easier to take up contributions separated widely in time. Even synchronous CMC has an asynchronous aspect. This is the case in our data. Although many contributions exhibit temporally adjacent coherence, participants are constructing shared representations over a period of an hour or so, and may at any time address an inscription that was created much earlier in this session. For example, at 14:36:10 of the transcript in Table 1, a hypothesis is taken up that was last manipulated at 14:33:10, but first created at 14:27:47 (not shown).

In summary, we cannot reduce the complexity of analysis by reducing the time-window in which one searches for relevance relations to adjacent contributions. Any contribution that was reified in the persistent medium could be taken up again. We need an alternative basic unit of interaction between participants that accommodates noncontiguous contributions.

3.1. The Uptake Graph

Based on the observation that collaboration is only possible when information is shared and transformed between participants, we began to work with the concept of “information uptake:” the event of a participant doing something with previously expressed information. This “doing” could take the form of a chat message or a manipulation of the graph. It could add to or modify the expressed information, or relate it to new information. Uptake could take up a participant's own prior contribution as well as those of others: by identifying both, we realized we could characterize the mixture of intrasubjective and intersubjective knowledge construction. Uptake is similar to the “thematic connections” of Resnick, et al. (1993), but allows for media as well as linguistic relationships. Like Resnick and colleagues, we try to keep our uptake graph close to the surface structure of the transcript, leaving theoretical inferences to the next level of analysis.

The second level of analysis thus begins with identifying the cases where a participant takes up contributions of previous participants. These uptake events can be expressed as a hypergraph $G=(V,E)$. The set of vertices V is the set of contributions visible in the media. E is a set of tuples $(\{c_1, \dots, c_n\}, c_u)$, $c_i \in V$, representing relations where contributions c_1 through c_n are “taken up” by contribution c_u . The graph is a directed acyclic graph: if the subscripts are time stamps then $u > i$, for $i = 1, \dots, n$. Informally, we want to find all the prior contributions that each contribution draws directly upon. V is partitioned into C_1 and C_2 (or more generally, $\{C_1 \dots C_m\}$), according to which participant made the contribution, for example, P1 or P2 in Table 1. If some of the contributions $\{c_1, \dots, c_n\}$ are made by a different participant than who made c_u (i.e., one of $c_1 \dots c_n$ is in a different partition than c_u), then there is intersubjective uptake, and the potential for collaboration exists.

To best serve as a starting point for analysis, this *uptake graph* would ideally be based on observable evidence of uptake and would avoid theoretical assumptions or inferences. Our motivation for grounding the uptake graph in observables is simply to conduct good empirical science. Our motivation for limiting (and preferably eliminating) the inferences required to construct the graph is twofold: (1) to reduce the complexity of an analysis by dividing up the task, separating the activity of constructing this grounded structure from the activity of conducting theoretical interpretations; and (2) to allow the graph to serve as a basis for comparison and integration of multiple theoretical interpretations, i.e., a boundary object for the study of collaboration.

To approach this ideal, each uptake relation was derived from notational relationships between visible media events as well as temporal contiguity where it seemed merited. For example, if c_u edited a prior contribution, or created a link between two prior contributions, or reused the wording of a prior contribution, or referenced the prior contribution through verbal or gestural deixis, then these prior contributions were included in $\{c_1, \dots, c_n\}$. These examples may be found in Table 1:

- *c_u edited a prior contribution:* At 14:33:10, P1 edits hypothesis H02, last edited by P1 at 14:28:22 (not shown).
- *c_u created a link between two prior contributions:* At 14:36:51, P2 takes up D14 (created by P1 at 14:36:36, P1) and D13 (created by P2 at 14:35:40)
- *c_u reused the wording of a prior contribution:* At 14:34:43, P1’s wording in D12 follows and thus takes up the wording of P1’s chat at 14:34:15.
- *c_u referenced the prior contribution through verbal or gestural deixis:* P1’s chat at 14:34:04 (“it’s not aluminum”) references an earlier idea (not in the table) that aluminum in the water causes ALS-PD, in relation to information in the page just displayed at 14:33:28 stating that aluminum is abundant throughout the world.

In general, the uptake relations must meet the criteria that the uptake was evidenced by observable manipulations or by reference to or reuse of informational content.

3.2. Visual Notation

Although the graph-theoretic notation is useful for providing a definition of the uptake graph, it would be cumbersome for conducting the analysis. A symbolic representation of uptake relations would not make the patterns of interaction salient. To serve our analysis, we invented a mixed tabular/diagrammatic notation (exemplified in figure 3 and subsequent figures). The tabular portion is based on the transcript but represents the activity (chat and changes to the representation) of participant 1 (P1) in the left hand column, and activity of participant 2 (P2) in the right hand column. A column in the middle is reserved for a diagrammatic (visual)

representation of the uptake graph. If $(\{c_1, \dots, c_n\}, c_u)$ is a tuple in the uptake graph (i.e., c_u builds on the information in $\{c_1, \dots, c_n\}$), then an n -tailed arrow is drawn from contributions $\{c_1, \dots, c_n\}$ to contribution c_u . The arrow is directed forward in time, as it shows the “flow” of information between actors (which may be the same or different participants) via the representation.



Figure 2. Key for uptake relations

The uptake instances were categorized in a manner reflected in the key of figure 2. As we are interested in collaboration through the representations rather than just individual use of the representations, we encoded this distinction: Dashed lines represent *intrasubjective uptake*: manipulation of items that were most recently manipulated by the *same* participant, while solid lines represent *intersubjective uptake*: manipulations of items that were most recently manipulated by the *other* participant. Therefore, solid lines represent potential collaborative knowledge construction in the sense of informational uptake from one actor to another. This coding is based on the actor who *most recently* manipulated an item in the persistent representation, not necessarily the originator of the item, because we wanted to capture the “back and forth” of co-manipulation of a representation. If references only went back to the original creation of an information item then it would not be possible to trace out dialectic interaction. For our particular analysis, color⁵ coding is used for a noninferential categorization of the uptake in terms of the type of action taken in the medium. These codes depend on the nature of the medium, so would differ from study to study.

A visualization of the uptake graph for a portion of the transcript of Table 1 is shown in Figure 3. Some of the transcript has been omitted to reduce the size of the figure, and focus on that portion needed to illustrate an interaction discussed later in the paper. Examining the uptake annotations, we see that both participants are bearing the burden of collaboration, as arrowheads go to both participants. Participant 1 exhibits more than twice the number of uptake acts (arrowheads) in this segment, but the combination of solid and dotted lines show that participant 2 is integrating information from both participants. We also see that all polarities of evidential relations are being considered (+, -, or ?, represented as green, red and grey). There is no chat in this segment: in fact, participants only chatted one more time, several pages later, on an unrelated hypothesis.

⁵ Color versions of the figures are available at <http://lilt.ics.hawaii.edu/lilt/papers/2006/Suthers-2006-RPTL-Figures/> or on request from the author.

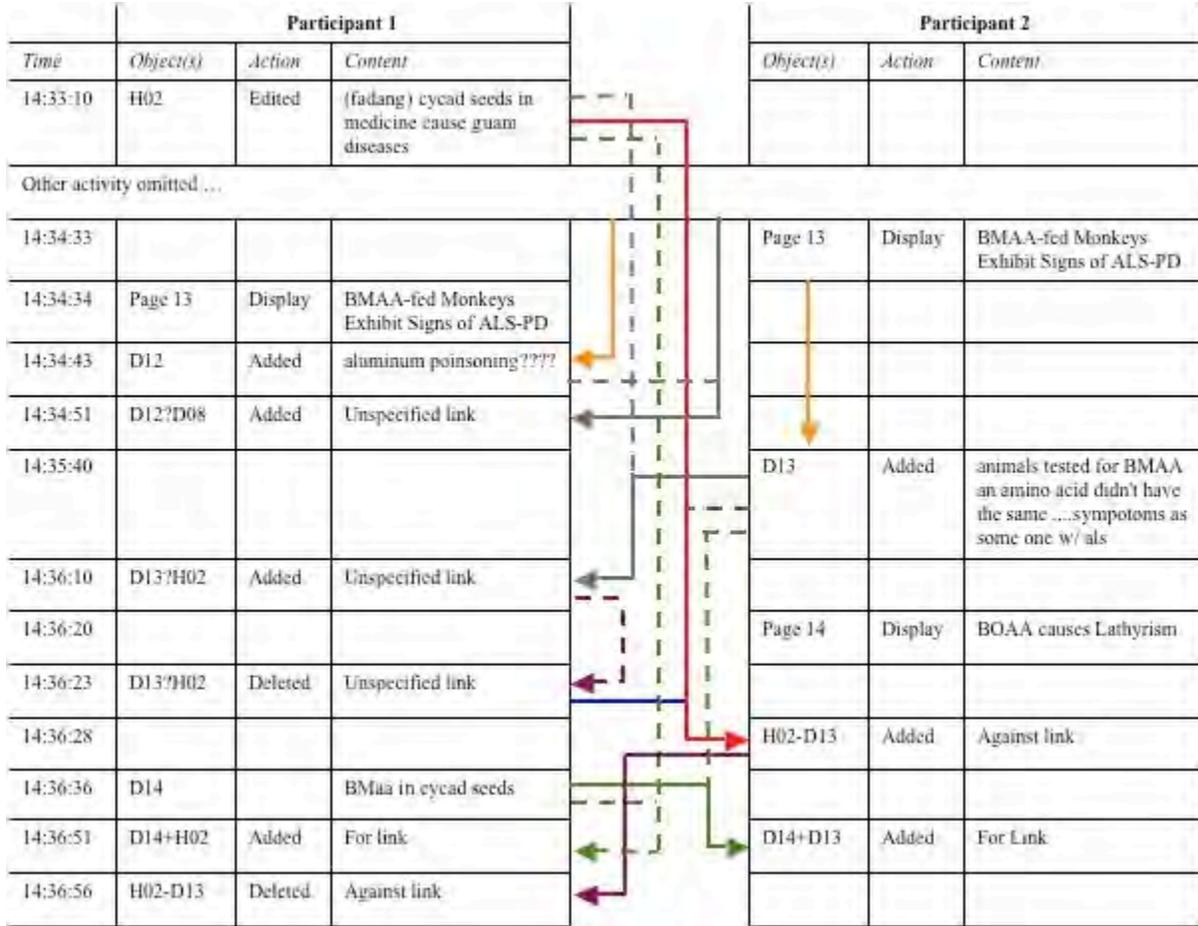


Figure 3. Uptake graph visualization for selected portions of the transcript for Session 3.

The uptake diagrams were constructed in PowerPoint™ because it allowed for construction of graphics overlaying a table in an elongated page size. This was a tedious process: development of specialized software support for partially automated construction and selective viewing of these graphs is on our agenda. Overviews of portions of the original PowerPoint visualizations for two transcripts are shown in Figure 4. (In these graphs, some of the intrasubjective uptakes are drawn on the outside edges to reduce crowding.) The bottom portion of the left hand graph corresponds to the interaction detailed in Figure 3. The bottom portion of the right hand graph will be detailed in Figure 7. Although one cannot read details at this granularity, the figure illustrates how trends and patterns of interaction may be identified:

Vertical lines indicate that participants are revisiting prior information. Revisited items are almost always information that is expressed in the persistent external representation (the evidence map). Therefore, an abundance of long vertical lines indicate that participants are taking advantage of the external representation's persistence, returning to previously encountered or expressed information. The many vertical lines in the right hand side of Figure 4 shows that participants in session 8 often reintroduced or reconsidered older ideas.

Arrow heads point to uptake acts, so the distribution of arrow heads across the two columns, along with solid and dotted lines, helps identify whether there is any asymmetry in participants' roles. This was exemplified in the discussion of Figure 3.

Color indicates the relationship between the items taken up. For example, red and green indicate that an inconsistency or consistency relation (respectively) is being noted. Therefore presence of these colors suggests the nature of the argumentation. Blue is used for revisions, and maroon for deletions. Red and maroon suggest that there may be conflict; green and blue suggest the accretion and refinement of ideas. Readers who have access to a color version of this figure (available from the author) will see a cluster of red and maroon at the bottom of the left hand side of Figure 4 suggesting conflict in session 3: this is why we chose to investigate the portion shown in Figure 3 in more detail. Session 8 shows primarily consistency relations, and no inconsistency relations, with revisions and deletions spread evenly throughout, suggesting cooperative construction of ideas but the possibility of a confirmation bias.

3.3. Analyzing Knowledge Construction

Once potential collaborations have been identified in the form of an uptake graph, the third step is to analyze this uptake graph in order to identify phenomena of interest such as collaborative knowledge construction. This level of analysis is necessarily specific to the purpose of the inquiry and the analyst's theoretical predilections. Properly done, the uptake graph should support multiple theoretical approaches. Stahl's (2006, chapters 9 and 15) model of collaborative knowledge building⁶ includes more than several cognitive and social processes, suggesting that an eclectic approach to analysis will indeed be needed. We found this model to be insightful yet somewhat overwhelming as the starting point for analysis, and therefore chose to use uptake as the basis for identifying the possible presence of knowledge construction. This section discusses some theoretical perspectives that can inform the interpretation of uptake acts.

In analyzing the examples in this paper, we first considered an influential theory of linguistic communication: Clark's model of grounding, a component of Contribution Theory (Clark & Brennan, 1991; Monk, 2003). We can restate grounding in terms of actions on a workspace representation such as an evidence map as follows: a participant expresses an idea in the representation; another participant acts on that representation in a manner that provides evidence of understanding the first participant's intent in a certain way; the first participant can choose to accept this action as evidence of sufficient understanding, or, if the evidence is insufficient, initiate repair. Under the grounding perspective, the analyst would look for *sequences of actions in which (1) one participant's action in a medium is (2) taken up by another participant in a manner that indicates understanding of its meaning, and (3) the first participant signals acceptance.*

Given the foregoing discussion of the nonadjacency of uptake in persistent media, we can immediately see that it is not sufficient to examine adjacent interlocutors' acts for (2) evidence of understanding or (3) the acceptance signal. Potentially any subsequent reply could constitute (2), so can be difficult to identify. Worse, the final signal of acceptance (3) is often implicit. It can consist merely of continuing the interaction rather than initiating repair of a breakdown. For example, P2's addition of the link D14+D13 (14:36:51) can be read as evidence for understanding P1's expression of D14. P1 does not acknowledge this link but does not challenge

⁶ See the first footnote on terminology.

it either. But implicitness is a property of interaction, not a limitation of the analysis method. More damaging, an analysis based solely on contribution theory at best can tell us only how people check that they act as if they have achieved mutual understanding, but does not inform us about the process by which this mutual understanding is reached (if it is). Furthermore, contribution theory is hampered by over-reliance on the notion of agreement. It is possible and even necessary to engage mutually through disagreement as well as agreement (Matusov, 1996). For example, does P1's deletion of P2's link in 14:36:56 indicate a lack of grounding, or disagreement? Therefore the theory was found to be of limited value in identifying the interactions that constitute collaborative knowledge construction. However, this perspective does suggest one way in which we might view interaction through representations as a form of nonverbal or semi-verbal conversation.

Socio-cognitive conflict theory (Doise & Mugny, 1984) and cognitive dissonance theory (Festinger, 1957) describe various ways in which learning can result from social interaction in which individuals encounter ideas that are different from their own. The individual is challenged to reconsider his or her beliefs, potentially leading to change, or to explain and justify those beliefs to others, leading to clarifications and elaborations that might not have otherwise taken place. Representations that externalize one's beliefs can make beliefs explicit enough for one's interlocutors to notice conflicts, thereby initiating a socio-cognitive process of learning through argumentation (Andriessen, Baker & Suthers, 2003). This noticing is especially likely to occur if multiple participants have externalized their beliefs in a representational system that makes conflicts explicit, and processes of elaboration and reconceptualization may also be externalized in interactions between participants that take place via shared representations. Under the socio-cognitive conflict perspective, we would want to *identify situations in which the externalization of ideas led to identification of differences (as well as commonalities) of interpretation that were subsequently taken up by at least one of the individuals involved*. In addition to overt verbal argumentation, clues that conflict is being addressed might include revision or deletion of the others' ideas (e.g., 14:36:51-14:36:56 in Figure 3) or creating an explicit conflict relation between inscriptions of one's own and others' ideas, if the representation provides for such relations. We will return to the argumentation found in Figure 3 shortly.

The foregoing perspective is useful as far as it goes, but limiting in that it treats participants as separate cognitive entities that interact via language and other notations, yet retains the locale of knowledge construction activity within the individual. A distributed cognition perspective (Hollan, Hutchins & Kirsch, 2002; Hutchins, 1995) suggests that cognitive activities such as knowledge construction are distributed across individuals and information artifacts through and with which they interact. In this perspective, the information-transformative and interpretive components of a cognitive activity can occur across multiple individuals via external representations. An individual can express a conception through a change to a representation that is shared with another individual, who subsequently takes up this information and adds to, transforms or interprets it in a new way, again resulting in a change to the representation that may be taken up by the first individual, where this distributed transformation can be viewed in the whole as an act of cognition. Knowledge construction, being a form of cognition, can also take place with and through external representations of various (visual and symbolic) forms not limited to language. Therefore, under the distributed cognition perspective we would *look for transformations of representations across individuals where those transformations can be interpreted as an intersubjective cognitive process*. Examples include sequences of actions in which participants merge, revise, and connect representations of ideas

contributed by other participants. Shortly an example of such an interaction will be provided based on session 8, the session visualized in the right hand side of Figure 4 that displays a prevalence of integrative activities.

The activity theoretic perspective (Cole & Engeström, 1993) considers how activity is embedded in a system that includes not only the self and the object or topic of interest, but also tools, one's community, one's role in this community, and the norms for behavior in the community. The concept of mediation is central. When we examine the relationship between any two elements of an activity system (the subject, object, tool, community, roles, rules), we can sometimes benefit from asking how a third element mediates the relationship between the first two, influencing the form the relationship takes. For example, external representations can mediate between individual and community by reifying aspects of prior practice in ways that facilitate replication (and transformation) of that practice. Similarly, external representations mediate collaborative inquiry when collaborators try to make sense of them (Roschelle, 1994). They can also influence one's interactions with others by suggesting specific epistemic activities (Collins & Ferguson, 1993) or facilitating or inhibiting cognitive activity (Blackwell & Green, 2003). Under an activity-theoretic perspective, we would analyze collaborative use of representations by looking for *ways in which the representation mediates (makes possible and guides) interactions between participants by virtue of its form*, especially where this guidance replicates prior community or culturally sanctioned patterns of action. For example, in our transcript we find that participants are concerned with the relationship between H02 and the collection of data represented by D13 and D14. This concern with relationships is encouraged and mediated by the "epistemic form" (Collins & Ferguson, 1993) of the evidence map, in which the fundamental activity is one of linking ideas by evidential relationships. The representation mediates the relationship between thought and action.

This viewpoint is consistent with the distributed cognition perspective, as well our own work on representational guidance (Suthers & Hundhausen, 2003). This work identified roles of external representations that are unique to situations in which a face-to-face pair is constructing and manipulating shared representations as part of a constructive activity, roles that suggest events to look for in an analysis. First, an individual who wishes to add to or modify a shared representation may feel some obligation to obtain agreement from one's group members, leading to negotiations about and justifications of representational acts. This discourse will include negotiations that would not be necessary in the individual case, where one can simply change the representation as one wishes. The creative acts afforded by a given representational notation may affect which negotiations take place. An implication for analysis is that we should look for *discussions initiated as participants prepare to act upon a representation*. Second, the constituents of a collaboratively constructed representation, having arisen from negotiations of the type just discussed, evoke in the minds of the participants meanings beyond what external observers might be able to discern by inspection of the representations alone. These constituents provide a resource for reference to ideas previously developed. In this manner, collaboratively constructed external representations facilitate subsequent negotiations; increasing the conceptual complexity that can be handled within group interactions and facilitating elaboration on previously represented information. An implication for analysis is that we should identify *ways in which participants use representations as a means of bringing ideas to the attention of others*.

4. Case Examples

The foregoing discussion included selected examples from our data. In this section, we provide more complete case examples. The examples provide existence proofs of the possibility that collaborative knowledge construction can be accomplished via shared representations while also illustrating the value of the methodology. First we complete the ongoing example of session 3, applying the level 3 analysis to uncover disagreement that takes place through manipulation of the graph. Then we examine a portion of session 8 in which participants interactively co-construct an interpretation of new data purely through the graphical medium, using the chat only for a final explicit acknowledgement of their accomplishment. Finally, we briefly discuss situations in which the chat rather than the graph was used for evaluation and interpretation of the data.

4.1. Session 3: Arguing Through the Evidence Map

The interaction of the pair in session 3 exemplifies how a conversation-like interaction can take place through manipulation of the evidence map, and how conflict can be identified and addressed (albeit not satisfactorily in this case) via manipulations of the map. The uptake graph is repeated in Figure 5, along with contextual information and a table providing our interpretation of the actions taken in the workspace. We skip the actions in 14:34:43 and 14:34:51, as they are wrapping up a previous interaction, but return to them at the end of the examples.

In this exchange, participants are exploring the implications of some new evidence for their second hypothesis (H02), that the cycad seeds cause the disease. Reading the uptake relations starting at 14:36:10, we find that this interaction has the form of a disagreement, summarized in the right hand column of the table by paraphrasing the actions on the evidence map as if they were a verbal conversation. At 14:36:10, P1 suggests the possibility of a relationship (D13?H02), but then retracts it 13 seconds later. At 14:36:28, P2 proposes that this is a negative relationship (H02-D13). As well as taking up H02 and D13 again, this act is essentially a modification of the deleted link (blue line) because it is relating the same elements. Then, after introducing some related data (D14 at 14:36:36), P1 proposes that this actually supports H02 (D14+H02, 14:36:51). At the same time (14:36:51), P2 is integrating the new data P1 introduced with her own interpretation (D14+D13). (Participants commonly use + to collect related data as well as for linking evidence to hypotheses.). Finally, at 14:36:56, P1 makes it clear that she is disagreeing with P2's proposed relationship (H02-D13), which has the opposite polarity of the one she just created (D14+H02), by deleting it. This case exemplifies a situation in which the externalization of ideas led to identifications of differences of interpretation that were subsequently taken up by the individuals for attempted resolution via manipulations of the evidence map alone, at least for short episodes. By this account, participants are engaging in a form of argumentation through the evidence map, without using the chat tool.

Upon closer examination the source of the disagreement can be seen to be an erroneous reading of the text. The text contains a double negative "Such symptoms are not unlike those of someone with ALS." P2 apparently read this as simple negation, writing that the animals "didn't have the same symptoms as some one w/als." This error accounts for P2's confidence that the data conflicts with H02. Apparently, the participants did not identify the source of their disagreement in this error of interpretation.

Context: Previously expressed by P1: H02 “(fading) cycad seeds in medicine cause guam diseases”

Participants have just read page 12 titled “BMAA-fed Monkeys Exhibit Signs of ALS-PD” reading “When scientists fed large doses of BMAA (an amino acid found in cycad seeds) to macaque monkeys, they observed the monkeys age before their eyes. After a few weeks' exposure to BMAA, some of the animals became weak. Over three months, some of the animals became apathetic, listless. Their hands trembled. They stooped and shuffled. Such symptoms are not unlike those of someone with ALS.”

Time	Participant 1				Participant 2		
	Object(s)	Action	Content		Object(s)	Action	Content
14:33:10	H02	Edited	(fading) cycad seeds in medicine cause guam diseases				
Other activity omitted ...							
14:34:33					Page 13	Display	BMAA-fed Monkeys Exhibit Signs of ALS-PD
14:34:34	Page 13	Display	BMAA-fed Monkeys Exhibit Signs of ALS-PD				
14:34:43	D12	Added	aluminum poisoning????				
14:34:51	D12?D08	Added	Unspecified link				
14:35:40					D13	Added	animals tested for BMAA an amino acid didn't have the same ...symptoms as some one w/ als
14:36:10	D13?H02	Added	Unspecified link				
14:36:20					Page 14	Display	BOAA causes Lathyrism
14:36:23	D13?H02	Deleted	Unspecified link				
14:36:28					H02-D13	Added	Against link
14:36:36	D14		Bmaa in cycad seeds				
14:36:51	D14+H02	Added	For link		D14+D13	Added	For Link
14:36:56	H02-D13	Deleted	Against link				

Time	Actor	Researcher's interpretation of the acts as a conversation
14:35:40	P2	<i>This looks like relevant data: “animals tested for BMAA an amino acid didn't have the same ... symptoms as some one w/ als”</i>
14:36:10	P1	<i>I think that has something to do with H02, but I'm not sure what.</i>
14:36:23		<i>Never mind.</i>
14:36:28	P2	<i>They conflict.</i>
14:36:36	P1	<i>It says that BMAA is in cycad seeds</i>
14:36:51	P2	<i>Right, that fits what I'm saying.</i>
14:36:51	P1	<i>So it's for the hypothesis.</i>
14:36:56		<i>You're wrong.</i>

Figure 5: A disagreement in session 3

4.2. Session 8: Graphical co-construction of an interpretation

The next example provides an example of co-constructive collaboration through the evidence map leading to a conclusion that is acknowledged verbally. The participants had previously represented a hypothesis (H02) that aluminum is the cause of the disease, and two data items, D05 and D06 concerning drinking water being the source of aluminum, these being linked by

consistency (+) links D05+D06 and D05+H02. After several pages concerning another possible disease agent, they encounter a new page indicating that ALS-PD patients have high levels of aluminum in their brains. The transcript begins at this point. The subgraph that resulted from the interaction is shown in figure 6. (We added the labels on the boxes. D06 is not visible because it was deleted.)

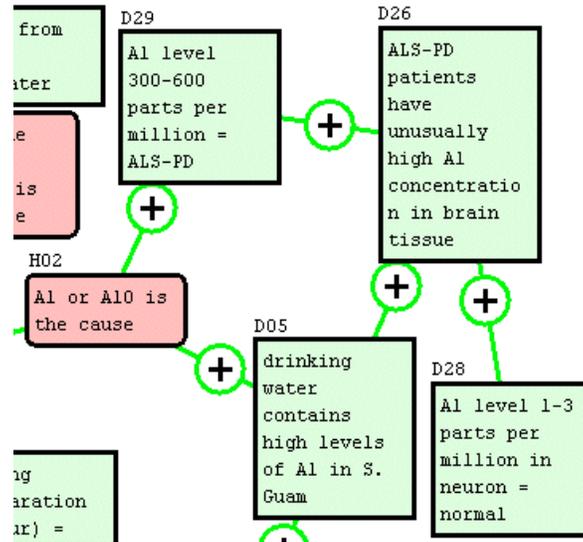


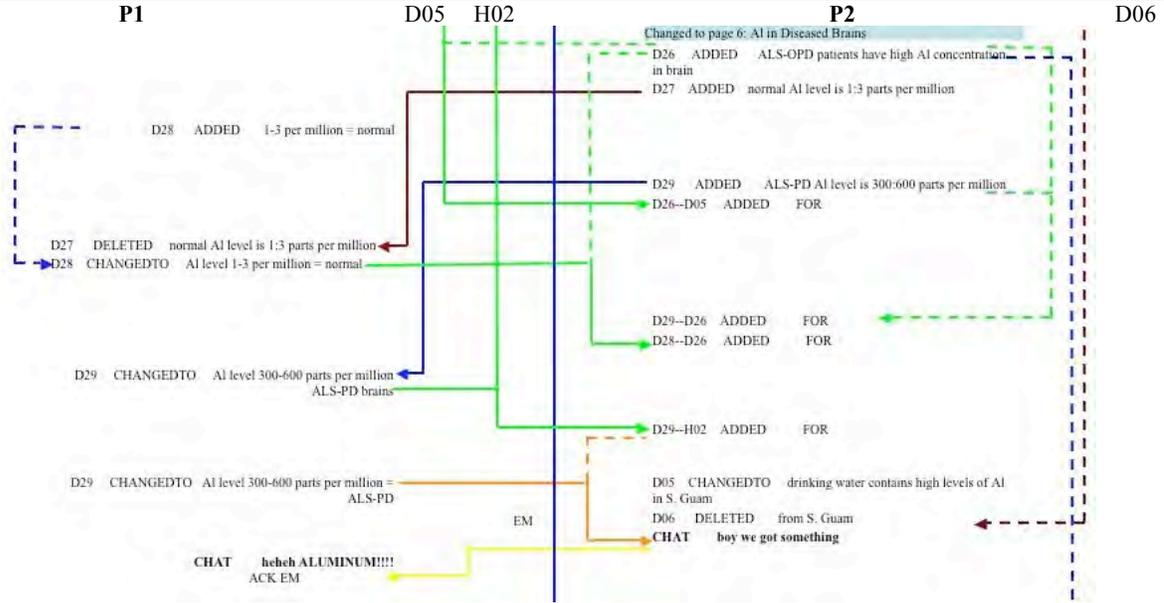
Figure 6. Fragment of graph constructed by participants in session 8

A simplified transcript and the original uptake graph we constructed are shown in Figure 7 (the original text notations are slightly different from those used in Figure 3). Taking an overview of the visualization of information uptake, we see that there is uptake through the evidence map in both directions: there are solid lines other than yellow with arrowheads going in both directions. Therefore, participants are collaborating through the evidence map; each is acting on information that was most recently provided or manipulated by the other. This visualized interactivity was why the present segment was selected for further analysis. We can see that the uptake involves integration through consistency links (green), deletion (maroon) and revision (blue); and that it draws upon material previously represented (lines going off the top of the image). This segment also exemplifies an asymmetric role division that was also seen in other pairs' sessions. P1 is adding and editing the content of the text boxes, while P2 is linking together information contributed by both P1 and P2 (of four links, one involves only P1's material, two bring P1 and P2's material together, and one involves only P2's own material).

Stepping through the annotated transcript, participants interacted as follows. At 15:49:51 onwards, P2 creates two data items D26 and D27 from the new information page. P1 is doing so at the same time, creating D28 at 15:50:35, which is redundant with D27. While P2 continues to work, P1 (apparently recognizing the redundancy) deletes P2's version (D27, at 15:51:09) and rewords his own version D28 (15:51:26) to include some information from D27 (that it is about aluminum). Parallel redundant activity followed by merging and cleanup is common in our online transcripts.

<i>Previous Objects</i>	H02	Al or AIO is the cause
	D05	drinking water contains high levels of Al
	D06	from S. Guam

Context: Participants have just read page titled “High Concentrations of Aluminum Found in Diseased Brains” which states: “Neuropathologist Daniel Perl X-ray probed the brain tissue of some ALS-PD patients. He found unusually high concentrations of aluminum in those brains. He says, “Normally, the background level of aluminum in a neuron is from one to three parts per million. In the diseased Guam brains we’re getting from three hundred to six hundred parts per million.””



<i>Time</i>	<i>Who</i>	<i>Act</i>	<i>Object(s)</i>	<i>Chat or graph content [spelling as given]</i>
15:49:51	P2	Added	D26	ALS-PD patients have high Al concentration in brain
15:50:20		Added	D27	normal Al level is 1:3 parts per million
15:50:35	P1	Added	D28	1-3 per million = normal
15:50:51	P2	Added	D29	ALS-PD Al level is 300:600 parts per million
15:51:08		Linked	D26+D05	
15:51:09	P1	Deleted	D27	
15:51:26		Modified	D28	Al level 1-3 per million = normal
15:52:21	P2	Linked	D29+D26	
15:52:24		Linked	D28+D26	
15:52:31	P1	Modified	D29	Al level 300-600 parts per million ALS-PD brains
15:52:47	P2	Linked	D29+H02	
15:52:52	P1	Modified	D29	Al level 300-600 parts per million = ALS-PD
15:53:25	P2	Modified	D05	drinking water contains high levels of Al in S. Guam
15:53:29		Deleted	D06	
		Moved	[various]	[repositions various objects for 44 seconds]
15:54:13		Chat		boy we got something
15:54:39	P1	Chat		heheh ALUMINUM!!!!

Figure 7. Co-construction of an interpretation in session 8

Meanwhile (15:50:51), P2 goes on to add one more data item D29 and link it to D05. The manipulation of D05 is a reintroduction of an item that has not been considered for a while (as is apparent from the line going off the top of the uptake graph): this exemplifies the utility of a visual representation for reminding participants of previous information and enabling them to reference it easily. D05 was originally created and was last manipulated by P1; therefore this

incident also illustrates one participant taking up information that had previously been contributed by another (as indicated by the solid line).

At 15:52:21, almost a full minute after P1's deletion (they might have been absorbing what each other had just done), P2 links D26 to both his own D29 and P1's recent contribution D28, forming a cluster of related data. While P1 cleans up the wording of P2's recent contribution (D29 at 15:52:31 and 15:52:52) to parallel that of D28 and make the connection to disease status more explicit, P2 (15:52:47) makes the evidential relationship to the aluminum hypothesis H02 explicit – again performing a reintroduction of an item originally introduced by P1. At 15:53:25, P2 joins P1's cleanup activities, by merging data items D05 and D06. After moving some things around, participants verbally acknowledge the shared interpretation that they have achieved by jointly manipulating the workspace: “boy, we got something”; “heheh ALUMINUM!!!!”

It is clear that participants were collaborating through the evidence map, each taking up information that was introduced by the other participant (evidence of grounding), and jointly constructing and transforming representations to arrive at a conclusion (evidence of distributed cognition). Although the role distribution is asymmetric, the collaboration constitutes a form of knowledge construction in which they use the evidence map notation to come to agreement on the structure of evidence and its implication for a hypothesis under consideration. The verbal chat is only used to acknowledge what has been accomplished through the evidence map. Apparently, participants feel the need for a more explicit verbal form of grounding to mark the completion of their negotiated interpretation. This concludes our case example, but brings us to one final observation.

4.3. Distribution of Collaboration across Chat and Evidence Map

Although the focus in these case examples has been on how knowledge construction may be accomplished through joint manipulation of a visual workspace such as an evidence map, both notational and linguistic media were used by all pairs. Most of the task-oriented interaction took place through the evidence map, although chat at times played a crucial role in supporting the communication. The evidence map was used primarily for what its representational primitives support: reporting and recording information gleaned from the source pages, proposing hypotheses, and indicating consistency and inconsistency relationships between these items. The evidence map was the primary means of accomplishing these communications, although there are a few examples of chat that could have been accomplished via the graphical notations provided. Participants discussed the problem extensively in the chat in only one of ten sessions. Several groups engaged in extended evaluative/interpretive discussions after reaching the final page, which announced that participants' “library research” was done.

Some pairs used chat primarily for social banter as they carried out task-oriented interactions in the graph. Typically this social use of chat was occasionally punctuated with task-oriented chat, such as role assignments (“you write the new part, ima edit the old one”) and coordination of page turning (“ready?” “next?”, etc.). Occasionally, brief chat exchanges during the session would focus on the value or interpretation of information, especially when a problematic situation arose. Often the interaction in such cases was multi-modal, involving use of both visual representations and chat. For example, in the portion of Table 1 omitted from the figure for Example 2 (14:33:24-14:34:51), participants discuss what to make of some new information that causes problems for one of their previous hypotheses. A brief verbal exchange is summarized by P1's action in the evidence map. This kind of movement from verbal discussion to visual representation was typical of the conversations in face-to-face studies (Suthers & Hundhausen, 2003), but less typical online, where participants more typically

generated a proposal in the workspace and then discussed it (Suthers et al., 2003). Often, a shift of domain reasoning from the workspace to the chat medium signaled an exceptional or problematic situation that requires collaborative meta-cognition, best accomplished via language due to its flexibility and self-referential power, supporting second-order modal statements about beliefs.

5. Conclusions

Previous research by the author had found that visual representations are used differently online than face-to-face, with more of the communicative function switching to visual representations online. The present work was motivated by a desire to understand, from a qualitative perspective, how participants use shared external representations to support their knowledge construction. A case-analysis was undertaken on transcript segments in which participants acted intensively on the graph. In order to understand this interaction, we developed a three-step method that identifies basic actions in the media; identifies “information uptake” events, building a graph of uptake relations; and then applies theory to interpret sequences and patterns of uptake in the graph as knowledge construction episodes. The uptake graph is constructed with minimal theoretical inference so that it may support the application of multiple theoretical orientations needed to fully understand online collaborative knowledge construction.

The present paper reports this method, and uses two case studies to illustrate its application, as well as to begin to address the original questions concerning the use of representations as communicative media. The first example showed how argumentation is possible through an evidence map, and the second example showed how agreement could be reached through joint manipulation of such a medium. Interaction through the evidence map displays many of the criteria for knowledge construction suggested by theory, including grounding by implicit uptake of the interlocutor’s actions in the graph, interactions that respond to and address differences of interpretation, and transformations of representations by multiple individuals leading to a joint solution. Examining how cognition is distributed through the visual and verbal representations, we found that in these cases much of the interaction concerning hypotheses and evidence (i.e., those supported by our particular tool) take place through the evidence mapping tool. Linguistic representations are used both as a back channel for social and task-coordination interactions, but are relied on for domain reasoning when critical events occur requiring second-order or modal evaluation of the propositions. Systematic quantification of the distribution of acts using a synthesis of qualitative and quantitative methodologies would be needed to test these observed trends, but it is clear that designers need to assume that collaboration will be distributed across all media that is both shared and modifiable.

Interaction through visual representations is different than language: although we demonstrated that true collaboration (and perhaps knowledge-building) is possible through visual as well as linguistic representations by interpreting actions on the former as if they were language acts, there is no reason to believe that the structure of interactions through a graph are necessarily isomorphic to linguistic discourse. At a minimum, coherence based on adjacency must be reconsidered.

Future work includes generalizing the notations beyond dyads and to asynchronous interaction; raising the level of analysis further to identify recurring patterns of interaction that take place in each notational medium; and clarifying the roles of each medium of interaction as well as how to coordinate the two effectively. As the work becomes more complex, computer tools in support of analysis will be necessary.

Some refinement to the methodology and theory is already underway. After conducting the work reported here, we realized that attention and attitude are also forms of uptake. By merely considering an information item, one is directing the attention of the group. Expressions of attitude influence what the group is going to do with the expressed information. Therefore our concept of uptake is no longer restricted to information uptake, although that was our focus when this analysis was undertaken. Also, in the analysis reported here, movement of digital objects was not analyzed, although we realize that spatial arrangements can express implicit interpretations. Finally, current work in an asynchronous setting has challenged assumptions about the mutual availability of expressed conceptions and required further work on our uptake notation.

It has been suggested that our results might have been different with different software interfaces, tasks, user populations, etc. This is undoubtedly true, and suggests lines of further investigation that we hope our methodology will enrich. We do believe that the results that it is possible to collaborate through media such as an evidence map and that there will be preferences for distribution of activity across graphical and linguistic media will stand, but we look forward to a research program that works out the details of how the affordances of computer media are appropriated by learners to accomplish their objectives and explores the implications for design (Suthers, 2006).

Although the present paper did not attempt to break new theoretical ground, we are currently refining a working definition that knowledge construction is evidenced by the *composition of interpretations* on a history that is simultaneously expanded by information seeking and transformations (Suthers, 2006). The act of interpretation may take the form of explicit sense-making commentary, but it may also take place through the transformation and integration of representations, as illustrated in this paper. Each node in the uptake graph is understood as corresponding to an act of reinterpretation of reifications of prior acts of interpretation in a constant interplay between participation and reification (Wenger, 1998). Then, *collaborative* knowledge construction takes place when multiple participants contribute to a shared history by building, commenting on, transforming and integrating prior acts of participation and their traces in the media. Suthers (2006) also proposes replacing “knowledge construction” and even “learning” with *intersubjective meaning-making* as the most productive analytic concept for further work, as “knowledge construction” has historical baggage and “learning” is a post-hoc judgment of the consequences of an activity, not the essence of what participants are trying to accomplish in the activity itself.

6. Acknowledgements

The author gratefully acknowledges the contributions of Chris Hundhausen and Laura Girardeau, who conducted the study sessions; Ravikiran Vatrappu, who assisted in the construction of the uptake graphs; and Nathan Dwyer, who has broadened and clarified my thinking through many discussions. Tim Koschmann and Gerry Stahl influenced the author’s reorientation towards qualitative analysis. This work was supported by the National Science Foundation under award 0093505.

7. References

Andriessen, J., Baker, M. & Suthers, D. (Eds.) (2003). *Arguing to Learn: Confronting Cognitions in Computer-Supported Collaborative Learning Environments*. Dordrecht: Kluwer.

- Blackwell, Alan & Green, Thomas (2003). Notational Systems--The Cognitive Dimensions of Notations Framework. In J. M. Carroll (Ed.), *HCI Models, Theories and Frameworks: Towards a Multidisciplinary Science*. San Francisco, Morgan Kaufmann: 103-133.
- Clark, H.H. & Brennan, S.E. (1991). Grounding in Communication. In L.B. Resnick, J.M. Levine and S.D. Teasley (eds.), *Perspectives on Socially Shared Cognition* (pp. 127-149). Hyattsville, MD: American Psychological Association.
- Cole, M. and Y. Engeström (1993). A cultural-historical approach to distributed cognition. In G. Salomon. (Ed.) *Distributed Cognitions: Psychological and Educational Considerations*. Cambridge, Cambridge University Press: 1-46.
- Collins, A. & Ferguson, W. (1993). Epistemic forms and epistemic games: Structures and strategies to guide inquiry. *Educational Psychologist*, 28(1), 25-42.
- Dillenbourg, P. (2005). Designing biases that augment socio-cognitive interactions. In R. Bromme, F. W. Hesse and H. Spada (Ed.s) *Barriers and Biases in Computer-Mediated Knowledge Communication—and How They May Be Overcome*. New York, NY, Springer: 243-264.
- Doise, W., and Mugny, G. (1984) *The Social Development of the Intellect*, International Series in Experimental Social Psychology, vol. 10, Pergamon Press.
- Dwyer, N & Suthers, D. (2005). A Study of the Foundations of Artifact-Mediated Collaboration. In T. Koschmann, D. Suthers & T. W. Chan (Eds.), *Computer Supported Collaborative Learning 2005: The Next 10 Years!* Mahwah, NJ: Lawrence Erlbaum Associates. pp. 135-144.
- Festinger, L. (1957). *A Theory of Cognitive Dissonance*, Stanford University Press, 1957.
- Goodwin, C., & Heritage, J. (1990). Conversation Analysis. *Annual Review of Anthropology*, 19, 283-307.
- Grosz, B. J. and C. L. Sidner (1986). Attention, Intentions, and the Structure of Discourse. *Computational Linguistics* 12(3): 175-204.
- Herring, S.C. (1999, January). Interactive coherence in CMC. In *Proceedings of the 32nd Hawai'i International Conference on the System Sciences (HICSS 32)*. (CD-ROM). Maui, Hawai'i: Institute of Electrical and Electronics Engineers, Inc. (IEEE).
- Hollan, J., E. Hutchins, & Kirsh, D. (2002). Distributed Cognition: Toward a New Foundation for Human-Computer Interaction Research. In J.M. Carroll (Ed.), *Human-Computer Interaction in the New Millennium*. New York, ACM Press Addison Wesley: 75-94.
- Hollan, J. & Stornetta, S. (1992). Beyond being there. *Proc. SIGCHI Conference on Human Factors in Computing Systems (CHI'92)*, Monterey, California, pp. 119-125.
- Hutchins, E. (1995). *Cognition in the Wild*. Cambridge, Massachusetts: The MIT Press.
- Lieberman, A. (2004) What You Should Know About Guam, Plants, Flying Foxes and Parkinson Disease. National Parkinson Foundation.
<http://www.parkinson.org/site/pp.aspx?c=9dJFJLPwB&b=99904>, accessed April 22, 2006
- Lyons, J. (1977). *Semantics, Volume 2*. Cambridge University Press.

- Matusov, E. (1996): Intersubjectivity Without Agreement. *Mind, Culture and Activity*. 3(1). 25-45.
- Monk, A. (2003). Common Ground in Electronically Mediated Communication: Clark's Theory of Language use. In J. M. Carroll (Ed.), *HCI Models, Theories and Frameworks: Towards a Multidisciplinary Science*. San Francisco, Morgan Kaufmann: 265-289.
- Mühlenbrock, M., & Hoppe, U. (1999). Computer Supported Interaction Analysis of Group Problem Solving. *Proc. Computer Support for Collaborative Learning Conference (CSCL '99)*, C. Hoadley & J. Roschelle (Eds.) Dec. 12-15, Stanford University, Palo Alto, California. Mahwah, NJ: Lawrence Erlbaum Associates.
- Olson, G. M & Olson, J. S. (2000). Distance Matters. In J. M Carroll (Ed) *Human-Computer Interaction in the New Millennium*. New York: ACM Press (2002). pp. 397-417.
- Roschelle, J. (1994, May). Designing for cognitive communication: Epistemic fidelity or mediating collaborative inquiry? *The Arachnet Electronic Journal of Virtual Culture*, 2(2).
- Sacks, H. (1984). Notes on methodology. In J. M. Atkinson & J. Heritage (Eds.), *Structures of Social Action* (pp. 21-27). Cambridge: Cambridge University Press.
- Sacks, H., Schegloff, E. A., & Jefferson, G. (1974). A Simplest Systematics for the Organization of Turn-Taking for Conversation. *Language*, 50(4), 696-735.
- Sanderson, P. M., & Fisher, C. (1994). Exploratory sequential data analysis: Foundations. *Human-Computer Interaction*, 9(3-4), 251-318.
- Stahl, G. (2006) *Group Cognition: Computer Support for Collaborative Knowledge Building*. Cambridge, MA: MIT Press.
- Stahl, G. (2006, this volume). Sustaining Group Cognition in a Math Chat Environment. *Research and Practice in Technology Enhanced Learning*, 1(2).
- Stahl, G., Koschmann, T, & Suthers, D. (2006). CSCL: An Historical Perspective. In R. K. Sawyer (Ed.). (2006). *Cambridge Handbook of the Learning Sciences*. Cambridge, UK: Cambridge University Press.
- Star, S. L. (1990). The structure of ill-structured solutions: boundary objects and heterogeneous distributed problem solving. In L. Gasser and M. N. Huhns (Eds.) *Distributed Artificial Intelligence: Vol. 2*. (pp. 37-54). San Francisco: Morgan Kaufmann.
- Suthers, D. (2006). Technology Affordances for Intersubjective Meaning-making: A Research Agenda for CSCL. To appear in *International Journal of Computer Supported Collaborative Learning* 1(2).
- Suthers, D, Connelly, J., Lesgold, A., Paolucci, M., Toth, E., Toth, J., and Weiner, A. (2001). Representational and Advisory Guidance for Students Learning Scientific Inquiry. In K. D. Forbus and P. J. Feltovich (Eds.) *Smart Machines in Education: The Coming Revolution in Educational Technology* (pp. 7-35). Menlo Park: AAI Press.
- Suthers, D., Girardeau, L. and Hundhausen, C. (2003). Deictic Roles of External Representations in Face-to-face and Online Collaboration. Designing for Change in Networked Learning Environments, *Proc. International Conference on Computer Support for Collaborative*

Learning, B. Wasson, S. Ludvigsen & U. Hoppe (Eds), Dordrecht: Kluwer Academic Publishers, pp. 173-182.

Suthers, D. & Hundhausen, C. (2003). An Empirical Study of the Effects of Representational Guidance on Collaborative Learning. *Journal of the Learning Sciences*, 12(2), 183-219.

Suthers, D., Vatrapu, R., Joseph, S., Dwyer, N & Medina, R. (2006). Representational Effects in Asynchronous Collaboration: A Research Paradigm and Initial Analysis. In *Proceedings of the 39th Hawai'i International Conference on the System Sciences (HICSS-39)*, January 4-7, 2005, Kauai, Hawai'i (CD-ROM), Institute of Electrical and Electronics Engineers, Inc. (IEEE).

Vygotsky, L. S. (1978). *Mind in Society: The Development of Higher Psychological Processes*. Cambridge, MA: Harvard University Press. (Originally published in 1930.)

Wenger, E. (1998). *Communities of Practice: Learning, Meaning and Identity*. Cambridge: Cambridge University Press.