Information Sharing is Incongruous with Collaborative Convergence: The Case for Interaction

Daniel Suthers, Richard Medina, Ravi Vatrapu, Nathan Dwyer
Laboratory for Interactive Learning Technologies
Dept. of Information and Computer Sciences, University of Hawai`i at Manoa
collaborative-representations@hawaii.edu

Abstract: Various authors have placed information sharing at the core of successful collaborative problem solving and learning. In this paper we report analyses of an experimental study that bring the sufficiency of an information sharing account of collaboration into question. One treatment group achieved greater convergence and integration of information in their handling of a complex problem, yet this same group shared *less* information in a hidden profile design. An additional analysis was conducted to assess whether interaction beyond information sharing accounts for the convergence and integration. Results are not conclusive, but suggest that further study is merited. The study illustrates how interaction may be quantified as a dependent variable in the quantitative experimental paradigm, and also illustrates a limitation of this paradigm. Interaction analysis will be needed to account for participants' accomplishments: a dialogue between methodological paradigms is in order.

Introduction

A central tenet of much research on group problem solving and learning in CSCL and related fields is that information sharing is the primary operative mechanism of effective group performance. For example, an influential theory of linguistic communication is concerned with the process by which interlocutors verify that they have successfully shared information (Clark & Brennan, 1991). In social psychology, a major (and productive) research strategy is the "hidden profile" (Stasser, 1992) in which information is distributed across participants and then group processes are tracked and evaluated in terms of how this information is shared. Common findings include the failure to share information and the failure to use information effectively once it has been shared (Dennis, 1996). Dennis (1996) states, "In order to reach a group decision, participants engage in three activities simultaneously …: information recall (either from memory or notes), information exchange (either giving or receiving information), and information processing (actually using the information …)." Thus, collaboration is characterized primarily in terms of the movement of pre-existing information between cognitive agents so that it may be properly assembled and evaluated. Similarly, Pfister (2005) tells us that "going from unshared to shared information is the gist of cooperative learning." While there is an important truth in this statement, let us examine its presuppositions: information exists that is first unshared, and then it becomes shared. Once this pre-existing information is shared, the important work of cooperation has been done. The focus is on the movement of information between individuals, but we might also consider how information is constructed in the interaction between individuals.

To be fair, perhaps all researchers who work in the information sharing paradigm would agree that there is more to collaborative learning and problem solving than information sharing alone, and that interaction between participants plays a role in building something new beyond the information that individuals held at the outset. Yet, much empirical work in CSCL (as well as some of its sister fields) remains focused on information sharing, while we lack an equally comprehensive research program on whether and how interaction adds value for collaborative learning beyond information sharing. Exceptions include (Baker, 2003; Enyedy, 2005; Koschmann, Zemel et al., 2005; Roschelle, 1992; Stahl, 2006). The strategy taken by this paper as a contribution to the ongoing methodological and theoretical dialogues within CSCL is to demonstrate that it may be profitable for those working in an experimental paradigm to examine interaction in order to account for quantitative results, and to seek alliances with those who work in analytic paradigms. As a topic of study, interaction has potential to unify our field by being the shared object of analysis between researchers in multiple methodological traditions.

The analyses presented in this paper were motivated by an interesting combination of empirical results obtained in an experimental study that was based on the hidden profile paradigm (Suthers, Vatrapu, Medina, Joseph,

& Dwyer, submitted). Pairs in one treatment condition performed better on measures related to collaborative knowledge construction: integration of multiple sources of information and convergence on similar solutions. From this, one would expect that the pairs in this treatment condition also shared more information. Problematically, the treatment conditions did not differ in information sharing as evidenced by the information that participants referenced in their essays, nor on their memory for facts one week later. Those results were based on measures of the products of the experimental sessions (essays and a post-test): more direct measures of information sharing were needed. For this, we turned to a different source of data: the session logs.

In the follow-up study reported in the present paper, we measured the information sharing that took place in the sessions by tracing information that was given to only one or the other participant at the outset. Surprisingly, we found that pairs in the higher performing condition shared *less* information in the session: a serious challenge to the information-sharing explanation of group performance. An alternative explanation was needed, for which we turned to interaction. In information sharing, a participant expresses something in some medium and this expression becomes available to another participant. The smallest way in which interaction can go beyond this basic act is a "round trip" of uptake: the second participant takes up that which was expressed by the first participant by forming a new, related expression, which then becomes available to the first participant. Accordingly, we measured interaction in terms of these round trips. By this measure, participants in the higher performing treatment condition (which shared less information) interacted more than participants in the other conditions, although the probability of this result was not low enough to reject the null hypothesis under traditional criteria. However, the incongruence of the information sharing along with the congruence of round trips and other measures to be discussed suggests that it is worth examining the practices by which participants integrate multiple sources of information and converge on common solutions.

The remainder of the paper serves to provide the reader with a more detailed account of how the results were obtained. First we briefly review the experimental context in which this work was done, and summarize the pattern of results on integration and convergence that indicated the need to conduct this follow-up study. Then we describe the two analyses of the present study: first, the information-sharing analysis, which showed an unexpected pattern across treatment conditions; then the round-trip analysis, which revealed a pattern that is congruent with the original results. We conclude with a discussion of both theoretical and methodological implications for further research in CSCL.

The Original Study

The original study was designed to test the hypothesis that conceptual representations (such as evidence maps) more effectively support collaborative knowledge construction than threaded discussion taken alone. A threaded discussion software environment, called "Text," provided the control condition. Since the viability of the hypothesis may be sensitive to the implementation chosen, two software environments represented the treatment condition: "Graph," in which all interaction took place in an evidence map with embedded notes; and "Mixed," in which an evidence map was used alongside a separate threaded discussion with a mechanism for referencing components of the graph.

The present paper is concerned with how well information sharing and interaction account for a pattern of results found in the prior study, rather than with the specific variables over which these results were obtained. We refer the reader to (Suthers, Vatrapu, Medina, Joseph, & Dwyer, in press; Suthers et al., submitted) for details of the study design, but summarize here enough for the reader to understand the source and nature of the data.

Method

Pairs of participants who were already aquatinted with each other were recruited from introductory natural science courses and assigned to one of three conditions (Graph, Mixed, Text) in a manner that was gender-balanced but was otherwise randomized. The groups were statistically equivalent on academic grade point average and standardized test scores. There were 60 participants forming 30 pairs, with 10 pairs in each of 3 treatment conditions. They were paid $50 US for their participation.

The software for the three conditions was designed with an asynchronous update protocol to simulate asynchronous interaction common in online learning (Mayadas, 1997). An action taken by one participant did not appear in the other participant's workspace until after the receiving participant "took a break" by playing a game of Tetris™.

Materials were prepared based on the professional literature concerning a complex public health problem: the disease known as "ALS-PD" that historically occurred in the native population on the island of Guam. The materials suggested 5 distinct possible causes of the disease, and provided mixed evidence for and against each cause. Relevant evidence was distributed in a hidden profile such that if participants did not share any information each participant would have evidence favoring a suboptimal disease hypothesis. Sharing was required to reject these hypotheses and construct a more complex explanation. In each dyad, Participant 1 (P1) received evidence for aluminum in the water and against genetic causes; Participant 2 (P2) received evidence against aluminum and for genetic causes; and both participants received evidence for and against cycad seeds as the source of a neurotoxin as well as crucial information about native diets that, when brought together, points to seed-eating bats as the vector by which this toxin gets into humans. The articles included distracter information as well as relevant evidence.

Participants conducted a "warm up" problem to become familiar with the software, and then began the main problem (ALS-PD). At the outset of the main problem and after each break, each participant was presented with a set of four articles on the disease. Participants were directed to use the computer workspace to share information with their partner, and were told that this was necessary to identify the correct cause of the disease and to perform well on the essay and post-test to be given at the end. At the conclusion of their problem solving, each individual was asked to write an essay detailing the disease hypotheses considered and the evidence for and against those hypotheses, and to identify the best explanation for the disease. One week after their session, participants were directed to take an online post-test. This test included questions that tested participants' memory for distracter information, memory for relevant information, and also tested for facts that required integration of multiple items of relevant information. "High integration" questions required integration of information that occurred far apart in the materials (in Suthers & Hundhausen, 2003's terms, there is a large "inferential span"). The questions were based on information given uniquely to one or the other participant, enabling us to assess the residue of information sharing.

Prior Results

Our analyses addressed *outcomes,* based on content analyses of the essays and scoring of the post-test; and *session processes,* based on quantitative analyses of elaboration on hypotheses. Details of these analyses are reported in (Suthers et al., in press; Suthers et al., submitted). The traditional criterion of $\propto \leq 0.05$ is used for statistics computed to test hypotheses. However, we are sympathetic with the view that probabilities express properties of the data to be reasoned about, and are not merely inputs to a mechanical binary decision procedure (Gigerenzer, 2004). Since this study seeks to uncover possible alternative explanations as well as test hypotheses, we report and interpret p values of 0.1 and below.

The treatment conditions did not differ under χ^2 in optimality of conclusion in the essays: relatively few participants in all conditions identified the bats-as-vector explanation for how the cycad toxin gets into humans. Pairs in the Graph condition were more likely to converge on the same (not necessarily optimal) conclusion than pairs in the other conditions ($\chi^2(2, N=30)=7.5$, $p=0.025$): see Figure 1. This suggested that Graph users may have shared more information, but analysis of essay contents did not back up this interpretation: participants in all conditions were equally likely to cite information that was originally given to their partner. Finally, Graph users performed significantly better than Mixed users on the high integration questions of the post-test ($F(2,57)=4.40$, $p=0.0167$), suggesting that they were able to more effectively bring relevant and distributed information together. However, comparison of participants' performance on memory for information that they received versus memory for information given to their partners yielded no statistically significant difference, again suggesting that information sharing was not the operative mechanism.

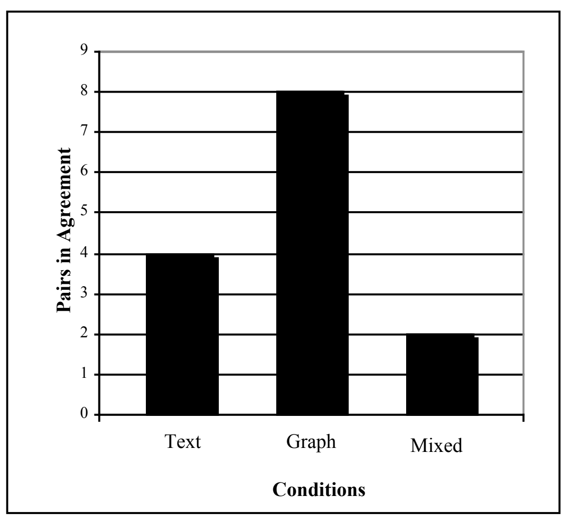

Figure 1. Convergence (Pair Agreement) Results

Process analyses in the prior study used one-way ANOVAs to assess the expression and manipulation of hypotheses. Graph and Mixed users expressed hypotheses significantly earlier in the sessions than Text users

($F(2,57)=10.14$, $p=0.0002$), and Graph users expressed more hypotheses than Text users ($F(2, 57)=4.73$, $p=0.0126$). Graph and Mixed users elaborated on hypotheses significantly more than participants in the Text condition ($F(2, 57)=6.86$, $p<0.0021$). These analyses counted individual acts in isolation, and did not directly address information sharing or interaction. However, the individual acts were undertaken in a shared medium, so these results may be indicative of differences in interaction. Therefore, interactive measures should be examined.

The Information Sharing Analysis

We have just summarized a pattern of results in which pairs of users of the Graph software scored higher on posttest questions requiring integration of information that was distributed between them, and more frequently converged on the same conclusions. Both of these results would be unsurprising if Graph participants shared more of the uniquely distributed information, yet their essays did not display any greater reliance on this information than those of the other conditions, and the post-test did not provide any evidence of differences in information sharing. However, the essays and post-test are only indirect measures of information sharing. We need to examine session data. Perhaps Graph participants achieved their integration and convergence by sharing more information during the session? We undertook an analysis to test this hypothesis.

Quantifying Information Sharing

Since our analyses are concerned with tracing out information that was uniquely provided to one or the other participant in the source materials, the unit of analysis is defined in terms of "information units" as expressed in the sentences and figures of the source materials. Participant 1 was provided with 226 information units whereas Participant 2 was provided with 229 information units. Both the participants received the same mission statement, end of study statement and a few overlapping information units. Controlling for these, there are 401 total information units that were uniquely provided to one or the other participant.

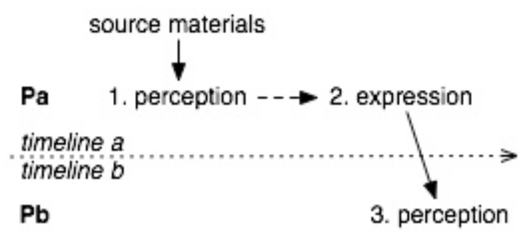

Figure 2. Uptake graph of an information sharing event. Solid arrow represents intersubjective uptake and dashed arrow represents intrasubjective uptake. In asynchronous interaction there are two timelines that interact only at workspace synchronizations.

Figure 2 schematizes an information sharing event in terms of an "uptake graph" (Dwyer, Suthers, & Vatrapu, submitted; Suthers, Dwyer, Vatrapu, & Medina, 2007). An information sharing event consists of the sequence in which (1) Pa perceives information that had been given uniquely to him or her (only unique information is considered so that the analysis will be both possible and relevant), (2) Pa expresses that information (in this study, by posting a message in the threaded discussion or creating or editing an object in the evidence map), and (3) Pb perceives that expression. In order to count such events, we coded expressive acts that were recorded in the session logs with codes identifying the initially unshared information units that were being expressed. It was not necessary to identify what was said about the information; only that it had been expressed. We then identified when the media expression became available in the workspace of the other participant, Pb, and identified evidence that Pb perceived that expression. The evidence for perception involved a log file entry indicating that the object was opened (it was necessary to open threaded discussion postings or new graph objects in order to read them). The total number of such events was summed for each pair. (The dyad is the unit of all analyses.)

Results

The content analysis for information sharing was performed in three phases. Both the coders followed the same coding procedure and used the same coding manual. In the first phase, two analysts independently carried out content analysis on 20% of the randomly selected experimental session content logs (20% of 30 sessions = 6 sessions of dyads, or 12 content logs were analyzed). We then calculated an inter-rater reliability metric (IRM) using the following formula:

$$IRM = Average(IRL_i) \qquad \text{where} \qquad IRL_i = \frac{|RaterACodes \cap RaterBCodes|}{|RaterACodes \cup RaterBCodes|} \bullet 100$$

for each log file i. The IRM computed after the first phase was a medium-low of 64.4. We then discussed the conflicts and made appropriate changes to the coding process. Each analyst incorporated these changes to content logs and this resulted in an improved IRM of 71.1. In the second phase of this analysis one analyst conducted content analysis of the remaining (60-12=48) logs. In the third phase, the second analyst conducted a content analysis on randomly selected 20% of the logs analyzed in phase two by the first analyst (10 logs). We obtained an IRM of 81.2.

Recall that 401 information units were uniquely provided to one or the other participant. These define the total possible information sharing events under this analysis. Tracing through Figure 2, in order to be perceived the information must first be expressed, so we report both events in order to provide the baseline number of units available for perception. More *expressions* of the information units were made in the Text condition compared to Mixed and Graph conditions. Closely following this pattern, more *perceptions* of these expressed information units were made in the Text condition compared to Mixed and Graph conditions (see Table 1 for descriptive statistics). These results are visualized in Figure 3. We then conducted a one-way analysis of variance (ANOVA) of perceptions of information units between the three treatment groups (Table 1 below). The probability of the observed differences in perceptions ($F(2, 27)=13.54$, $p<0.0001$) would be very low if there were no differences between conditions on information sharing. Importantly, the Bonferroni 95% confidence interval showed that Graph had *fewer* information sharing events than Text.

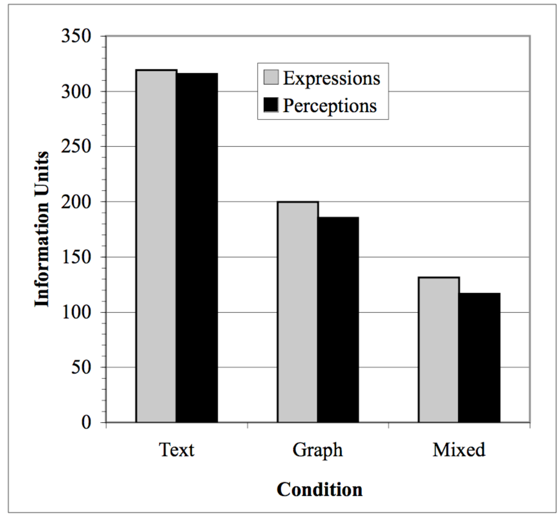

Figure 3. Information sharing results

Table 1. One-way ANOVA of perception of information units

n	30				
Session Perceptions by Condition	**n**	**Mean**	**SD**	**SE**	
Graph	10	185.100	95.919	30.3324	
Mixed	10	116.500	82.453	26.0739	
Text	10	315.700	81.753	25.8526	
Source of variation	**SSq**	**DF**	**MSq**	**F**	**p**
Condition	204809.867	2	102404.933	13.54	<0.0001
Within cells	204143.500	27	7560.870		
Total	408953.367	29			
Contrast	**Difference**	**Bonferroni 95% CI**			
Graph v Mixed	68.600	-30.657	to 167.857		
Graph v Text	-130.600	-229.857	to -31.343	(significant)	
Mixed v Text	-199.200	-298.457	to -99.943	(significant)	

Discussion

From this analysis, it seems implausible that Graph's convergence and integration outcomes were due to greater information sharing during the session. The distributions in Figure 1 and Figure 3 are completely different. It would have been problematic enough for an information sharing account if there were no significant differences between groups. The result that the Graph participants actually shared *fewer* information items than Text users is more problematic. What could Graph users be doing that led to their greater integration and convergence?

It is possible to provide an individual account of these results based on considerations of representational guidance. Although less information was shared in Graph, perhaps the crucial information *was* shared (consistent with Suthers & Hundhausen's (2003) finding of greater selectivity in the use of evidence maps) and the graph itself enables each individual to get an overview of and then integrate this information (Novak, 1990). Perhaps also the graph makes the same conclusion visually obvious to individual participants, leading them to converge without needing to interact with each other. The lack of similar benefits for Mixed users could be explained by the multi-representational workspace, which separated information that needed to be integrated and made it less likely that participants would look at the same thing (Ainsworth, Bibby, & Wood, 1998). We believe that these explanations identify pertinent factors, but that they do not rule out the possibility that Graph participants are also achieving integration and convergence through collaborative interaction, as well as individual access to the representations. We turn next to a more direct test of this hypothesis.

The Round Trip Analysis

"Interaction" is potentially a complex idea: it includes the basic act that we are calling "information sharing" and extends to diverse forms of discourse. To conduct a quantitative analysis we need to identify the simplest possible unit of interaction that is distinguishable from information sharing. Given that we have defined information sharing as the expression by a participant Pa of an idea related to a topic that is perceived by the interlocutor Pb, the next interactive step that can be taken beyond information sharing is for the interlocutor Pb to express a related idea that is then perceived by the originating participant Pa. In this "round trip," intersubjectivity forms: the subject has expressed and seen his or her expression interpreted by the other. Therefore we set out to define and count round trips.

Quantifying Round Trips

In general, a round trip involves the sequence of events shown in Figure 4: (1) participant Pa expresses an idea in a shared medium; (2) this expression becomes available to participant Pb, and Pb perceives the expression; (3) Pb expresses a related idea in the medium; (4) this second expression becomes available to and is perceived by Pa. One can imagine writing a grammar for interaction patterns such as that shown here and using automated support to recognize these patterns in a session protocol (Olson, Herbsleb, & Rueter, 1994). That is one of the intended applications of the uptake graph transcript notation (Dwyer et al., submitted; Suthers et al., 2007).

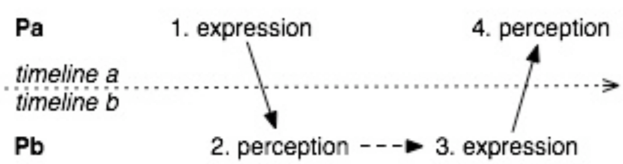

Figure 4. Uptake graph for a round trip event

However, it turns out that a restriction on the present analysis enables a much simpler procedure for identifying round trips with the present data. In order to place this analysis on the same foundation as the information sharing analysis, we decided to include only round trips that involved an information item that was uniquely given to the originating participant (Pa). This is not a severe restriction on our measure of interactivity: it includes most of the information that is needed to reason about the disease. We then realized that, because of this restriction, round trips could be counted by identifying only events of type (4) instead of the entire path (1→4), for the following reasons. The log files only recorded perceptions of expressions that were created by the other person. Therefore, for each perception such as (4) involving a given topic there must exist an expression (3) by Pb involving this topic. If a given participant Pa (4) perceived an expression that was concerned with information that had only been provided to him or herself, then that expression (3) could only have been created by Pb by virtue of having previously perceived (2) an expression of that information by Pa (1), completing the round trip pattern. Therefore we know that (1), (2) and (3) exist without needing to identify them. The analysis was further simplified by the fact that the information sharing analysis had already coded all expressions with the source material information units they addressed. We wrote a database query that transferred this coding from the expression event (3) to the perception event (4).

There are some ambiguities in what constitutes a round trip. For example, suppose that more than one event of type (4) is labeled with the same topic Are these multiple round trips, or one? We are concerned with measuring the degree of interactivity, and it is certainly more interactive if an interlocutor continues to express further thoughts on a topic than if it is addressed once and then ignored. Iterative elaboration, questioning, accumulation of evidence, etc. are good for knowledge building. Therefore, we counted each event of type (4) as a new round trip.

Based on this reasoning, we wrote queries on the database generated by the information sharing analysis in which we counted the events in which each participant accessed an expression (from their partner) of a topic related to information that was introduced in their own materials. We then summed the number of such events per pair, and compared the average round trips between the treatment groups. Other round trips are possible involving other topics and forms of response that are not captured by our analysis. However, this analysis produces a measure that is directly proportional to the totality of interactions that involve the critical information originally given to one participant.

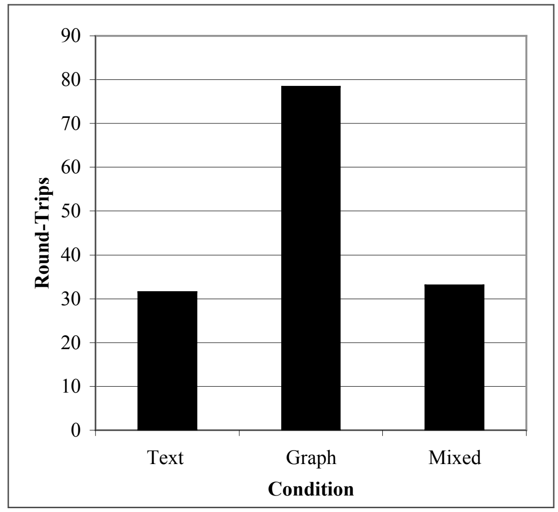

Figure 5. Round trips

Results

The results are visualized in Figure 5. More round trips were made in the Graph condition compared to Mixed and Text conditions, following the pattern of Figure 1. A one-way ANOVA (Table 2) on number of round trips between the three treatment groups suggests that these results are not likely if the groups were equivalent on interactivity ($F(2, 27)=3.03$, $p=0.0648$). Yet the pairwise differences did not fall within the Bonferroni 90% confidence interval.

Discussion

By the traditional criteria of $\propto \leq 0.05$, we cannot reject the hypothesis that there was no difference in round trips. Yet these results on round trips do not rule out the hypothesis that some aspect of interaction plays a role. We believe that the combination of results—more elaboration in the graph condition, a pattern of round trips that is unlikely yet congruent with the pattern of convergence we seek to explain, and an incongruent pattern of information sharing—rules out information sharing as an adequate explanation and is sufficient to suggest that interaction is worthy of further study as the basis for knowledge integration and convergence in collaborative learning.

Table 2. One-way ANOVA for round-trip analysis.

Round-Trips by Condition	**n**	**Mean**	**SD**	**SE**
Graph	10	78.400	71.070	22.4743
Mixed	10	33.000	31.875	10.0797
Text	10	31.500	30.974	9.7948

Source of variation	**SSq**	**DF**	**MSq**	**F**	**p**
Condition	14210.067	2	7105.033	3.03	0.0648
Within cells	63236.900	27	2342.107		
Total	77446.967	29			

Contrast	**Difference**	**Bonferroni 90% CI**	
Graph v Mixed	45.400	-3.137	to 93.937
Graph v Text	46.900	-1.637	to 95.437
Mixed v Text	1.500	-47.037	to 50.037

Conclusions

To summarize, in the context of a study designed to evaluate the value of conceptual representations for enhancing collaborative knowledge-building, we found greater convergence and integration in one condition, yet participant's essays and post-tests did not differ on information that would need to be shared. Further analysis showed that participants in the higher convergence and integration condition shared *less* information. This result brings the adequacy of an information sharing analysis of collaborative learning into question. Like Dennis (1996), we found that technologies that enable people to share more information do not necessarily lead to effective use of that information. Alternative explanations were considered. A strong candidate is level of elaboration: the group that shared the most information elaborated on hypotheses the least. The possibility that interaction between participants is behind the convergence was also considered. An analysis of interaction operationalized as "round trips" addressing unique information items was inconclusive from a hypothesis testing perspective but showed strong trends that call for further investigation. The round trip is the most minimalist definition of interaction beyond unidirectional information sharing: it is possible that more specific forms of interaction account for the difference in convergence. Given that there were significant results in the timing and number of hypotheses expressed and the elaboration on these hypotheses, and given that essays did not differ on information they relied on, it seems reasonable to expect that interaction was not focused on the information that we traced in this analysis, but rather on the hypotheses. Therefore, our next analysis will examine round trips focused on hypotheses rather than data (information units).

(Notes to Reviewers: (1) This analysis will require further annotation of the session data, but may be available before the publication version of this paper is due. (2) Close to the paper deadline we came across the closely related study by (Fischer & Mandl, 2005). We will incorporate comparison of results in the final version of this paper.)

The study is limited in that it was not specifically designed to test the hypothesis that round trip interaction rather than information sharing causes the results, but we feel that the results are compelling enough to pursue a program of further research. Some challenges face such a program within the experimental paradigm. In order to study the relative contributions of information sharing versus interaction, one must have a situation in which these activities vary. One approach would be to try to manipulate information sharing and interaction directly. Participation could be scripted (Weinberger, Reiserer, Ertl, Fischer, & Mandl, 2005) or otherwise modeled or constrained to require participants in one condition to share all the information given to them (but not interact), while participants in another condition are required to not only share information but also reply to each other's contributions. A problem with this approach is the artificiality of interaction that may result from scripting (Dillenbourg, 2002). Another approach, exemplified by this paper, would be to find variables that define treatment conditions known to differentially affect the dependent variables of interest (learning outcomes), and then conduct analyses of the information sharing and interaction that occurs "naturally" within the settings of these experimental conditions. Analyses can correlate measures of information sharing and of interaction with the dependent variables at the smallest granularity of the pair (or group) rather than at the granularity of treatment conditions.

As we have written elsewhere (Suthers et al., 2007), quantitative methods of "coding and counting" has two major limitations: (1) *"coding,"* the attempt to assign meaning to participants' acts as isolated units, loses the meaning constructed through interaction; and (2) *"counting"* aggregates individual practices into averages that may not represent any instance of actual practice. We have attempted to address some of these problems in the round trip analysis of this paper. Generalizing, the strategy is to define dependent variables in terms of sequences that capture some interesting aspect of the interaction. Through this strategy, experimentalism can become more relevant to the study of interaction. However, there are limitations exemplified by this paper. In order to set up an experimental analysis using interaction patterns such as the round trip, we had to define in advance the patterns of interest. The phenomena that experimentalism can subject to its formal methodology are restricted to those that are anticipated in advance. Also, our analysis has not identified what pairs are accomplishing through interaction. Without directly examining participants' interaction, we can only reason about the data and guess new possibilities to be tested. For example, we have just speculated that participants interacted most effectively about hypotheses, and have planned another analysis that finds and counts hypothesis-centered round trips. But why not just look directly at the data and see what participants are doing? We have done this, intensively analyzing several sessions over the period of a year in the process of developing an analytic methodology described elsewhere (Dwyer et al., submitted; Suthers et al., 2007). Yet the experimental approach will remain valuable in addressing questions of the influences of design features and the reliability of findings.

Practitioners of descriptively analytic paradigms can contribute to the research program by examining successful (and unsuccessful) instances of collaborative knowledge construction to identify the interactional practices by which integration and convergence are accomplished. The present study has provided a minimalist definition of a minimal unit of interaction, the round trip. It has not examined the nature of the round trip. In what way does the second participant interpret the first participant's contribution? What are the attentional, affective, informational, and interpersonal dimensions of this interpretation (Bronckart, 1995)? How does it frame future interaction and move the collaboration forward (Wells, 1999)? Given a choice of semiotic resources, what affordances of the medium of interaction are appropriated to enact contributions at each of these levels? Are there regularities we can exploit in such communicative use of media (Dwyer & Suthers, in press)? Nor has it examined interactional structures larger than the round trip. How are acts of intersubjective and intrasubjective uptake composed over time, and what is accomplished interactionally through such composition? Although the empirical work of this paper is conducted within the quantitative experimental paradigm, and demonstrates how interactional variables may be utilized in such a paradigm, we believe that further interactional analysis of the type advocated by recent writers (Dwyer et al., submitted; Koschmann, Stahl, & Zemel, 2005; Koschmann, Zemel et al., 2005; Stahl, submitted; Suthers, 2006) will be valuable to understand what is leading to quantitative results such as these: a dialogue between methodological paradigms is in order.

References

Ainsworth, S. E., Bibby, P. A., & Wood, D. J. (1998). Analyzing the costs and benefits of multi-representational learning environments. In M. W. van Someren, P. Reimann, H. P. A. Boshuizen & T. de Jong (Eds.), Learning with Multiple Representations (pp. 120-134). Amsterdam: Elsevier Science, Ltd.

Baker, M. (2003). Computer-mediated argumentative interactions for the co-elaboration of scientific learning tasks. In Andriessen, Baker & Suthers (Eds.), Arguing to Learn: Confronting Cognitions in Computer-Supported Collaborative Learning Environments. (pp. 47-78.). Dordrecht: Kluwer.

Bronckart, J. P. (1995). Theories of action, speech, natural language, and discourse. In J. V. Wertsch, P. D. Rio & A. Alvarez (Eds.), Sociocultural Studies of Mind (pp. 75-91). New York: Cambridge University Press.

Clark, H. H., & Brennan, S. E. (1991). Grounding in communication. In L. B. Resnick, J. M. Levine & S. D. Teasley (Eds.), Perspectives on Socially Shared Cognition (pp. 127-149): American Psychological Association.

Dennis, A. R. (1996). Information exchange and use in group decision making: You can lead a group to information, but you can't make it think. MIS Quarterly, 20(4), 433-457.

Dillenbourg, P. (2002). Over-scripting CSCL: The risks of blending collaborative learning with instructional design. In P. A. Kirschner (Ed.), Three worlds of CSCL- Can we support CSCL? (pp. 61-91). Heerlen: Open University Nederland.

Dwyer, N., & Suthers, D. D. (in press). Consistent Practices in Artifact-Mediated Collaboration. International Journal of Computer Supported Collaborative Learning, 1(4).

Dwyer, N., Suthers, D. D., & Vatrapu, R. (submitted). A Methodology for Analyzing Intersubjective Meaning-Making, submitted to Computer Supported Collaborative Learning. Rutgers.

Enyedy, N. (2005). Inventing mapping: Creating cultural forms to solve collective problems. Cognition and Instruction, 23(4), 427-466.

Fischer, F., & Mandl, H. (2005). Knowledge Convergence in Computer-Supported Collaborative Learning: The Role of External Representation Tools Journal of the Learning Sciences, 14(3), 405-441.

Gigerenzer, G. (2004). Mindless statistics The Journal of Socio-Economics, 33, 587-606.

Koschmann, T., Stahl, G., & Zemel, A. (2005). The Video Analyst's Manifesto (or The Implications of Garfinkel's Policies for Studying Practice within Design-Based Research. In S. Derry & R. Pea (Eds.), Video research in the learning sciences. Mahwah, NJ: LEA.

Koschmann, T., Zemel, A., Conlee-Stevens, M., Young, N., Robbs, J., & Barnhart, A. (2005). How do people learn: Member's methods and communicative mediation. In R. Bromme, F. W. Hesse & H. Spada (Eds.), Barriers and Biases in Computer-Mediated Knowledge Communication (and how they may be overcome) (pp. 265-294). Amsterdam: Kluwer Academic Press.

Mayadas, F. (1997). Asynchronous Learning Networks: A Sloan Foundation Perspective. Journal of Asynchronous Learning Networks, 1, http://www.aln.org/alnweb/journal/jaln_issue1.htm#mayadas.

Novak, J. (1990). Concept mapping: A useful tool for science education. Journal of Research in Science Teaching, 27(10), 937-949.

Olson, G. M., Herbsleb, J. D., & Rueter, H. H. (1994). Characterizing the sequential structure of interactive behaviors through statistical and grammatical techniques. Human-Computer Interaction, 9, 427-472.
Pfister, H.-R. (2005). How to support synchronous net-based learning discourses: Principles and perspectives. In R. Bromme, F. Hesse & H. Spada (Eds.), Barriers and Biases in Computer-Mediated Knowledge Communication--And How They May be Overcome. Dordrecht: Kluwer.
Roschelle, J. (1992). Learning by Collaborating: Convergent Conceptual Change. The Journal of the Learning Sciences, 2(3), 235-276.
Stahl, G. (2006). Group Cognition: Computer Support for Collaborative Knowledge Building. Cambridge, MA: MIT Press.
Stahl, G. (submitted). Meaning making in CSCL: Conditions and preconditions for cognitive processes by groups, Computer Supported Collaborative Learning 2007. Rutgeres.
Stasser, G. (1992). Information salience and the discovery of hidden profiles by decision-making groups: A "thought experiment". Organizational Behavior and Human Decision Processes, 52(1), 156-181.
Suthers, D. D. (2006). Technology affordances for intersubjective meaning-making: A research agenda for CSCL. International Journal of Computers Supported Collaborative Learning, 1(3), 315-337.
Suthers, D. D., Dwyer, N., Vatrapu, R., & Medina, R. (2007). An abstract transcript notation for analyzing interactional construction of meaning in online learning. In Proceedings of the 40th Hawai`i International Conference on the System Sciences (HICSS-34), January 3-6, 2007, Waikoloa, Hawai`i (CD-ROM): Institute of Electrical and Electronics Engineers, Inc. (IEEE).
Suthers, D. D., & Hundhausen, C. (2003). An experimental study of the effects of representational guidance on collaborative learning. Journal of the Learning Sciences, 12(2), 183-219.
Suthers, D. D., Vatrapu, R., Medina, R., Joseph, S., & Dwyer, N. (in press). Beyond threaded discussion: Representational guidance in asynchronous collaborative learning environments. Computers & Education.
Suthers, D. D., Vatrapu, R., Medina, R., Joseph, S., & Dwyer, N. (submitted, July 16-21, 2007). Conceptual representations enhance knowledge construction in asynchronous collaboration. Paper presented at the Computer Supported Collaborative Learning 2007, Rutgers.
Weinberger, A., Reiserer, M., Ertl, B., Fischer, F., & Mandl, H. (2005). Facilitating collaborative knowledge construction in computer-mediated learning environments with cooperation scripts. In R. Bromme, F. W. Hesse & H. Spada (Eds.), Barriers and Biases in Comptuer-Mediated Knowledge Communication -- And How They May Be Overcome. New York: Springer.
Wells, G. (1999). Dialogic Inquiry: Toward a Sociocultural Practice and Theory of Education. New York: Cambridge University Press.

Acknowledgments

David Burger, Sam Joseph and Niels Pinkwart contributed to the design of the prior study on which this study is based. This work was supported by the National Science Foundation under CAREER award 0093505.