Group cognition: The collaborative locus of agency in CSCL

Gerry Stahl
College of Information Science & Technology
Drexel University, Philadelphia, USA
www.cis.drexel.edu/faculty/gerry
Gerry.Stahl@drexel.edu

Abstract. CSCL faces the challenge of not only designing educational technologies and interventions, but of inventing analytic methodologies and theoretical frameworks appropriate to the unique character of collaborative learning as an interactional group accomplishment. This paper argues that thinking in CSCL settings should be primarily analyzed at the small-group unit of analysis, where contributions coming from individual interpretive perspectives are interwoven into group cognition. The collaborative discourse is the agent of knowledge building that requires computer support and curriculum design. Groups can think; with the help of CSCL in the next decade, they may be able to overcome the limitations of the individual mind.

Keywords: agency, shared knowledge, intersubjectivity, dialogism, group cognition, discourse, AI, emergence, Turing, Searle, Dreyfus

In the past decade, CSCL has grown willy-nilly out of various theoretical and methodological traditions that are mutually incompatible, but that each seem to contribute important insights. As is typical in exciting new fields, CSCL research has demonstrated—perhaps above all else—that relatively straight-forward extensions of traditional approaches are inadequate for addressing the intertwining issues raised by CSCL. Researchers in CSCL have come to the field from diverse disciplines and have brought with them disparate methodological traditions. If CSCL wants to become a truly international and multidisciplinary endeavor in the next decade, it needs to develop its own theoretical framework, one appropriate for defining the phenomena and methods of a unique field that transcends academic and cultural boundaries of the past.

At CSCL ’03, I claimed that in situations of collaborative learning, the building of knowledge or the construction of meaning is a group process (Stahl, 2003). It produces artifacts (words, texts, pictures, tools) with group meaning. This meaning should be conceived of at the small-group unit of analysis, even though this shared meaning necessarily involves interpretation and contributions by individuals.

In this paper, I want to push this analysis further and ask, Can collaborative groups think? Answering this question in the affirmative, I want to propose a concept of group cognition (Stahl, in press). A theory of *collaborative learning as group cognition* locates the locus of agency for CSCL in the group, not in the individual, where other theories of learning seek it.

FROM AI TO CSCL

Turing (1950) famously posed the question, “Can machines think?” For the 50 years since then, the field of artificial intelligence (AI) was largely driven by the quest for computer-based (artificial) cognition (intelligence). In recent years, this quest has migrated into the development of technologies that aid or augment human intelligence. As the collaborative technologies of CSCL become more important, the trend may be even more to design computationally-intensive media to support communication among people, making their—human but computer-mediated—group efforts more intelligent.

It has become increasingly clear that computers do not “think” in anything like the way that people do. As has been repeatedly stressed in the past decade or two, human cognition is essentially situated, interpretive, perspectival and largely tacit. Computer symbol processing has none of these characteristics. Computers manipulate information that does not have meaning for the computer, but only for the people who configured or use the computer. Without meaning, there is no need or possibility to reference a situation, interpret symbols, view from a perspective or link to tacit background understanding. It is only the combination of computers with people that think in a meaningful way with the help of computer manipulation of information.

In this paper, I pose a question analogous to the classic AI question: Can groups think? In keeping with the priorities of CSCL, I am interested in the potential of small groups that are collaborating effectively with technological mediation. At CSCL ’02 I argued that collaborative knowledge building was a central phenomenon for CSCL (Stahl, 2002b), and at CSCL ’03 I extended the argument by claiming that meaning-making in collaborative contexts took place primarily at the small-group unit of analysis (Stahl, 2003). Perhaps the question of group cognition can help to set an agenda for future work in CSCL, much as Turing’s question propelled AI research in the past. Perhaps CSCL can provide a positive answer to the question, taking advantage of what AI learned in the process of arriving at its negative conclusion. After all, many technological pursuits within CSCL have been inspired by AI. In the following, we consider three important efforts to determine if computers can think, and apply their considerations to the question of whether small groups of people collaborating together can, under propitious conditions, be said to be thinking as a group. First, let us address a primary stumbling block to thinking about groups as thinking agents.

A GROUP DOES NOT HAVE A BRAIN

The common sense objection to attributing thought to small groups of people is that groups do not have something like a “group brain” the way that individual people have brains. It is assumed that cognition requires some sort of brain—as a substrate for the thinking and as an archive for the thoughts.

Thought as software. The idea of a substrate for thinking was developed in its extreme form in AI. Here, the analogy was that computer hardware was like a human brain in the sense that software runs on it the way that thinking takes place in the brain. Software and its manipulation of information was conceptualized as computations on data. Projecting this model back on psychology, the human mind was then viewed in terms of computations in the brain. Originally, this computation was assumed to be symbol manipulation (Newell & Simon, 1963), but it was later generalized to include the computation of connection values in parallel distributed processes of neural network models (Rumelhart & McClelland, 1986).

Thought as content. Thought has also traditionally been considered some kind of mental content or idea-objects (facts, propositions, beliefs) that exist in the heads of individual humans. For instance, in educational theory the application of this view to learning has been critically characterized as the pouring of content by teachers into the container heads of students (Freire, 1970). Again, this has its analogy in the computer model. Ideas are stored in heads like data is stored in computer memory. According to this model, the mind consists of a database filled with the ideas or facts that a person has learned. Such a view assumes that knowledge is a body of explicit facts. Such facts can be transferred unproblematically from one storage container to another along simple conduits of communication. This view raises apparent problems for the concept of group cognition. For instance, it is often asked when the notion of group learning is proposed, what happens to the group learning when the members of the group separate. To the extent that group members have internalized some of the group learning as individual learning, then this is preserved in the individuals’ respective heads. But the group learning as such has no head to preserve it.

Groupware as group memory. One tact to take in conceptualizing group cognition would be to argue that groupware can serve as a substrate and archival repository for group thought and ideas. Then, one could say that a small group along with its appropriate groupware, as an integrated system, can think.

Discourse as cognition. The view that will be proposed here is somewhat different, although related. We will view *discourse* as providing a substrate for group cognition. The role of groupware is a secondary one of mediating the discourse – providing a conduit that is by no means a simple transfer mechanism. Discourse consists of material things observable in the physical world, like spoken words, inscriptions on paper and bodily gestures. The cognitive ability to engage in discourse is not viewed as the possession of a large set of facts or ideas, but as the ability to skillfully use communicative resources. Among the artifacts that groups learn to use as resources are the affordances of groupware and other technologies. The substrate for a group’s skilled performance includes the individual group members, available meaningful artifacts (including groupware and other collaboration tools or media), the situation of the activity structure, the shared culture and the socio-historical context. So, in a sense, the cognitive ability of a group vanishes when the group breaks up, because it is dependent on the interactions among the members. But it is also true that it is not simply identical to the sum of the members’ individual cognitive abilities because (a) the members have different abilities individually and socially—according to Vygotsky’s (1930/1978) notion of the zone of proximal development as the difference between these—and (b) group cognitive ability is responsive to the context, which is interactively achieved in the group discourse (Garfinkel, 1967). Both of these points make sense if one conceives of the abilities of members as primarily capacities to respond to discursive settings and to take advantage of contextual resources, rather than conceiving of intelligence as a store of facts that can be expressed and used in logical inferences. To the extent that members internalize skills that have been developed in collaborative interactions or acquire cognitive artifacts that have been mediated by group activities, the members preserve the group learning and can bring it to bear on future social occasions, although it might not show up on tests administered to the individuals in isolation.

In the following, we want to explore the sense in which we can claim that small groups can think or engage in group cognition. We will successively take up the three major arguments of Turing, Searle and Dreyfus about whether computers can think, applying their considerations to group cognition.

A TURING TEST FOR GROUPS

In a visionary essay that foresaw much of the subsequent field of AI, Turing (1950) considered many of the arguments related to the question of whether machines could think. By machines, he meant digital computers. He was not arguing that the computers that he worked on at the time could think, but that it was possible to imagine computers that could think. He operationalized the determination of whether something is thinking by assessing whether it could respond to questions in a way that was indistinguishable from how a thinking person might respond. He spelled out this test in terms of an imitation game and predicted that an actual computer could win this game by the year 2000.

The original imitation game is played with three people: a man and a woman, who respond to questions, and an interrogator who cannot see the other two but can pose questions to them and receive their responses. The object of the game is for the interrogator to determine which of the responders is the woman, while the man tries to fool the interrogator and the woman answers honestly.

Turing transposed this game into a test for the question of whether computers can think, subsequently called the Turing Test:

> I believe that in about 50 years' time it will be possible to programme computers, with a storage capacity of about 10^9, to make them play the imitation game so well that an average interrogator will not have more than 70 per cent. chance of making the right identification after five minutes of questioning. (p. 442)

The test reduces the question of whether a computer can think to the question of whether a (properly programmed) computer could produce responses to a human interrogator's probing questions that could not be distinguished from the responses of a (thinking) human.

It is generally accepted that no computer passed the Turing test by the year 2000. Computer programs have been developed that do well on the test if the interrogator's questions are confined to a well-defined domain of subject matter, but not if the questions can be as wide-ranging as Turing's examples. The domain of chess is a good example of a well-defined realm of intelligent behavior. A computer did succeed in beating the best human chess player by around 2000. But interestingly, it did so by using massive numbers of look-ahead computations in a brute-force method, quite the opposite of how human masters play.

Can a group pass the Turing test? Turing argued that his test transformed the ambiguous and ill-defined question about computers thinking into a testable claim that met a variety of objections. His approach has proven to be appealing, although it is not without its critics and although it has not turned out to support his specific prediction. We will now see what we can borrow from the Turing test for the question of whether collaborative groups can think.

Suppose an interrogator communicated questions to a thinking individual person and to a collaborating small group of people. Could the group fool the interrogator into not being able to distinguish to a high probability that the group is not a person? Clearly, a simple strategy would be for the group to elect a spokesperson and let that person respond as an individual. There seems to be no question but that a group can think in the same sense as an individual human according to the Turing test.

In a sense, the Turing test, by operationalizing the phenomenon under consideration puts it in a black box. We can no longer see how thoughts (responses to the interrogator) are being produced. It is reminiscent of the limitation of many quantitative CSCL studies of learning. An operational hypothesis is either confirmed or denied, but the mechanisms of interest are systematically obscured (Stahl, 2002a). We do not really learn much about the nature of thought or learning – whether by individuals, groups or computers – by determining whether their results are indistinguishable or not. One would like to look inside the box.

A CHINESE ROOM FOR GROUPS

Searle's (1980) controversial Chinese room argument takes a look inside the box of an AI computer … and he is disappointed. Writing in an article on "Minds, Brains and Programs," Searle reviews many leading views on whether computers can think, attracts even more views in commentaries, and ends up leaving most readers in more of a quandary than when they started.

Searle's argument revolves around a thought experiment that can actually be traced back to Turing's paper. In describing a model of computers, Turing starts out by saying that a digital computer is "intended to carry out any operations which could be done by a human computer" (Turing, 1950, p. 436). By "human computer" he has in mind a person who follows a book of fixed rules without deviation, doing calculations on an unlimited supply of paper. In a digital computer, the book of rules, paper and human are replaced by software, digital memory and computer processor. Searle reverse-engineers the computer to ask if digital computers think by asking the same question of the "human computer" that Turing imagined. In his thought experiment, Searle imagines that he is the human who follows a book of fixed rules to do computations on paper.

The key move that Searle makes is to note that the computer follows the rules of its software *without* interpreting them. To get a feel of the computer's perspective on this, Searle specifies that the symbols coming into the computer and those going out are all in Chinese. As Searle (who knows no Chinese) sits inside the computer manipulating these symbols according to his book of rules in English, he has no idea what these symbols mean. The software that he executes was cleverly programmed by someone who understood Chinese, so the outputs make Chinese sense as responses to their inputs, even though Searle who is manipulating them inside the computer has no understanding of this sense. From the outside, the computer seems to be behaving intelligently with Chinese symbols. But this is a result of the intelligence of the programmer, not of the human computer (Searle) who is blindly but systematically manipulating the symbols according to the program of his rule book in English.

According to Searle's thought experiment, a computer could, for instance, even pass the Turing test without engaging in any thoughtful understanding whatsoever. Human programmers would have written software based on their understandings, human AI workers would have structured large databases according to their understandings and human interrogators or observers would have interpreted inputs and outputs according to their understandings. The computer would have manipulated bits following strict rules, but with no understanding. The bits might as well be in an unknown foreign language.

Searle's reformulation of the question is whether the instantiation of some AI software could ever, by itself, be a sufficient condition of understanding. He concludes that it could not. He argues that it could not because the computer manipulations have no *intentionality*, that is they do not index any meaning. If a sequence of symbols being processed by the computer is supposed to represent a hamburger in a story about a restaurant, the computer has no understanding that those symbols reference a hamburger, and so the computer cannot be described as intelligently understanding the story. The software programmer and the people interacting with the computer might understand the symbols as representing something meaningful, but the computer does not. Searle distinguishes the perspective of the computer from that of its users, and attributes understanding of the processed information only to the users. He says of machines including digital computers that "they have a level of description at which we can describe them as taking information in at one end, transforming it and producing information as output. But in this case it is up to outside observers to interpret the input and output as information in the ordinary sense" (Searle, 1980, p. 423).

Searle concludes that there is necessarily a material basis for understanding, that no purely formal model like a software program can ever have. He says that he is able to understand English and have other forms of intentionality

> because I am a certain sort of organism with a certain biological (i.e., chemical and physical) structure, and this structure, under certain conditions, is causally capable of producing perception, action, understanding, learning and other intentional phenomena. And part of the point of the present argument is that only something that had those causal powers could have that intentionality. (p. 422)

For Searle, "intentionality" is defined as a feature of mental states such as beliefs or desires, by which they are directed at or are about objects and states of affairs in the world.

Putting Searle into a group. Searle is quite convinced that computers cannot think in the sense proposed by strong AI advocates. Do his arguments apply to groups thinking?

Applying Searle's thought experiment, analysis and conclusions to the question of whether a collaborative group could think is tricky because of the shift of locus of agency from a single physical object to a group of multiple objects, or subjects. What would it mean to remove the individual Searle from his hypothesized computer and to put him into a collaborative group? It would make no sense to put him into a Chinese-speaking group. But we are not asking if every possible group can be said to think, understand or have intentional states. Can it be said of *any* collaborative group that it thinks? So we would put Searle into a group of his English-speaking peers. If the group started to have a successful knowledge-building discourse, we can assume that from Searle's insider position he might well agree that he had an understanding of what was being discussed and also that the group understood the topic.

Would he have to attribute understanding of the topic to the group as a whole or only to its members? If the utterances of the members only made sense as part of the group discourse, or if members of the group only learned by means of the group interactions, then one would be inclined to attribute sense-making and learning to the group unit. This would be the attribution of intentional states to the group in the sense that the group is making sense of something and learning about something—i.e., the group is intending or attending to something.

Another move that Searle considers with his human computer experiment is to have the person who is following the rules in the book and writing on scraps of paper then internalize the book and papers so that the whole system is in the person. In Searle's critique of Turing, this changes nothing of consequence. If we make a similar move with the group, what happens? If one person internalizes the perspectives and utterances of everyone in a collaborative group, that person can play out the group interactions by himself. This is what theoreticians of dialog—e.g., Bakhtin (1986) and Mead (1934/1962)—say happens when we are influenced by others. Vygotsky (1930/1978) sees this process of internalization of social partners and groups as fundamental to individual learning. When one plays out a debate on a topic by oneself, one can certainly be said to be thinking. So why not say that a group that carries out an identical debate, conceivably using the same utterances, is also thinking?

The only issue that still arises is that of agency. One might insist on asking *who* is doing the thinking, and be looking for a unitary physical agent. The group itself could be spread around the world, interacting asynchronously through email. Perhaps a collaboration takes place over time such that at no one time are all the members simultaneously involved. Where is the biological basis for intentionality, with its causal powers that Searle claims as a necessary condition for intentionality, understanding and thought? Certainly, one would say that thought went into formulating the individual emails. That can be explained as the result of an individual's biology, causality, intentionality, understanding, etc. But, in addition, the larger email interchange can be a process of shared meaning-making, where the meaning is understood by the group itself. Comments in a given email may only make sense in relation to other emails by other members.

The group may rely on the eyes of individuals to see things in the physical world and it may rely on the arms of individuals to move things around in the physical world, because the group as a whole has no eyes or arms other than those of its members. But the group itself can make group meaning through its own group discourse. The interplay of words and gestures, their inferences and implications, their connotations and references, their indexing of their situation and their mediating of available artifacts can take place at the group unit of analysis. These actions may not be attributable to any individual unit—or at least may be more simply understood at the group level.

BEING-IN-THE-WORLD AS GROUPS

The third "critique of artificial reason" that we want to consider is that of Dreyfus (1972; 1986; 1991). Dreyfus agrees with Searle that AI has emerged from the attempt to push a specific philosophic position too far, to the detriment and confusion of AI. Dreyfus calls this extreme position "representationalism" and argues that it ignores much of what accounts for human understanding. It in effect reduces our complex engagement in the world, our sophisticated social know-how and our subtle sense of what is going on around our embodied presence to a large database of symbols and books of explicit rules:

> Rationalists such as Descartes and Leibniz thought of the mind as defined by its capacity to form representations of all domains of activity. These representations were taken to be theories of the domains in question, the idea being that representing the fixed, context-free features of a domain and the principles governing their interaction explains the domain's intelligibility … mirrored in the mind in propositional form. (Dreyfus, 1992, p. xvii)

Representationalism reduces all knowing, meaning, understanding, cognition, intelligence to the possession of sets of facts, ideas or propositions. It matters little whether these explicit formulations of knowledge are said to exist in an ideal world of non-material forms (Plato), as purely mental thoughts (Descartes), as linguistic propositions (early Wittgenstein) or stored in database entries (AI). Wittgenstein's early *Tractatus*, which reduces philosophy to a set of numbered propositions, begins by defining the world as "the totality of facts, not of things" (Wittgenstein, 1921/1974, § 1.1). From here, via the work of the logical positivists, it is easy to conceive of capturing human knowledge in a database of explicit representations of facts—such as Searle imagined in his books of programmed instructions for manipulating Chinese symbols.

The problem with representationalism, according to Dreyfus, is that it ignores the diverse ways in which people know. The consequence that Dreyfus draws for AI is that it cannot succeed in its goal of reproducing intelligence using just formal representations of knowledge. Dreyfus highlights three problems that arose for AI in pursuing this approach: (1) sensible retrieval, (2) representation of skills and (3) identification of relevance.

Retrieval. The AI approach has proven unable to structure its knowledge-bases in a way that supports the drawing of commonsensical inferences from them. For instance, as people learn more about a topic, they are able to infer other things about that topic faster and easier, but as a computer stores more facts on a topic its retrieval and inference algorithms slow down dramatically.

Dreyfus details his critique by focusing on a large AI effort to capture people's everyday background knowledge and to retrieve relevant facts needed for making common sense inferences. Dreyfus argues that the logic of this approach is precisely backward from the way people's minds work:

> The conviction that people *are* storing context-free facts and using meta-rules to cut down the search space is precisely the dubious rationalist assumption in question. It must be tested by looking at the phenomenology of everyday know-how. Such an account is worked out by Heidegger and his followers such as Merleau-Ponty and the anthropologist Pierre Bourdieu. They find that what counts as the facts depends on our everyday skills. (Dreyfus, 1992, p. xxii)

Skills. AI representations cannot capture the forms of knowledge that consist in skills, know-how and expertise. People know how to do many things—like ride a bike, enjoy a poem or respond to a chess position—that they are unable to state or explain in sentences and rules. The effort within AI to program expert systems, for instance, largely failed because it proved impossible to solicit the knowledge of domain experts. An important form of this issue is that human understanding relies heavily upon a vast background knowledge that allows people to make sense of propositional knowledge. This background knowledge builds upon our extensive life experience, which is not reducible to sets of stored facts.

> Human beings who have had vast experience in the natural and social world have a direct sense of how things are done and what to expect. Our global familiarity thus enables us to respond to what is relevant and ignore what is irrelevant without planning based on purpose-free representations of context-free facts. (p. xxix)

Relevance. A fundamental interpretive skill of people is knowing what is relevant within a given situation and perspective. This sense of relevance cannot be programmed into a computer using explicit rules. This ability to focus on what is relevant is related to people's skill in drawing inferences and builds on their expert background knowledge.

> The point is that a manager's expertise, and expertise in general, consists in being able to respond to the relevant facts. A computer can help by supplying more facts than the manager could possibly remember, but only experience enables the manager to see the current state of affairs as a specific situation and to see what is relevant. That expert know-how cannot be put into the computer by adding more facts, since the issue is which is the current correct perspective from which to determine which facts are relevant. (p. xlii)

Dreyfus emphasizes that *facts* are not what is immediately given in human experience and understanding. Rather, what is to count as a fact is itself mediated by our skills, our situation in the world and our perspective as embodied and engaged.

Dreyfus' critique shows that computers cannot think in the most important ways that people do. Arguing on the basis of a Heideggerian analysis of human being-in-the-world as situated, engaged, perspectival, skilled and involved with meaningful artifacts, Dreyfus provides the basis for understanding the failure of computers to pass the Turing test and to exhibit the kind of intentionality that Searle argues is a necessary condition of cognition. Explicit, propositional, factual knowledge is not an adequate starting point for analyzing or duplicating human cognition. There are a number of factors that come first analytically and experientially: tacit know-how, practical skills, social practices, cultural habits, embodied orientation, engaged perspective, involvement with artifacts, social interaction, perception of meaningfulness and directedness toward things in the world. Heidegger's (1927/1996) analysis of human existence, for instance, begins with our being involved in the world within situational networks of significant artifacts. Our relationship to things as objects of explicit propositions and our expression of factual propositions are much later, secondary products of mediations built on top of the more primordial phenomena. Similarly, Merleau-Ponty (1945/2002) stresses our orientation within a meaningful social and physical space structured around our sense of being embodied. Because AI representations lack the features that are primary in human cognition and try to reduce everything to a secondary phenomenon of factual propositions, they ultimately fail to be able to either imitate human cognition to the degree envisioned by Turing or to capture the sense of understanding sought by Searle.

Being-with-others in groups. We now turn to the question of whether the proposed notion of group cognition fares any better against these standards than did the AI notion of computer cognition.

Clearly, the individual members of a group bring with them the skills, background and intentionality to allow a group to determine what are the relevant facts and issues. But in what sense does the group as a whole have or share these? We do not define the group as a physical collection of the members' bodies. The group might exist in an online, virtual form, physically distributed across arbitrary spatial and temporal distances. Rather, the group exists as a discourse, perhaps recorded in a video, chat log or transcript. This group discourse can reflect such tacit skills, commonsense background knowledge and intentionality.

Group discourse is engaged in a group activity, embedded within a context of tacitly understood goals and situated in a network of meaningful artifacts. The discourse itself exhibits intentionality. It builds upon tacit background knowledge of the experiential world. It adopts—sometimes through involved group processes of negotiation and enactment—perspectives that determine relevance. So groups can think in much the same situated, engaged way that individuals do.

GROUP DISCOURSE AS EMERGENT THINKING

This paper has argued that small collaborative groups—at least on occasion and under properly conducive conditions—can think. It is not only possible, but also quite reasonable to speak of groups as engaging in human cognition in a sense that is not appropriate for applying to computer computations, even in AI simulations of intelligent behavior. When we talk of groups thinking, we are referring not so much to the physical assemblage of people as to the group discourse in which they engage.

To some social scientists, such as Vygotsky, the group level (which he calls social or inter-subjective) is actually prior in conceptual and developmental importance to the individual (intra-subjective) level. So why does the notion of group cognition strike many people as counter-intuitive? When it is recognized, it is generally trivialized as some kind of mysterious "synergy." Often, people focus on the dangers identified by social psychologists as "group think"—where group obedience overrides individual rationality. At best, the commonsensical attitude acknowledges that "two heads are better than one." This standard expression suggests part of the problem: thought is conceived as something that takes place inside of individual heads, so that group cognition is conceived as a sum of facts from individual heads, rather than as a positive cognitive phenomenon of its own.

An alternative conceptualization is to view group cognition as an *emergent* quality of the interaction of individual cognitive processes. The emergence of group cognition is different from other forms of emergence. Conversation is the interaction of utterances, gestures, etc. from a small number of people. The interaction can be extremely complex. It involves the ways in which subsequent utterances respond to previous ones and anticipate or solicit future ones. Individual terms carry with them extensive histories of connotations and implications. Features of the situation and of its constituent artifacts are indexed in manifold ways. Syntactic structures weave together meanings and implications. Effective interpretations are active at many levels, constructing an accounting of the conversation itself even as it enacts its locutionary, perlocutionary and illocutionary force (Searle, 1969).

Yes, small groups can think. Their group cognition emerges in their group discourse. This is a unique form of emergence. It differs from statistical, simple-rule-governed and social emergence. It is driven by linguistic mechanisms. Understanding group cognition will require a new science with methods that differ from the traditions of AI, psychology and educational research—methods based on the interactional subtleties of conversational discourse rather than on statistical regularities.

GROUP COGNITION AND CSCL

Many methodologies popular in CSCL research focus on the individual as the unit of analysis and locus of agency: what the individual student does or says or learns. Even from the perspective of an interest in group cognition and group discourse, such methods can be useful and provide part of the analysis, because group thinking and activity is intimately intertwined with that of the individual members of the group. However, it is also important and insightful to view collaborative activities as linguistic, cognitive and interactional processes at the group level of description. This involves taking the group as the unit of analysis and as the focal agent. One can then analyze how a group solves a problem through the interplay of utterances proposing, challenging, questioning, correcting, negotiating and confirming emergent group meaning. One can see how a group does things with words that have the force of accomplishing changes in the shared social world. Some things, like electing an official, can only be done by groups—although this obviously involves individuals. Other things, like solving a challenging problem, may be done better by groups than by individuals—although the different perspectives and considerations are contributed by individuals.

CSCL is distinguished as a field of inquiry by its focus on group collaboration in learning; it makes sense to orient the methods of the field to thinking at the small-group unit of analysis. This may require re-thinking—as a research community—our theoretical framework, such as our conceptualization of "cognition" that we have inherited from the representationalism of cognitive sciences and learning sciences oriented overwhelmingly toward the individual.

Group interactions may be characterized as "cognitive" because they display the requisite characteristics of sequentiality, accountability and sense making—not because they are extensions of individual cognition. Group cognition is a phenomenon at the small-group unit of analysis, not a derivative of either individual thinking or community-level establishment of cultural resources. It is the source of knowledge constructed collaboratively—and is therefore an appropriate foundation for CSCL.

Individual learning enters the picture secondarily. Because collaboration requires shared understanding, processes of group cognition generally ensure that all participants keep pace with the group, to the extent needed for the group discourse's practical purposes. This causes individuals to develop and alter their interpretations of constructed meaning and perhaps internalize cognitive artifacts based on the products of group cognition, such as meaningful texts.

The exploration of empirical case studies of small-group knowledge-building discourse are needed to help to describe in both concrete and theoretical ways how group cognition is accomplished as a linguistic achievement. Rigorous conversational analysis of multiple studies will lead to an improved understanding of the methods that participants use to constitute and structure group interaction and to engage in collaborative problem solving.

THE NEXT DECADE

The Internet offers the potential to join individual minds together effectively across time and space, thereby overcoming the limitations of individual cognition. Networked computers not only allow global access to information, but could also facilitate collaborative knowledge building within online communities. However, numerous case studies in CSCL have found that even in virtual environments intentionally designed to support knowledge building, discussions are generally limited at best to the sharing of personal opinions. Commercial systems provide media for generic communication or transmission of information, but no specific support for the phases of more involved collaboration. Driven by the market-place demands of corporate users and educational institutions, the designs of these systems aim to structure and control individual access and usage rather than to scaffold group cognition.

We need to better conceptualize collaborative knowledge building as a set of group processes. This will lead to the analysis of group cognition as a phenomenon of small-group discourse. Contributions to collaborative knowledge-building discussions do not typically express meanings that already existed in mental representations of individual participants. The utterances are indexical, elliptical and projective in the sense that they contribute to meaning at the group unit of analysis by virtue of their embeddedness in the group situation, discourse and activity. The meaning and the knowledge are originally constructed through group cognition. Individual cognition may later result from internalization or retrospective accounts. Accordingly, evidence of collaborative learning is to be found in the brief episodes of shared meaning making in which group knowing is constituted, rather than in traces of lasting capabilities of individuals, which are subject to numerous psychological factors.

In particular, conversation analysis (Psathas, 1995; Sacks, 1992) can serve as a methodology for making group cognition visible. Methodologically rigorous interpretations (Koschmann, Stahl, & Zemel, 2005) can analyze intersubjective interactions like turn-taking, knowledge negotiation, adjacency pairs and conversational repair. Through such analysis, we can see that the basic components of collaborative knowledge building are not actions of individuals, but are methods of small-group activity. Through them, shared meanings are proposed, adopted and refined. The processes of group cognition incorporate contributions by individuals, based on individual interpretations of the emerging and evolving group meanings. But these individual utterances are essentially fragmentary; they only become meaningful by virtue of their contributing to the group context. That is why computer support for collaborative knowledge building must be centrally concerned with group cognition.

The cycle of software prototyping, conversation analysis and theoretical reflection must be iterated repeatedly. Many innovations of CSCL systems will have to be developed and tried out, building a whole field of technology for use in supporting specific group methods of collaboration. The interactions that take place online in these and other contexts must be analyzed systematically, in order to catalog methods that people use to accomplish their group work, learning, communicating and thinking. The technology and the analyses should be conceptualized within a vocabulary adequate for making sense of them. A theory of group cognition may provide a starting point for this.

The comprehension of how thinking takes place at the small-group locus of agency will guide the design of more effective computer support for collaborative knowledge building. Then the potential of group cognition can blossom around the world. This will require a global effort, itself a major instance of group cognition. This defines the task of CSCL in the next decade.

REFERENCES

Bakhtin, M. (1986). *Speech genres and other late essays* (V. McGee, Trans.). Austin, TX: University of Texas Press.

Dreyfus, H. (1972). *What computers cannot do*. New York, NY: Harper and Row.

Dreyfus, H., & Dreyfus, S. (1986). *Mind over machine: The power of human intuition and expertise in the era of the computer*. New York, NY: Free Press.

Dreyfus, H. (1991). *Being-in-the-World: A commentary on Heidegger's Being and time, division I*. Cambridge, MA: MIT Press.

Dreyfus, H. (1992). *What computers still can't do: A critique of artificial reason*. Cambridge, MA: MIT Press.
Freire, P. (1970). *Pedagogy of the oppressed*. New York, NY: Continuum.
Garfinkel, H. (1967). *Studies in ethnomethodology*. Englewood Cliffs, NJ: Prentice-Hall.
Heidegger, M. (1927/1996). *Being and time: A translation of Sein und Zeit* (J. Stambaugh, Trans.). Albany, NY: SUNY Press.
Koschmann, T., Stahl, G., & Zemel, A. (2005). The video analyst's manifesto (or the implications of Garfinkel's policies for the development of a program of video analytic research within the learning sciences). In R. Goldman, R. Pea, B. Barron & S. Derry (Eds.), *Video research in the learning sciences*. Retrieved from http://www.cis.drexel.edu/faculty/gerry/publications/journals/manifesto.pdf.
Mead, G. H. (1934/1962). *Mind, self and society*. Chicago, IL: University of Chicago Press.
Merleau-Ponty, M. (1945/2002). *The phenomenology of perception* (C. Smith, Trans. 2 ed.). New York, NY: Routledge.
Newell, A., & Simon, H. A. (1963). GPS, a program that simulates human thought. In A. Feigenbaum & V. Feldman (Eds.), *Computers and thought* (pp. 279-293). New York, NY: McGraw Hill.
Psathas, G. (1995). *Conversation analysis: The study of talk-in-interaction*. Thousand Oaks, CA: Sage.
Rumelhart, D. A., & McClelland, J. L. (1986). *Parallel distributed processing: Explorations in the microstructure of cognition. Volumes 1 & 2*. Cambridge, MA: MIT Press.
Sacks, H. (1992). *Lectures on conversation*. Oxford, UK: Blackwell.
Searle, J. (1969). *Speech acts: An essay in the philosophy of language*. Cambridge, UK: Cambridge University Press.
Searle, J. (1980). Minds, brains and programs. *Behavioral and Brain Sciences, 3*, 417-424.
Stahl, G. (2002a). Rediscovering CSCL. In T. Koschmann, R. Hall & N. Miyake (Eds.), *CSCL 2: Carrying forward the conversation* (pp. 169-181). Hillsdale, NJ: Lawrence Erlbaum Associates. Retrieved from http://www.cis.drexel.edu/faculty/gerry/cscl/papers/ch01.pdf.
Stahl, G. (2002b). *Contributions to a theoretical framework for CSCL*. Paper presented at the International Conference on Computer Supported Collaborative Learning (CSCL '02), Boulder, CO. Proceedings pp. 62-71. Retrieved from http://www.cis.drexel.edu/faculty/gerry/cscl/papers/ch15.pdf.
Stahl, G. (2003). Meaning and interpretation in collaboration. In B. Wasson, S. Ludvigsen & U. Hoppe (Eds.), *Designing for change in networked learning environments: Proceedings of the international conference on computer support for collaborative learning (CSCL '03)* (pp. 523-532). Bergen, Norway: Kluwer Publishers. Retrieved from http://www.cis.drexel.edu/faculty/gerry/cscl/papers/ch20.pdf.
Stahl, G. (in press). *Group cognition: Computer support for collaborative knowledge building*. Cambridge, MA: MIT Press. Retrieved from http://www.cis.drexel.edu/faculty/gerry/mit/.
Turing, A. M. (1950). Computing machinery and intelligence. *Mind, 59*, 433-460.
Vygotsky, L. (1930/1978). *Mind in society*. Cambridge, MA: Harvard University Press.
Wittgenstein, L. (1921/1974). *Tractatus logico philosophicus*. London, UK: Routledge.