
Can Collaborative Groups Think?

Can the thinking in collaborative groups really be attributed to the group as a whole? What forms of cognition can small groups engage in? What is the relationship of group cognition to the cognition of the individuals who are in the group? The question suggests a widening of the notion of cognition from one conventionally based on the individual human. It is hard for many people to accept this rethinking of thinking.

A similar attempt to extend the definition of cognition with AI eventually concluded that computers cannot think because (a) they could not convincingly imitate intelligent human behavior, (b) they could not understand the symbols they manipulated, and (c) they could not act intelligently in the world. Applying these criteria to small groups, however, suggests that groups can think as entities distinct from their individual human members. Many people became accustomed to considering computers as potential cognitive agents during the half century in which the claims of AI were debated. Perhaps this receptivity to a broader definition of cognition can now be transferred to human group cognition, even as it has been denied to computers.

It is important to look at cognition by the group as a unit of analysis, both because individual thought is derived in crucial ways from group cognition and because the meaning created by groups often only makes sense within the group discourse context as the analytic unit. The group discourse is thus seen as an important source of shared meaning and as the locus of collaborative cognitive processes. This provides an evocative framework for conceptualizing CSCW and CSCL.

From AI to CSCL

Turing (1950) famously posed the question, “Can machines think?” For 50 years after that, the field of artificial intelligence (AI) was largely driven by Turing’s framing of 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 CSCW and 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 (see part I of this book).

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 people (who understand meaning) with

computers (that help manipulate information) that it can be said that computers are involved in thinking.

In this chapter, I pose a question analogous to the classic AI question: can groups think? In keeping with the priorities of CSCW and CSCL, I am interested in the potential of small groups that are collaborating effectively with technological mediation. Chapter 15 argued that collaborative knowledge building was a central phenomenon for collaboration, and chapter 16 extended the argument by claiming that meaning making in collaborative contexts took place primarily at the small-group unit of analysis. Perhaps the question of group cognition can help to set an agenda for future work in computer-supported collaboration, much as Turing's question propelled AI research in the past. CSCW and CSCL may 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 CSCW and CSCL have been inspired by AI.

In the following, I want to explore the sense in which small groups of people collaborating together can, under propitious conditions, be said to be thinking as a group or engaging in group cognition. I start with the simpler issue of whether small groups can *learn*, drawing on a study by Cohen and colleagues. Then, we take up the three major arguments by Turing, Searle and Dreyfus about whether computers can think, applying their considerations to group cognition.

Can Groups Learn?

Learning is considered closely related to thinking. Perhaps a first step in addressing the question of group cognition would be to ask if groups can learn. This seems like a concrete question capable of being operationalized and explored empirically.

We saw in chapter 13 that the group of students working with *SimRocket* learned something as a whole—the group learned to see the list structure as a “paired configuration.” Utterances from individuals indicated that they did not see the list this way before the moment of collaboration, but afterward they were able to appreciate this structure (Steven at 1:24:46) and to use it to compare rockets (Brent, Jamie, Chuck at 1:26:46).

Educational researchers often prefer to argue for hypotheses based on statistically significant differences between experimentally controlled conditions, applying results of pre- and post-tests. Where discourse is brought in as data, utterances are coded and the number of utterances in certain categories are quantitatively compared across conditions. An interesting analysis using such a methodology was recently conducted about whether groups can learn by a team of educational researchers at Stanford. We now turn to their findings.

An Empirical Test

Their research has been reported under the title, “Can Groups Learn?” (Cohen *et al.*, 2002). They designed an experiment to assess the role of small groups as learners.

The authors note that social scientists who have studied groups often claim that a group is greater than the sum of its individual parts. However, they continue,

There has been a tendency in assessment to regard the potential for performance in a group as the sum total of the amount of information, skills and abilities that *individuals* bring to that group. Through the creative exchange of ideas, groups can solve problems and construct knowledge beyond the capacity of any single member. Thus it is possible to talk about the concept of *group learning* that is a result of the interaction of the group members and is not attributable to one well-informed person who undertakes to create the product or even to a division of labor in which different persons contribute different pieces of the product. (p. 1046)

Here the authors have identified the problem to be one of assessment at the group unit of analysis, which is distinct from treating the group learning as measurable by the sum of individual member learning. They suggest that groups can build knowledge as a group because of the potential of group discourse. The authors define the concept of group learning as something that is not attributable to one person or to a division of labor resulting in multiple individual products. (Unfortunately the authors characterize the building of group knowledge at the individual unit of analysis, as “the creative exchange of ideas,” where ideas presumably come from the individuals and it is ambiguous whether the creativity is conceived as that of the group discourse or that of the individual participants.)

A carefully planned experiment was conducted with 39 small groups of four or five students. The groups were located in five different sixth-grade classrooms. The 163 students were from a linguistically, ethnically and racially diverse student body in California, including children of immigrant workers. All five of the teachers were highly trained in instructional strategies of complex instruction and the classes were tested to confirm that the teachers and students were all equally proficient at engaging in group work. The class work involved a week-long focal unit on the Egyptian afterlife, including group discussions, rehearsed group skits and individual written reports.

The experimental design was an ingenious attempt to distinguish individual and group phases of learning and to assess them separately in order to see their relationship. The controlled variable was that only three of the five teachers gave explicit instructions to the students in the groups. For instance, for a skit about the heart, the evaluation criteria were, “Skit includes at least 2 sins, 2 virtues and 1 spell; Skit gives good reasons for whether or not the deceased entered the afterlife; Skit is well rehearsed and believable.” All classrooms participated in skill-builder exercises designed to improve the general quality of group discussion, but the explicit evaluation criteria were only included in the exercises for three groups—the experimental condition. The unit included a number of group discussions and activities like the skits. At the end, students wrote individual essays, where they were clearly instructed to use what they learned in the group activities.

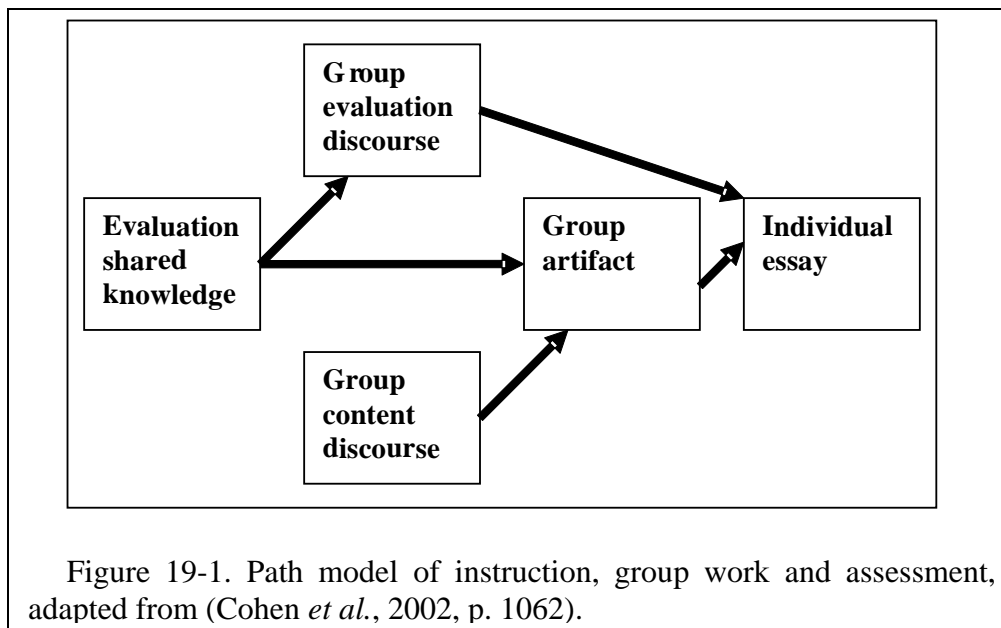
The following measurements were taken:

1. a pre-test of individual knowledge of the subject matter of the Egyptian afterlife,
2. whether or not the group was trained with the explicit evaluation criteria,
3. the percent of group discussion coded as evaluating the group product,
4. the percent of group discussion coded as related to the content of the group product,
5. the percent of group discussion coded as off-topic,
6. the quality of the group product (e.g., a skit), based on the explicit evaluation criteria,
7. the quality of the individual essay (language skills were factored out) and
8. a post-test of the student’s knowledge of the subject matter.

A path model of the causal relations between instructional variables and these assessment measurements was constructed by running several regression analyses.

Interestingly, pre-test scores (1) were not a predictor of the quality of the group products. The better group products (6) were the result of focused group discussion (3 & 4) and shared awareness of evaluation criteria (2), not the result of superior individual knowledge brought to the group. The same proved true for the individual essays (7)! Prior individual knowledge was not a predictor of the quality of the individual essay. The better individual essays (averaged for each group) were the result of the group's discourse on evaluation considerations of the group product and the quality of the group product itself.

Perhaps equally surprisingly, the experimental condition of providing training on the explicit evaluation criteria (2) had no *direct* effect on the individual cognitive performance (7). The individual effect was mediated by the group work! "Evaluation criteria had no effect on essay score. It is through the increase in self-assessment (talk that is evaluation of product) and through the superior product that evaluation criteria affect the final essay" (p. 1062). Figure 19-1 represents all of the statistically significant causal relationships with arrows. Individual knowledge as measured by pre- and post-tests was not a significant determinant, even of the individual performance, let alone of the group performance. Furthermore, the difference between the control and the



experimental difference only exerts an effect by means of its effect on the group performance. In the terms of the experiment's paradigm, this demonstrates an empirically quantifiable phenomenon of group learning as distinct from the sum of individual member learning.

Figure 19-1 goes approximately here

Curiously, the authors turn the conclusion around at the end of their paper to argue that individual assessments can be used for assessment at the group level, as though

teachers and others are primarily concerned about group performance: “The fact that individual performance in the last analysis was affected by both quality of group product and self-assessment shows that teachers can feel confident about using individual assessment to measure the instructional outcome of group work. Individuals greatly benefit by exposure to the discourse and the creative process in groups” (p. 1066). At least this addresses the more usual concern that people have for group learning benefiting the individual students.

How Do Groups Learn?

Given its title and its opening remarks about group learning being assumed by social scientists but rarely assessed, the research paper based on the study described above seems to be positioning itself as a study of group learning. The authors, in a series of three central hypotheses of the experiment, speculate about the mechanisms of group learning. For instance, in motivating their first hypothesis, the authors suggest: the use of explicit criteria and feedback from the whole class should “prepare groups to produce better group products”; it should “improve the quality of the discussion and thus promote group learning”; it should “make the group more task focused” (p. 1048). In connection with hypothesis II, they note that certain conditions are ordinarily necessary for group work to have a positive effect on individual learning, but that when such conditions are met, for instance, “when there is a true group task that cannot be well done by one individual, the process of creating a group product will add to the understanding and ability to articulate knowledge on the part of all group members” (p. 1049). Finally, hypothesis III is, “The better the quality of the group discussion and product, the better will be the individual performance of group members” (p. 1050).

These suggestions concerning group learning and its relationship to individual learning are reasonable and are largely confirmed by the empirical support discovered for the path model shown in figure 1. This confirmation, coupled with the disconfirmation of a significant correlation between pre-test scores and group performance, leads the authors to conclude that learning “came about through reciprocal exchange of ideas and through a willingness to be self-critical about what the group was creating. ... Learning arose from the group as a whole” (p. 1064).

But these specific characterizations of what happened are unjustifiable interpretations based largely on the authors’ assumptions as expressed in their hypotheses. The method they followed obfuscated the discourse mechanisms at work in the group sessions (similar to the studies critiqued in chapter 10). There is no data analysis concerning exchanges of ideas or willingness of students to be self-critical. These are interpretations imposed on the evidence by the researchers, not grounded in their empirical analyses. These are interpretations that do not emerge from a rigorous interpretive methodology such as that of video analysis as proposed in chapter 18.

Statistical correlations can at best indicate that one condition caused another (to a high probability). They are rarely able to show the mechanisms at work. The authors of the study speculate in their hypotheses that certain mechanisms are at work, such as discussion quality, task focus, increased understanding, articulation of knowledge, and product quality. Then, in their conclusions they repeat their assumptions that mechanisms like reciprocal exchange of ideas and willingness to be self-critical are responsible for superior performances. However, their methodology does not allow them to make these

factors visible to us in the sense discussed in chapter 18. They count how many utterances in a group expressed ideas related to the topic and how many expressed opinions critical of content ideas, but the utterances are treated as self-contained units, attributable to individuals. Even the authors' wording betrays this attribution of the utterances to the individual students: "reciprocal exchange" and "willingness." The study concludes that the experimental effects are not attributable to individuals because (a) the study aggregated all the discourse and (b) the aggregated pre-tests did not correlate with the aggregated essays. This is rather indirect evidence. All sorts of relationships could be hidden in all this aggregation.

Another complicating factor is the nature of the experimental condition. The independent variable was training in evaluation criteria. But these evaluation criteria were to be applied to the group work, not to the individual reports. So the fact that the condition impacted the group work directly and that the discovered individual effects were consequently mediated by the group work should be no surprise.

It would be interesting to be able to look at the group interactions carefully and see how the training in evaluation criteria actually played itself out in the different groups. How does the "reciprocal exchange of ideas" take place interactionally—through question and answer pairs or through key terms gradually accruing more meaning? Do some students play leading roles in developing or critiquing ideas or do the ideas emerge and evolve through intense, indexical, elliptical, projecting, mutually completing gestures and utterances (as we saw in chapter 12) that would not be interpretable in isolation? Statistical parameters give us little insight into the nature of group learning and its intimate relationship with individual learning. They can be useful in locating episodes in the data where group learning seems to be talking place, so that qualitative micro-analysis can then be applied to this data to understand it better.

Nevertheless, this paper does define a sense in which group learning can be defined, operationalized and quantified. It concludes that there is an important phenomenon of group learning in this sense, and demonstrates its presence in a real school setting. Having reviewed this empirical argument that groups can learn, we now ask if groups can think. To explore this more philosophical question, we investigate whether we can adopt the arguments from AI concerning whether computers can think. We address the main reaction against the idea of group cognition and then turn to the arguments of Turing, Searle and Dreyfus about what it means to think.

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. However, this argument is not entirely satisfactory.

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. (It may be considered ironic that Turing's most famous proposal is based on deceptions about gender, considering the circumstances of the tragic end of his life.)

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 percent 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.

Turing, who was largely responsible for working out the foundations of computation theory, specified what he meant by a computer in terms of a “discrete state machine.” This is a theoretical machine that is always in one of a number of well-defined states and that moves from one to another of these states based on a state-change table that specifies a new state given any legal input signal and the current state.

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 we saw in chapter 10 of many quantitative CSCL studies of learning. An operational hypothesis is either confirmed or denied, but the mechanisms of interest are systematically obscured. 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 computer as a discrete state machine, 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 an executive, store and control—or, in modern terms, software, digital memory and computer processor. Searle reverse-engineers the computer to ask if digital computers consisting of software, memory and processors think by asking the same question of the “human computer” that Turing imagined being asked of the digital computer. 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 argumentative 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 (written in English, of course), 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.

According to Searle’s “thought experiment” (note that from the start Searle modestly characterizes his own behavior as thinking) 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, which 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 unit of analysis 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. Such a group would not meet the hermeneutic precondition of shared background knowledge and would not be a collaborative success. 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 (1986a) 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 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.

Searle, who wrote the ground-breaking text on speech acts (Searle, 1969), has overlooked in his discussion of thinking the power of language itself to be the agent of thought. This may not affect his critique of AI (for in outputting words, computers do not engage in intentional speech acts, except in the eyes of others), but it is crucial for our question of group thinking. For when we say that a group thinks, we are not postulating the group as a unitary physical object but are focusing on the unity of the group's discourse: the fact that effective collaborative discourse is best understood at the level of the group interaction rather than by focusing on the contributions of individual members. The group discourse has a coherence, and the references of the words within it are densely, inextricably interwoven. Furthermore, the group can act by means of its speech acts.

Although Searle sounds like he is making a materialist argument for biological structures and causal properties that do not map directly to collaborative groups, his discussion is primarily one about language. It has more to do with the nature of Chinese and programming languages than it does with hamburgers and neurons. He is basically arguing that computers do not understand their programming languages the way people understand their mother languages. The difference has more to do with the languages than with the computers or people. Even if a human computer executes a "story understanding" AI program in a software language, there will be no understanding of the story, there will only be a lot of rote following of meaningless rules.

Searle's case hinges on the argument that knowing a lot of rules about something is not equivalent to understanding it. For instance, in manipulating rules related to stories about visits to restaurants, the rules about symbols for menus, ordering, food, eating, paying, tipping etc. do not make for an understanding of restaurant stories because the

computer does not understand that the symbol “hamburger” represents a hamburger—the symbol manipulation lacks intentionality. Let us consider this argument further.

Why would someone ever have thought that having large sets of rules constitutes understanding? Perhaps they thought this because we often learn by being given explicit rules. For instance, I learned German by memorizing rules about word order, endings, spellings, uses and relationships. So why wouldn't Searle understand Chinese after internalizing all the rules? Searle might respond that German was my second language and that I learned it by relating it to my mother tongue, but that I learned English the way that Vygotsky's infant learned the pointing gesture—by interacting with the world and other people, with intentionality (Vygotsky, 1930/1978).

What about more abstract understanding than that of restaurants? Don't we learn chess by learning rules for legal moves, strategies and common positions? Surely being able to see and move the little wooden pieces is not of the essence. We can play chess blindfolded or online, and computers can play chess better than we can. It is, indeed, interesting that computers do not display the kind of expert understanding—or “professional vision” (Goodwin & Goodwin, 1994)—that the best chess players do (Dreyfus & Dreyfus, 1986). But that does not mean (a) that they have no understanding or (b) that understanding does not come through internalizing rules and therefore might not come to Searle inside his Chinese room.

Consider students learning mathematics, for example basic algebra. Algebra is not about hamburgers or apples and oranges. It is about symbols and rules for manipulating the symbols. In learning algebra, students learn algorithmic procedures for manipulating mathematical symbols. They learn to stick to the fixed rules rigidly and to carry out the manipulations quickly. Once they know a book full of rules and can carry out the manipulations strictly according to the rules, we say that they have learned algebra, that they know algebra, that they understand algebra and that they can think algebraically.

Some teachers of mathematics in recent years might say that memorization of procedural rules is not enough. Students need to be able to talk about the math. They should be able to demonstrate a “deep understanding” of the math. But what does a deep understanding consist of here? Well, they might say, the student should be able to explain how she arrived at her answer. Perhaps she should be able to solve a given problem by a number of alternative methods, thereby exploring the nature of the problem. But isn't this just a matter of internalizing more rules and being able to state them? The given problem can be solved by applying various sets of rules and manipulations, and one can express these rules in knowledgeable-sounding utterances. Perhaps knowing how to select the right rules to solve a given problem is a sign of mathematical understanding. But software often includes rules for making such decisions. In fact, problems in logic and mathematics can be solved by computer programs quite a bit better than by ninth graders. What is it that these programs do not understand that the ninth graders do?

Perhaps the answer to this question will have to wait for the results of future empirical studies of collaborative discourse involved in math problem solving. Rather than speculating on this matter, we should look closely using the methods of video analysis or conversation analysis to see just what goes on in the discourse of groups who display a deep understanding of the mathematics they are collaborating on and the discourse of groups who display patterns of manipulating mathematical symbols with little understanding. Such an approach can get behind the comparison of outputs to inputs

(e.g., an algorithmic solution to a given math problem) to make visible the reasoning that goes on within the problem-solving group. In this way, the thinking of groups would provide a window on how individuals think.

An investigation of the thoughtful understanding and the meaning making that takes place in the events simulated in AI programs or quantified in educational experiments—but lost through the behavioristic or operationalizing procedures of simulating or quantifying what takes place—might get at the nature of collaborative thought and human deep understanding. In the end, Searle recommends that AI purge itself of the approach already established by Turing of ignoring the phenomena that are not immediately observable by certain experimental methods:

The Turing test is typical of the tradition in being unashamedly behavioristic and operationalistic, and I believe that if AI workers totally repudiated behaviorism and operationalism much of the confusion between simulation and duplication would be eliminated. (Searle, 1980, p. 423)

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 and 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 a knowledge-base in a way that supports the drawing of commonsensical inferences from it. 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 Lenat's Cyc project, 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 (retrieval) and builds on their expert background knowledge (skills).

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)

In all three points, 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. So we need to ask whether such a group discourse reflects such tacit skills, commonsense background knowledge and intentionality.

Recall a key utterance from the group discourse in chapter 12:

1:22:05	Brent	This one's different
---------	-------	----------------------

This utterance reveals intentionality. The deictic phrase, "this one," indexes some part of the simulation list artifact. The attribute, "different," which the utterance associates

with its subject, connotes background knowledge. The attribution of difference is necessarily from a specific perspective. Any two things can be considered different from some perspective of relevance (Rittel & Webber, 1973). To make this utterance is to assume a particular perspective and to assume that it is part of the group perspective. The fact that others did not agree with the utterance at first signals that this perspective had not yet been established as a shared group perspective. It precipitates an intense moment of collaboration in which the students repair the breakdown of the group perspective and establish the perspective proposed by this utterance through group negotiation and clarification. A close conversation analysis shows how subtle this particular perspective was and how the group had to go through a complex learning process in order to adopt it.

Similarly, look at the utterance from 20 seconds later that consolidated the group perspective and moved on within that perspective:

1:22:21	Jamie	Yeah. Compare two n one. So that the rounded n- (0.1) no the rounded one is better. Number one.
---------	-------	---

Here we see again the group intentionality in how the list artifact is being indexed. Now the specific detail of the artifact is named: “two n one.” In addition, the discussion of which rockets to compare, with its question of determining which nose cone performs better, is re-located within the larger context of the design situation.

It should now be clear that the group discourse is itself 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.

Group Discourse as Emergent Thinking

This chapter 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 intersubjective) 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. Here, one can choose to view things at

the individual unit of analysis where traditional individual cognition takes place *or* at the group unit of analysis. The individual mechanisms are taken as primary and the group phenomena are seen as emergent. This is not the view of group discourse as primary and individual thought as a mediated, internalized, derivative version of the primary social cognition. However, it is still worth considering this emergent conception.

Emergence occurs on various scales; it has quite different characteristics and mechanisms in these different guises (Johnson, 2001). We will distinguish three scales: large scale statistical emergence, mid-level adaptive system emergence and small-group emergence.

Statistical Emergence

Chemical properties can often be viewed as emergent phenomena that arise out of large numbers of particles, each following laws of physics. For instance, thermodynamic phenomena involving billions of atoms exhibit higher level characteristics, such as molecular movement appearing as heat. The lower-level behaviors of the individual molecules are covered by the laws of physics. But their mass interactions exhibit qualities such as temperature and pressure, which are studied by chemistry rather than quantum mechanics. This transformation of individual motions to group qualities can be modeled by statistical analysis. The distinction in levels of analysis gives rise to distinct sciences, each with their own methodologies: biology cannot be reduced to chemistry, or chemistry to physics.

Adaptive System Emergence

Thousands of ants each following simple rules of behavior and interaction exhibit meta-level behaviors, such as efficient work organization and group foraging strategies. The lower-level rules are biologically evolved through success at the meta-level. These connections can be modeled by parallel computational systems encoding simple rules, such as *StarLogo*, *SimCity* and *AgentSheets*. In such systems, there are simple units whose behavior follows small sets of simple rules. The rule-governed behaviors interact in ways that allow groups of these units to follow patterns of behavior and to adapt to their context. In this way, group-level behaviors emerge from interactions at the lower level.

Small Group Emergence

Traditional human social interaction typically takes place among up to 150 people. It differs from the other kinds of emergence in that it does not involve statistically significant numbers of individuals and the rules they follow or the rules by which they interact are not simple. Small-group interaction is governed by very complex, subtle, interpreted, negotiated, mutually constituted rules. These depend on:

- Biologically evolved capabilities of human brains to interpret the behavior of other people, to recognize individuals and to maintain models of their minds.
- Culturally transmitted social practices that have accumulated over millennia.
- Language as a medium for conducting social interaction.
- Language as a tool for interpreting social interaction.

- Education, training and experience of a lifetime.

As we saw in the critique of AI, the determinants of human group behavior cannot even be made explicit and stated as rules. Nevertheless, human groups exhibit behaviors that cannot be predicted from an understanding of the individuals involved. For instance, families, neighborhoods, villages and cities emerge with complex structures and behaviors. Sociology cannot be reduced to psychology, let alone to biology.

Discourse

The emergence of group cognition is somewhat distinct from the emergence of social phenomena as discussed above. Conversation is the interaction of utterances, gestures, etc. from a small number of people. Often it involves only two people. Internal discussion or thought is generated by one person, although it may incorporate multiple internalized perspectives. The interaction can, nevertheless, 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 from 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 representationalism approach of AI.

Group Cognition and CSCL

Many methodologies popular in CSCL research focus on the individual as the unit of analysis: 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 an 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 oriented overwhelmingly toward the individual.