Chapter 9

Conclusion

The primary goal for this book was to help move discussion from a vague *notion* of collective intelligence to a reasonably well-defined *concept*. The issue was constrained to groups of limited size and duration doing intellectual tasks. By focusing on conceptual work, I was able to consider collaborative groups as a form of information processing system, analogous to Newell's and Simon's IPS model of individual cognition. The bulk of the discussion, then, became an examination of that system and its main components.

A collective memory system was described that includes subsystems that provide a collective long-term memory for tangible knowledge, built and maintained in a computer system, and for intangible knowledge, carried in the heads of the human beings that comprise the group. The memory system also includes working memory for both types of information.

A collective processor was described next. It includes the fine-grain operations used by groups to develop, access, and maintain the information stored in the memory system. Three basic types of processes were discussed: computer-, group-, and conference-mediated cognition cycles, depending on the situation in which they occur and the technology being used. The collective processor as a whole was described as a loosely coupled distributed system that includes multiple independent processors, joined by communications and social networks.

Collective strategy enables coherence in collaborative work. Individual processes occur not in isolation but in purposeful sequences. We can think of these strings of operations as being similar to statements in a language intended to accomplish a goal or to communicate a message. The system responsible for generating sequences of operations is analogous to the grammar individuals use to generate strings of words. Consequently, one way to study strategic behavior is to identify the rules for a grammar that can parse the

sequences of operations observed in groups. Such a grammar would represent an analytic model of collective strategy.

The discussion concluded by considering two metacognitive issues: collective awareness and collective control. Many projects are too large and too complex to be understood by any one person. Yet, we would like to see groups produce work with the same integrity and consistency sometimes found in work produced by a single good mind working alone. By developing thick, overlapping areas of shared knowledge, perhaps through slow build/fast review cycles, groups may be able to piece together a form of collective, but distributed awareness that is sufficiently coherent to achieve this goal. Control must also be distributed over a group. Otherwise, information will not flow across boundaries, and the group and its work will be brittle. However, although many decisions can, and probably should, be made by consensus, authority must ultimately be centralized in order to resolve disagreements and to preserve the integrity of the group's work.

What I hope has emerged from this discussion is an *image* of collective intelligence. It is a highly abstract image that references other systems and models as well as the physical world. However, I have tried to sketch it in enough detail so that the reader can "see" it in the mind's eye — both the shape of the whole and its main components. I have also tried to describe it "in motion" so that it can be seen as a dynamic system, in the process of building large conceptual structures. But, it is only a sketch — more suggestive than definitive. A great deal of research will be needed to flesh out its outline and to correct its mistakes. Nevertheless, I hope the discussion, even in its current form, has shown that a theory of collective intelligence is possible.

Perhaps that work can be motivated by considering the potential value a theory of collaboration could have. In the preface, I identified two primary audiences for this book: people involved in research in collaboration behavior and systems, and, second, people who work collaboratively or are simply interested in the topic. Let's look at this issue for both groups.

Currently, research in Computer Supported Cooperative Work is highly fragmented. Most teams focus on only a small part of the problem with little thought about how their work fits into a larger whole or how contextual factors impact their results. Yet, patterns of behavior observed in one situation may not appear under other

circumstances. For example, behavior in an early design meeting held in a group's normal working environment may be quite different from behavior that occurs in a laboratory setting or in a down stream review session. The same may be true for groups working on small, contrived problems versus large, actual problems in which they have a vested interest. If the research community saw itself as working toward a general theory of collaboration, individual projects would be encouraged to think more about how their part of the problem relates to work going on elsewhere. And, as that theory emerged, it would provide continuity across disparate studies.

This problem of isolation in the CSCW research community is not more pronounced than in other areas, but it is more ironic. Indeed, if the topic of research is *cooperation*, should not those in the field be working cooperatively, themselves? Working on a general theory of collaboration is tantamount to working within a paradigm. But, as Kuhn (1962) observed, the emergence of a paradigm marks the mature stage of a discipline. We may not be at that point yet in collaboration research. Indeed, not everyone believes a theory of collective intelligence is even possible. But, regardless of one's position on that issue, it is not too early to identify a common set dimensions that encompass all research in this field. By identifying those dimensions, and the points that fall along them, we can begin to see where individual projects lie, how research done by one group relates to that of another, and ways groups could work together more effectively.

To illustrate this idea of research dimensions, consider the five types of information built and used by collaborative groups: target and instrumental products, shared and private intangible knowledge, and ephemeral products. Each of these five types can be associated with a point along an *information type* dimension.[8] A number of studies — previous or potential — can then be viewed as examining movement of information along this dimension. For example, considering how intangible knowledge becomes part of the artifact can be viewed as a study of the transformation processes that map types found at one point to types associated with a different point. A related question is how the artifact, which represents general tangible knowledge with respect to the group, is comprehended by a new member and, thus,

---

8 This dimension is *ordinal*, consisting of identifiable states and implying no concept of distance between points. Other dimensions are *scalar*, implying a measurable distance relationship between points.

transformed into private intangible knowledge relative to that person. Similarly, one could consider the roles ephemeral products play in both of these processes.

Other dimensions include the spectrum of tools groups use, the duration over which behavior is observed and characterized, physical space, the size and coherence of the group, and the various processes groups use to do their work. Possible points along these dimensions include the following:

Information Type

intangible private
intangible shared
ephemeral
tangible instrumental
tangible target

Tools

noncomputer
computer applications
database and/or distributed file system
hypermedia
data communication
computer conferencing
audio/visual communication
transparent wide area network
intelligent agents
observational

Time

10 sec
100 sec (several minutes)
1k sec (~15 minutes)
10k sec (several hours)
100k sec (several days)
1m sec (several weeks)
10m sec (several months)
100m sec (several years)
1b sec (life's work)

Space

office
cluster
floor
building
site
1-hour travel
1-day travel

People

individuals
informal coalitions
teams
collection of teams
group as a whole

Processes

cognitive/conceptual
metacognitive
social interaction
mediated
organizational
collective/distributed
algorithmic

Dimensions become more interesting when they are considered relative to one another, like *x*,*y* axis. Two dimensions can define a plane of points, whereas three or more can define spaces. Each point in one of these planes or spaces can be associated with a particular set of issues. For example, we could ask which tools and functions provided by a collaboration support system are used to work with which particular types of information? This question involves two dimensions: *tools* and *information types*. Each point represents a particular combination of tool and information type, and a value for a given point might indicate whether the particular combination was observed in a given group. Similarly, considering whether particular tools are more or less useful for groups located in the same cluster of offices versus a widely distributed group involves the *tools* and *space* dimensions. If a study further broke the question down by type of

information, it would involve all three dimensions. We could continue expanding the issue by asking whether the effects are found at all stages of a project or only at certain times; whether the distribution of processes in groups differ according to the size and organization of the group; and so on. Thus, dimensions can be used individually, in pairs, and in higher order combinations.

If researchers could agree on a common set of dimensions, projects could use them to suggest ways their respective agendas might be extended incrementally to make their research more comprehensive and to identify other groups with whom they might collaborate or share data. Conversely, a common set of dimensions could also be used to suggest ways large studies might be divided among several research groups.

An agreed upon set of dimensions could also be helpful for developers of collaboration systems. Most current systems support a limited range of activities, such as shared asynchronous work (e.g., Lotus Notes) or group editing (e.g., Aspects). No commercial system supports the full range of tasks groups carry out, and I am aware of only one research system other than ABC that supports both synchronous and asynchronous work (i.e., SEPIA). Consequently, switching between individual and collective work is cognitively disruptive and requires moving from one computing context to another, with no guarantee that data built in one can be imported into the other. However, if developers viewed requirements against a backdrop of potential functions, the specific tasks those functions support, and the varied conditions under which groups work, it could help them see ways in which their systems could be extended to make them more useful or easier to use. By identifying major seams in the design space, a common set of dimensions could also facilitate developing standards for collaboration systems. Standards, in turn, could have a major impact on future designs, enabling data to be passed easily across boundaries and leading to more modular software that could be used in combination.

Thus, a long-term commitment to building a theory of collective intelligence and to building a more structured awareness of its associated research space could help the CSCW community work with greater continuity and coherence as well as inform system development.

A theory of collective intelligence could also benefit people who work collaboratively but are not researchers or system builders. I have already mentioned one potential benefit — better support tools. Other benefits are related to the ways in which a theory could help groups work more productively.

The issue boils down to the proposition that if the people who form groups had a better understanding of the collaborative process, they would be able to work together more effectively. That understanding would include a mental image of the overall process, strategies to guide them through it, and a view of their own roles in that process, including relationships with other individuals and/or subgroups. It should also help them see, at any given moment, where a project is relative to its long-term agenda and schedule. This proposition assumes that knowledge and awareness will result in better performance, both by individuals and the group as a whole. This may or may not be true, but it is an hypothesis worth testing. In the meantime, I would argue the point by analogy.

A large part of learning any knowledge-construction skill is learning strategies, particularly after one has mastered the rudiments of the skill. Implicit in a strategy is a comprehensive view of the task. I suppose one could argue that specific strategies can be regarded as individual large-grain processes to be used in particular circumstances without the individual having to think about the overall task, but that seems a perverse point of view. Instead, the working assumption in most skill development training seems to be that people need an overview of the process in addition to specific operations they can perform. I am simply arguing that groups need the same thing with respect to collaboration.

We should note, however, that collaboration, as a skill, differs from most conventional skills. It is inherently a second order phenomenon. That is, it is a metaskill that includes first order skills as components. For example, when one takes training in writing, one does so in order to write better documents or to write them more easily. Thus, there is a direct relationship between the skill — writing — and the results of applying that skill — a document. Collaboration is different. It is a layer of behavior that lies on top of conventional skills. For example, groups frequently co-author documents. Writing is the primary skill that actually produces the document; collaboration is the secondary skill — important, but instrumental.

This is as it should be. Nevertheless, it would be a healthy development if a theory of collective intelligence could lead to pedagogy for developing collaborative skills. Students at all levels could be taught collaboration skill, in addition to learning basic, first order skills. A collaboration curriculum could include instruction and practice in specific modes of activity. Thus, groups could go through explicit brainstorming activities, learn how to give and receive oral briefings, perform structured reviews of one another's work, do cross editing on co-authored documents to produce a consistent tone and voice, and so on. They could also learn strategies for using these various modes. And they could learn how to make effective use of both the computer and conventional tools available to them. Although collaboration training could be tied to specific content domains, such as writing or programming, it could also be taught as an abstract skill in which the specific task or content area is left up to the individual team.

Training in collaboration will, no doubt, encourage that way of working. When this happens, it is likely to be more a response to changing conditions than a driving force. Collaboration is already the predominant mode of work in many industries and in many cultures. That trend will probably accelerate. Although this may lead to greater productivity, it may also challenge deeply held values. In some societies, individuality has been a cornerstone of their cultural heritages. They honor rugged individuals, tell stories about their deeds, and train people to work alone. Consequently, much of an individual's self-image and self-worth may be tied up in his or her sense of uniqueness.

At least one response to this (potential) problem is suggested by the concept of collective intelligence, itself. Individual conceptualization lies at the heart of group-mediated cognition. Indeed, we saw in chapter 7 that group-mediated cognition is driven by the insights of individuals that are then shared with the group. But groups must also make decisions, build substantial bodies of shared knowledge, and work from consensus. And the individuals in a group must know when it is time to exercise free and unrestrained individuality and when it is time to close ranks and move on. Collaboration is, thus, a rhythmic process that oscillates between individual and collective behavior, individual and collective thinking. To be sure, it constrains individuality in some circumstances, but behavior has always been constrained, whether the individual realizes it or not, by the culture, by standards of acceptable behavior, and by

laws. But, along with the constraint posed by collaboration comes a compensating benefit — the sense of being part of a group and contributing toward a common goal. Thus, all things in their season. A theory of collective intelligence can help people strike the right balance between individual, independent thought and deed and collective, interdependent forms of the same.

As we look toward the future, the technologies that are emerging promise collaboration and cooperative interaction on a worldwide scale. That much seems certain. What is less clear is to what purposes we will put that technology. No doubt part of it will be devoted to popular entertainment — the proverbial 500 channels of television — and to chat groups — the equivalent of international CB radio. I see nothing inherently wrong with such uses. But we should also use these resources in our work and in our efforts to solve some of the many problems the world faces. It is to those efforts that a theory of collaboration can contribute most. Thus, developing a theory of collective intelligence is more than a matter of casual or academic interest. It could mean the difference between being run over by the technology and harnessing it for worthwhile purposes. It is a goal worth pursuing.