Chapter 1

Introduction

The notion of *collective intelligence* (CI) is that a group of human beings can carry out a task as if the group, itself, were a coherent, intelligent organism working with one mind, rather than a collection of independent agents.

The idea — referred to by several different terms — has been around for some time, but with recent interest in collaborative and cooperative work, it is being heard more often. Usually, it carries with it a bit of blue sky or is part of a throw-away line. For example, a grant proposal might suggest that the computer system the project is building to support collaborative work might eventually lead to a form of collective cognition by its users. But what, exactly, does that mean? What mode of thinking would constitute collective intelligence? What would be its characteristics? Would we recognize it if we saw it or experienced it?

In this discussion, I examine the idea of collective intelligence in order to try to pin it down and put some flesh on its bones. Thus, I hope to move discussion from a vague *notion* of collective intelligence to a *concept* that is reasonably well-defined. In the long-term, perhaps those working in this field can eventually build a *theory* of collective intelligence that is sufficiently precise so that it can be tested and refined. If such a theory existed, it could have a number of useful consequences. For example, if we really understood how groups of individuals can occasionally and under particular circumstances meld their thinking into a coherent whole, we would have a better idea of how to build computer and communications systems to support them, how to train other groups to work this way, and how to organize projects and institutions to promote this mode of work. I hope this discussion is a first step toward these goals.

Not everyone believes such a theory is possible. For example, Allen Newell (1990) argued that it is impossible for any group to function as a coherent rational agent. His objection, which I discuss in more detail later, is based on the rate at which information can be transferred from one human being to another. He argued that the bandwidth is insufficient to permit the various members of a group to all share the same knowledge — a condition he believed would be required to achieve what I call collective intelligence. Newell's objection is an important one that must be answered. Although I cannot refute his premise, I try to build a path around that roadblock.

Constraints

Group intellectual activities take place on many different scales. I frame the issue narrowly in order to make the discussion as concrete as possible. Perhaps later these constraints can be relaxed and the concept extended to a broader range of groups.

The discussion is limited to *intellectual* groups that are building some type of concrete conceptual object, such as a technical report, a marketing plan, a computer system, a legislative bill, or an airplane design. Excluded from the discussion, then, are groups that are primarily social, those carrying out manual tasks, or collaborations that produce aesthetic objects.

The discussion is limited to groups that range in size from three to four individuals to a handful of such groups working together on a single project. Excluded, then, are two-person collaborations that often have highly personal dynamics that don't extend to larger groups and, at the other extreme, projects that involve hundreds or even thousands of people. However, within this band of 3–30 people we can consider many of the problems encountered by groups of all sizes as well as the first extrapolation from a single group to a collection of groups.

Third, the discussion is limited to groups working on tasks that last from several weeks to several years. Durations within this range are long enough to raise problems of conceptual coordination and integration of ideas and materials, yet they are sufficiently bounded that the work is not viewed as ongoing and the group becomes institutionalized or bureaucratic in its operations.

Fourth, I distinguish between *collaboration* and *cooperation*. Collaboration carries with it the expectation of a singular purpose and a seamless integration of the parts, as if the conceptual object were produced by a single good mind. For example, a well-done collaborative document has a clear purpose or message. The reader is unable to tell from internal cues which chapters or sections were written by which authors. The sections are also consistent with one another, and one section shows appropriate awareness of the contents of the other sections.

Cooperative work is less stringent in its demands for intellectual integration. It requires that the individuals that comprise a group or, for larger projects, a set of groups carry out their individual tasks in accord with some larger plan. However, in a cooperative structure, the different individuals or groups are not required to know what goes on in the other parts of the project, so long as they carry out their own assigned tasks satisfactorily.

For example, the various teams of biologists that are currently mapping the human genome normally concentrate their research on a single chromosome or portion of a chromosome (DHHS, 1990). Although it could be advantageous, one team does not necessarily have to monitor work going on in other portions of the DNA structure in order to achieve its goals, nor is one team required to reconcile its methods and results with those of other groups. Such integration may eventually come — indeed, we see glimpses of this as newly articulated genes are mapped against various diseases and abnormalities. But, for now, although work within groups may be collaborative, work among groups in this field is more separate and diverse, albeit still cooperative.

I make this distinction between cooperation and collaboration to further limit the discussion. It seems to me far easier to imagine a concept of collective intelligence existing within a collaborative project than in one that is cooperative. Indeed, I suggest that collective intelligence is a requirement for effective collaboration, at least as a goal or boundary condition. Consequently, I limit the rest of this discussion to collaborative groups.

To summarize the constraints outlined so far, I examine a concept of collective intelligence by considering how collaborative groups ranging in size from 3 to 30 individuals working together for periods of several weeks to several years can produce an intellectual product that represents accomplishment of the group's main goal so that the product has the characteristics we would expect had it been produced by a single good mind.

I make one final assumption. The discussion is constrained to collaborative groups that use a computer and communication system as an integral part of their work. A requirement for collective intelligence is achieving a critical level of coherence in the work of the group. Although I admit the possibility in the abstract that a group might achieve this level of coherence without using a computer system, I cannot personally envision how large groups could coordinate their efforts and integrate the products develop by their individual members to the degree required for CI without such a system. Thus, I assume that CI is a form of intellectual behavior that is at least partially enabled by the technology. Later, when we understand the phenomenon better, it may be possible to relax this constraint and observe or develop CI in groups working without computer assistance.

Intelligence Amplification

The view of collective intelligence as a form of behavior made possible by some form of mediating computer system places it within the general tradition of *intelligence amplification* (IA). This perspective takes the position that computer systems can be developed that partially mirror human mental functions; thus, by increasing the capacity or speed of operation of those functions, these systems can thereby increase or amplify the mental capacity of the human user working with them. As a result, quantitative increases in specific functions may produce qualitative differences in intellectual behavior, making the computer a necessary but not sufficient tool for enabling this mode of thinking.

Vannevar Bush (1945) is generally credited with originating the idea of intelligence amplification. Writing before the first commercial computers were developed, Bush described a hypothetical desk-like device he called the *memex* that would be implemented using microfilm technology. It would permit a human user to store vast quantities of data, add new information, but, most important, add cross-references at the bottom of any microfilm page that could be instantly followed to some other page. Thus, the human user could construct large networks of semantic relationships within the memex, drawing together vast quantities of data and then quickly and associatively move from one intellectual context to another. One could argue that the book — or at least a library of books — could similarly extend the capacity and precision of human long-term memory and that books do, in fact, include similar cross-references. Bush's innovation lay in the speed with which associative links could be followed to access new material — a second or two versus the minutes or even hours required to move from one printed volume to another.

It makes sense to talk about Bush's memex as an amplifying device in the following sense. He identified several key architectural features of human intelligence — long-term memory, semantic relationships, and associative access — and then provided within his memex — at least in theory — their operational counterparts, but with greater capacity (the microfilm store with its embedded semantic relationships) and comparable speed (associative access). Thus, Bush believed his device could amplify a specific set of human mental functions. No one has yet built a complete memex as Bush described the device. However, using more familiar computer technology, Doug Engelbart was the first to build a memex-like system (Engelbart, Watson, & Norton, 1973). In recognition of the goal to supplement human intelligence, Engelbart called one version of his system *Augment*. Today, many of the features first described by Bush and first built by Engelbart are routinely found in contemporary hypertext systems, some of which are discussed in chapter 3.

Just as IA systems make possible a type of mental behavior that would not be possible without them, so, I suggest, a particular type of collaboration support system may enable a form of collective mental behavior that would not be possible without it. These systems, I suspect, will be based on principles analogous to those for IA systems, but with important distinctions and extensions. We normally form collaborative groups for two reasons. First, the task is too large and/or there is not enough time for it to be done by one person. Second, no individual possesses all of the skills and/or knowledge required. However, when we (necessarily) assemble a group to overcome these problems, we inherently create other problems. Because the intellectual construct being developed by the group is likely to be too large to be known in its entirety by any one individual, it may lack intellectual integrity. Rather than being a structure that is deeply principled and elegantly simple — as we expect of the work of our best individual minds — it may emerge as an awkward assembly of incongruous pieces. Indeed, we have come to expect this of groups, as indicated by the old joke that a camel is a horse produced by a committee.

A computer system that can help groups approximate a CI will have to include, as a minimum, functions that help them perceive and address the overall structure and integrity of their work. It must include tools to help groups establish and maintain the internal consistency and coherence among the various information products they produce through the individual hands of their various members. Thus, it will have to amplify intellectual skills that are (relatively) strong in individuals but less so within groups. It may also include additional tools to facilitate access and version control, communication, joint work, and other group behaviors. But it cannot neglect the more basic requirements of intellectual integrity, coherence, and consistency.

Overview

The approach I take in building a concept of collective intelligence is to consider collaboration as a type of information processing activity. Thus, I look at several Information Processing System (IPS) models and architectures of individual cognition, identify key components and functions within them, and then identify constructs within collaborative groups that are recognizable as extrapolations of these components and functions. I should point out that there is no inherent reason to believe that a collective intelligence should necessarily resemble familiar models of individual cognition; it could have an entirely different structure. But, if we can see a resemblance between the construct identified as CI and commonly accepted models of human cognition, to which we attribute intelligence, then we are likely to be willing to attribute intelligence to that construct, as well. On the other hand, if the structure identified as CI were entirely different, it would require more extensive justification to extend the claim of intelligence to it.

The volume is dividend into two parts. Part I discusses foundation concepts that are used in Part II to build a concept of collective intelligence and to inform that discussion. Chapter 2 considers the range of activities found in collaborative groups as a result of differences in size, scale of work, task domain, and so forth, by considering three different collaboration scenarios. A simple model of basic information types and the flow of information as one type is transformed into another is also presented. Chapter 3 discusses computer support for collaboration. It reviews a range of system features that fall within the general category of CSCW systems and identifies key features needed to develop the different information types noted in chapter 2. It also describes one particular system in more detail that serves as the reference system for the rest of the discussion. Chapter 4 discusses IPS models and architectures in order to identify key components. It describes both general models/architectures as well as specific IPS models for particular tasks and particular circumstances (i.e., human–computer interaction).

Part II tries to build a concept of collective intelligence. Chapter 5 discusses the different memory systems found in computer-supported collaborative groups that can be recognized as extrapolations of human memory systems. These constructs function as a form of *collective memory* for collaborative groups. Chapter 6 identifies several types of conceptual processing found in collaborative groups that are analogous to individual conceptual processing. They form a concept of *collective processing*. Chapter 7 considers issues of *collective strategy* within large multigroup collaborative projects. Chapter 8 examines issues of *collective awareness* and *control*. Chapter 9 concludes the discussion by identifying a set of research dimensions that could provide a framework through which to view and relate a broad range of research and development in the field. It also looks briefly at implications for a theory of collective intelligence.