Collective Intelligence

in

Computer-Based Collaboration

John B. Smith

Department of Computer Science
The University of North Carolina
Chapel Hill, NC 27599-3175
919-962-1792 jbs@cs.unc.edu

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 [refs*], 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 will 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 help move discussion from a vague *notion* of collective intelligence to a *concept* that is reasonably well-defined. In the long-term, I hope that those of us working in this field can eventually build a *theory* of collective intelligence that is sufficiently precise that it can be tested and refined. Such a theory 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, as well as how to organize projects and institutions to promote this mode of work. I hope this discussion will be a first step toward these goals.

Of course, not everyone believes such a theory is possible. For example, Allen Newell has recently argued that it is impossible for any group to function as a coherent rational agent [Newell, 1990.i.name:Newell, Allen,;]. His objection, which I will discuss in more detail below, is based on the rate at which information can be transferred from one human being to another. He argues that this bandwidth is insufficient to permit the various members of a group to all share the same knowledge -- a condition he believes to be necessary in order to realize what I call here a Collective Intelligence. Newell's objection is an important one that must eventually be answered. While I can't refute his basic premise with respect to bandwidth, I will try to show that we can construct an alternative path around that roadblock.

Constraints

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

The discussion will be 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, an airplane design, etc. Excluded from the discussion, then, are groups that are primarily social, those carrying out manual tasks, or collaborations that produce aesthetic objects.

The discussion will be limited to groups that range in size from 3-4 individuals to a handful of such groups working together on a single project. Excluded, then, are two-person collaborations and, at the other extreme, projects that involve hundreds or 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 will be 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 group does not become 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 will have a clear purpose or message. The reader will not be able to tell from internal cues which chapters or sections were written by which authors. The sections will also be consistent with one another, and one section will show 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 single 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 aren't 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. Considerable coherence is given to this field through a central databank where researchers register their results [refs*]. However, while it could be advantageous, one team doesn't 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, while 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 in order 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 or coordinated. Indeed, I will suggest that collective intelligence is a *requirement* for effective collaboration, at least as a goal or boundary condition. Consequently, I will limit the rest of this discussion to *collaborative*, as opposed to cooperative or coordinated, groups.

To summarize the constraints outlined so far, I will examine a concept of collective intelligence by considering how collaborative groups ranging in size from 3-30 individuals working together for periods of several weeks to several years can produce an intellectual product that represents the 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 will be 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. While 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 will assume that CI is a form of intellectual behavior that is at least partially induced 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 cognition 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 develop 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 agent for inducing this mode of thinking.

Vannevar Bush is generally credited with originating the idea of Intelligence Amplification [Bush, 1945]. 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 then be instantly followed. 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, however, 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 basic human mental functions. No one has yet built a complete memex, as Bush described the device. However, Doug Engelbart, using more familiar computer technology, 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 [refs*].

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 type of collective mental behavior that would not be possible without it. These systems, I suspect, will be based on principals 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 a single individual. A second reason is that no single 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. Since 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 my emerge as an awkward assemblage of incongruous pieces. Indeed, we have come to expect this, as indicated by the characterization of a camel as typical of the handiwork of a committee or group. Two other related problems concern the internal consistency of large intellectual objects developed by groups and their coherence.

A computer system that can help a group approximate a CI will have to include, as a minimum, functions that help groups overcome these problems. That is, it must provide tools to help groups perceive and address issues of overall structure and integrity. And it must include tools to help groups establish and maintain the internal consistency and coherence among the different information products they produce through the individual hands of its 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 these more basic requirements.

Collective Intelligence is, thus, a mode of intellectual behavior that is partially induced by a particular type of computer system. I will return to this issue in Chapter 3 when I briefly review collaboration support systems and then describe one particular system that will serve as representative of the group for this discussion.

Overview

The approach I will take in building a concept of Collective Intelligence will be to consider collaboration as a type of information processing activity. Thus, I will look at several Information Processing System (IPS) models and architectures of individual cognition, identify key components or functions, and then identify constructs within collaborative groups that are recognizable extrapolations of these components or 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 book is dividend into two large sections. In Part 1, several foundation concepts are discussed that are then used in Part 2 to build a concept of a Collective Intelligence and to inform that discussion. Part 1 is comprised of Chapters 2-4. Chapter 2 considers the range of activities found in collaborative groups as a result of differences in size, scale of work, task domain, etc. by considering three different collaborative scenarios. However, in spite of these differences, similarities can be found in all three groups in terms of a simple model of basic information types and the flow of information from one type to another. Chapter 3 discusses computer support for collaboration. It reviews the range of system features that fall within the general category of CSCW systems and then identifies key features needed to develop the different information types noted in Chapter 2. It also describes in more detail one particular system that will be presumed in the remainder of the discussion. Chapter 4 discusses IPS models and architectures of cognition in order to identify key components needed to define a concept of Collective Intelligence. It begins by discussing general models/architectures and then specific IPS models for particular tasks and for human-computer interaction.

Part 2, comprised of Chapters 5-10, is concerned with the actual construction of a concept of Collective Intelligence. Chapter 5 discusses several different kinds of memory systems found in models of individual cognition and then identifies analogous constructs that can be recognized as extrapolations of these systems that are found in computer-supported collaborative groups. These constructs are, thus, identified as a form of Collective Memory for collaborative groups. Chapter 6 focuses on conceptual processing. As was done for memory, analogous constructs for several different kinds of conceptual processing are identified for collaborative groups to form a concept of Collective Processing. The next two chapters discuss metacognitive issues. Chapter 7 considers issues of Collective Strategy within large multi-group collaborative projects. Chapter 8 examines issues of Collective Awareness and Control. In each of these chapters, specific issues for further research are identified. Chapter 9 discusses 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. Using it, researchers would be able to see other projects or studies that are "closely" related to theirs as well as ways in which small incremental additions to their research designs might make their work more comprehensive. Chapter 10 provides a brief conclusion and a further look toward the future.