

## Preface

Like many books, this one has grown in part out of personal experience. For the first 10 or 12 years of my professional life, I tended to work alone. My research interests centered on computer-assisted text and natural language analysis. I was fortunate enough to hold an academic appointment at a large university. At the time, my work was regarded as unusual if not eccentric. It was tolerated by my colleagues, even supported, but none of them shared those interests. Of course, I had professional colleagues at other institutions whom I saw at conferences, and we shared a number of good times talking and arguing ideas. But, for the most part, I had to generate my own motivation and to rely largely on my own ideas.

When I came to the University of North Carolina (UNC) some 10 years ago, one of the joys of that move was the chance to work closely with others. My first collaborators were Steve Weiss, a computer scientist; Marcy Lansman, a cognitive psychologist; Jay Bolter, a classicist; plus five or six graduate students, most of whom were from computer science but also including several from cognitive psychology. Since that time, our group has grown to include some 10 faculty and nearly twice that many graduate students, and we have expanded the multidisciplinary character of the group to include anthropologists.

Our earliest collaborations were concerned with hypermedia computer systems and their application to technical and scientific writing. From the beginning, we believed that if we could understand more clearly the cognitive process of writing, then we should be able to build computer systems consistent with that process. That is, if we could identify key mental activities that comprise expository writing, then we should be able to build corresponding features into our computer systems to support and, we hoped, enhance those same activities.

Our interests in writing and in individual computer users continue, but the focus of our research began to shift several years ago toward issues of collaboration. Thus, what was first a way of working became a topic of research. Of course, our studies of collaboration have gone beyond our own group, but I have benefited greatly by being able to observe as well as participate in an evolving group and to

test ideas about collaborative theory against collaborative practice, firsthand.

We were initially drawn in this direction by a desire to better understand collaborative writing and to adapt our writing system so that it could support groups jointly authoring documents. We have since extended the task domain to software development. However, the shift from writing to software development and from individual users to collaborative groups is not so large a step as it may first appear.

Writing documents is a metatask for many different forms of intellectual work. This is true in two respects. First, many intellectual activities ultimately produce a document that records and communicates the results of that activity. For example, when one plans a research project, one normally expresses that plan as a document. Similarly, software development usually involves writing a number of different documents; these include requirements, specifications, the design or architecture of the system, test plans, and user instructions. In fact, it's hard to imagine a substantial intellectual task that doesn't involve writing some form of document as an integral part of that task.

A second reason writing can be considered a metatask lies in the underlying mental processes it draws on and the constraints those processes are subject to. Working out the abstract structure of ideas that is the content of a document is as much a part of the "writing process" as expressing those ideas in words and sentences. Many intellectual tasks involve working out some form of preliminary plan and then building or expressing in detail the individual components that make up the plan. But, in doing so, we inevitably discover problems or subtleties not foreseen during the planning stage. Consequently, we must go back and revise the plan, which, in turn, may lead to additional changes needed to make the rest of the plan consistent with those changes. Thus, the processes of planning, representing, evaluating, and revising are used both iteratively and recursively, as work descends from abstraction into detailed expression or physical realization. As a result, this *overall* process of knowledge-construction is the same, regardless of whether the content is ultimately expressed in words, diagrams, code, or physical material.

Many of these same processes also operate within groups. For example, groups brainstorm together, particularly at the beginning of a new project. They also build plans together, review, and, at times,

edit and revise together. But, of course, the individual members that make up the group also work alone between joint work sessions. What is the overall pattern of behavior in a group as its members move back and forth from individual to collective work and from one process to another? If we expect groups to develop products that are coherent, internally consistent, and have intellectual integrity, then we need to understand this overall collaborative process.

What we would like to have, eventually, is a process model of collaboration similar to the process models that have been developed for tasks performed by individuals. We would like that model to be sufficiently general so that it applies to different task domains and to groups working in different organizational contexts. If I am correct in believing that conceptual construction tasks of all kinds draw on a common set of underlying processes and constraints, a general theory of collaboration is possible, although it is likely to differ substantially from models of individual cognition.

I do not get so far as actually defining a model of collaboration in this discussion, but that is the direction in which I am headed. What I do hope to accomplish is to sketch in some detail an image of collaboration as an information processing activity and to describe a framework for research that can help us build on one another's work in a way that may eventually lead to such a theory.

The term I use for this view of collaboration is *collective intelligence* (CI). Because much of the discussion is an attempt to define this concept, I do not describe it in more detail here but, instead, comment on several of its attributes. I associate CI with a set of goals or boundary conditions for collaboration. In trying to identify the essential processes and constraints that operate within groups, I am inevitably drawn toward describing patterns of behavior that can help groups produce products that have greater coherence and consistency. Thus, the discussion includes at various points a sense that some particular way of working may be preferable to another. However, I do not mean to advocate any specific behavior. Biomechanics provides an analogy for the fine line I am trying to walk. By analyzing the different muscle, bone, and ligament structures found in the human body, experts have been able to devise training programs that enhance or strengthen these different systems with the result that athletes who follow these regimens can run faster, jump higher, or throw further than they could otherwise. Similarly, if we understood the intellectual and social systems that operate within collaborative groups, we might eventually be able to train groups and

build tools to help them so that they could accomplish their tasks better than they could otherwise. But, because we do not yet have that knowledge, optimization and advocacy must remain topics for future consideration.

In the past, much more emphasis has been placed on the social dimensions of groups than on the cognitive. For example, prior studies have commonly examined patterns of influence within groups, qualities of leadership, the role of status, and so on. In this discussion, I have emphasized the cognitive and conceptual aspects of collaboration and the technology to support them. Thus, the discussion draws heavily on research from cognitive science and computer science. I have done this because collaborative groups are also fundamentally intellectual when their goal is to produce a conceptual artifact. Although I have also tried to show that cognitive processes are often closely, even fundamentally entwined with social processes, the discussion is imbalanced, in part to make up for this deficit on the cognitive side. I hope that future treatments of collaboration will establish the proper balance between the two.

To do so will require integrating theories and research from a number of disciplines, including anthropology, ethnography, management science, organizational theory, economics, sociology, social psychology, speech communication, and the study of small groups. Although I wish I had had the time and background to synthesize this material, apart from being impractical, I'm not sure, on reflection, it would have been desirable. Scholars who have spent a lifetime working in these fields can bring a depth of knowledge and discipline-based understanding that no single individual could ever assimilate. I hope that this discussion will pique their interest, that they will see ways in which they can expand or correct it, and that they will join an on-going *collaborative* enterprise to develop a valid theory of collective intelligence.

I have written this book with two primary audiences in mind. The first is anyone interested in collaboration who would like to think more about how groups operate and how technology may now and in the future affect the thinking, the social interactions, and the products produced by collaborative groups. The rapid pace of development in computer networks, distributed systems, and communications makes it increasingly possible for people to interact with one another, even when they are widely separated geographically. When we take into account the increasing tendency of organizations to encourage people to work in more flexible combinations with one another, we can

expect collaboration to become the predominant form of intellectual work. Consequently, this first group of potential readers includes just about anyone who does or is interested in intellectual work.

The second group of readers are my colleagues engaged in research and development in collaboration theory, studies, and systems. I have tried to sketch a comprehensive view of collaboration as a kind of intelligent organism. At present, the field of computer-supported cooperative work (CSCW) is driven largely by the technology and by enthusiasm, with the result that our activities and our systems are scattered and piecemeal. I believe it is not too early to attempt a more inclusive view of collaboration, comparable in scale and hope to the shift now taking place in cognitive science as they move from partial models to comprehensive architectures. Of course, we don't have a comparable body of prior research to build on, but we can generate this base of knowledge more quickly and more efficiently if we are aware of how our individual projects fit within a larger whole. What I have offered here is a candidate whole. Thus, my goal is to provide us with *something* to talk about, to debate, to correct, and perhaps, in part, to confirm.

There is a third group of readers who may find this discussion of interest. At present, CSCW is a highly multidisciplinary field, drawing concepts and techniques from a number of different disciplines, including computer science, human-computer interaction, cognitive psychology, anthropology, ethnography, sociology, organizational theory, small group theory, composition theory, economics, and, no doubt, other areas of which I am unaware. As a result, scholars and researchers working in these more mature "ancestor" disciplines may find some of the results emerging from CSCW, including this discussion, at least tangentially related to their own interests.

Although the discussion refers to a number of collaboration studies and support systems, it is not intended to be a comprehensive overview of work in the field. Rather, it is a "think piece" that tries to provide a concrete image of an abstract concept. As I gauge the rate of progress in CSCW, I estimate that it could have a useful half-life of about 5 years. When that time comes, I hope someone else will write a similar book that sketches another possible future from that vantage point.

Finally, I want to acknowledge some of those who have contributed to the ideas presented here. I have been fortunate over the past 10 years to have many outstanding collaborators. At this point, it is impossible to disentangle which ideas originated with me and which with them. A more accurate way to think about what took place is that we all participated in a process of intellectual evolution in which a large, amorphous conceptual structure gradually developed. What I have discussed here draws on that structure and is formed from my understanding of that larger community property we all share and own.

My earliest collaborators were Steve Weiss, Marcy Lansman, Jay Bolter, and Gordon Ferguson. I am indebted to them for their contributions to the earlier research on writing and writing systems, described later, out of which grew our later work on collaboration.

That original group was joined by Don Smith, Kevin Jeffay, Dotty Holland, Dana Smith, Peter Calingaert, and Hussein Abdel-Wahab to form the UNC Collaboratory Project. Reports of their work can be found throughout the book.

More recently, the UNC group has expanded to include David Stotts and Prasun Dewan. Although we have worked together only a short time, I want to acknowledge their contributions, along with that of my other UNC colleagues, to the discussion about future collaboration systems and the research that will be needed to achieve them.

During this period, I have been fortunate to work with a number of graduate students. It is they who have written the systems the Project takes credit for. Many of them have also played key roles in both the architecture and the conceptual basis of those systems.

I also want to acknowledge a debt to my colleague, Fred Brooks, from whom I learned the value of thinking hard in order to think simply.

A number of readers made valuable suggestions for improving the manuscript. Thanks to the students who read and discussed an early draft during a seminar on collaboration and to Don Smith for his helpful comments on several sections. Thanks, also, to Jessika Toral who began the index and to Claire Gingell who completed it and helped with many other tasks involved with preparing the manuscript. Brian Ladd prepared many of the illustrations, showing a real flair for the visual, for which I am grateful. I also want to thank Amy Pierce for her willingness to take a chance with this book, for her patience in

allowing me to complete it at my own pace, and for her encouragement throughout.

I am particularly indebted to Jan Walker for performing one of her patented very close, very thoughtful readings of the entire manuscript. Her comments led to a number of corrections and improvements. The problems that remain are, of course, the responsibility of the author.

Portions of the research done at UNC and reported in this book were sponsored by the National Science Foundation and the International Business Machines Corporation. I especially want to acknowledge Larry Rosenberg's efforts and leadership in establishing the Coordination Theory and Collaboration Technology program at NSF; without its support, this book and much of the work on which it is based could not have been done. I am also grateful to University of North Carolina for a research leave during which the book was begun.

Finally, I wish to acknowledge my oldest and best collaborator, Catherine Smith. She has lived these ideas with me, contributed to their development, and has never refused to read a section and give me her thoughtful views, even under the most trying circumstances.

*John B. Smith*