


Avr (8:22:28 PM): $A = 1/2bh$
Avr (8:22:31 PM): I believe
pin 805 (8:22:35 PM): yes
pin 805 (8:22:37 PM): i concue
pin 805 (8:22:39 PM): concur*
Avr (8:22:42 PM): then find the area of
Avr (8:22:54 PM): oh, wait
Sup (8:23:03 PM): the base and height a
Avr (8:23:11 PM): no

GROUP COGNITION

Computer Support for Building Collaborative Knowledge



Gerry Stahl

Brent: This one's different
Jamie: Yeah, but it has same nose
Chuck: Pointy nose cone=
Steven: =Oh, yeah=
Chuck: =But it's not the same engine
Jamie: Yeah it is,
Brent: =Yes it is,
Jamie: Compare two n one

Group Cognition

Acting with Technology

Bonnie Nardi, Victor Kaptelinin, and Kirsten Foot, editors

Tracing Genres through Organizations: A Sociocultural Approach to Information Design

Clay Spinuzzi, 2003

Activity-Centered Design: An Ecological Approach to Designing Smart Tools and Usable Systems

Geri Gay and Helene Hembrooke, 2004

The Semiotic Engineering of Human Computer Interaction

Clarisse Sieckenius de Souza, 2004

Group Cognition: Computer Support for Building Collaborative Knowledge

Gerry Stahl, 2006

Group Cognition

Computer Support for Building Collaborative Knowledge

Gerry Stahl

The MIT Press
Cambridge, Massachusetts
London, England

© 2006 Massachusetts Institute of Technology

All rights reserved. No part of this book may be reproduced in any form by any electronic or mechanical means (including photocopying, recording, or information storage and retrieval) without permission in writing from the publisher.

MIT Press books may be purchased at special quantity discounts for business or sales promotional use. For information, please e-mail <special_sales@mitpress.mit.edu> or write to Special Sales Department, The MIT Press, 55 Hayward Street, Cambridge, MA 02142.

This book was set in Sabon by SNP Best-set Typesetter Ltd., Hong Kong.
Printed and bound in the United States of America.

Library of Congress Cataloging-in-Publication Data

Stahl, Gerry.

Group cognition : computer support for collaborative knowledge building / Gerry Stahl.
p. cm.—(Acting with technology)

Includes bibliographical references and index.

ISBN 0-262-19539-9 (hc : ak. paper)

1. Computer-assisted instruction. 2. Computer networks. I. Title. II. Series.

LB1028.5.S696 2006

371.33'4—dc22

2005052047

10 9 8 7 6 5 4 3 2 1

Contents

Series Foreword vii

Introduction: Essays on Technology, Interaction, and Cognition 1

I Design of Computer Support for Collaboration

Studies of Technology Design 25

1 Share Globally, Adapt Locally 31

2 Evolving a Learning Environment 47

3 Armchair Missions to Mars 65

4 Supporting Situated Interpretation 81

5 Collaboration Technology for Communities 93

6 Perspectives on Collaborative Learning 119

7 Groupware Goes to School 155

8 Knowledge Negotiation Online 177

II Analysis of Collaborative Knowledge Building

Studies of Interaction Analysis 193

9 A Model of Collaborative Knowledge Building 201

10 Rediscovering the Collaboration 213

11 Contributions to a Theory of Collaboration 227

12 In a Moment of Collaboration 245

13 Collaborating with Relational References 257

III Theory of Group Cognition

Studies of Collaboration Theory	277
14 Communicating with Technology	285
15 Building Collaborative Knowing	303
16 Group Meaning / Individual Interpretation	331
17 Shared Meaning, Common Ground, Group Cognition	347
18 Making Group Cognition Visible	361
19 Can Collaborative Groups Think?	385
20 Opening New Worlds for Collaboration	409
21 Thinking at the Small-Group Unit of Analysis	431
Notes	469
References	479
Name Index	499
Subject Index	503

Series Foreword

The MIT Press Acting with Technology series is concerned with the study of meaningful human activity as it is mediated by tools and technologies. The goal of the series is to publish the best new books—both research monographs and textbooks—that contribute to an understanding of technology as a crucial facet of human activity enacted in rich social and physical contexts.

The focus of the series is on tool-mediated processes of working, playing, and learning in and across a wide variety of social settings. The series explores developments in postcognitivist theory and practice from the fields of sociology, communication, education, organizational studies, science and technology studies, human-computer interaction studies, and computer-supported collaborative work. It aims to encompass theoretical frameworks developed through cultural-historical activity theory, actor-network theory, distributed cognition, ethnomethodology, and grounded theory.

In *Group Cognition: Computer Support for Building Collaborative Knowledge*, Gerry Stahl challenges us with the provocative notion that “small groups are the engines of knowledge building.” He notes that research on learning has focused on either individual cognition or the larger community. Based on his extensive experience in teaching and system building, Stahl points to the “decisive role of small groups” in learning. Stahl’s contribution is to alert us to the need for a theoretical representation of small groups and their pivotal role in group cognition. He explores this theme in varied ways—empirical, theoretical, philosophical—each persuasive and thoughtful in its own way.

Stahl pushes hard on the notion of group cognition, proposing that we view discourse as a “substrate for group cognition.” *Discourse* is defined broadly to include spoken words, inscriptions, and body language. Using these notions, Stahl hopes to position cognition in that zone of small groups where he feels it belongs, moving

it away from individual “brains” but not too far into less precise entities such as community. Stahl notes that adopting such a notion would change education profoundly in terms of classroom practice, testing, assessment, and teacher training.

Group Cognition is a welcome addition to the Acting with Technology series. It is highly recommended for readers interested in education, human-computer interaction, and computer-supported collaborative work.

Group Cognition

Introduction: Essays on Technology, Interaction, and Cognition

The promise of globally networked computers to usher in a new age of universal learning and sharing of human knowledge remains a distant dream; the software and social practices needed have yet to be conceived, designed, and adopted. To support online collaboration, our technology and culture have to be reconfigured to meet a bewildering set of constraints. Above all, this requires understanding how digital technology can mediate human collaboration. The essays gathered in this volume document one path of exploration of these challenges. They include efforts to design software prototypes featuring specific collaboration-support functionality, to analyze empirical instances of collaboration, and to theorize about the issues, phenomena, and concepts involved today in supporting collaborative knowledge building.

The studies in this book grapple with the problem of how to increase opportunities for effective collaborative working, learning, and acting through innovative uses of computer technology. From a technological perspective, the possibilities seem endless and effortless. The ubiquitous linking of computers in local and global networks makes possible the sharing of thoughts by people who are separated spatially or temporally. Brainstorming and critiquing of ideas can be conducted in many-to-many interactions, without being confined by a sequential order imposed by the inherent limitations of face-to-face meetings and classrooms. Negotiation of consensual decisions and group knowledge can be conducted in new ways.

Collaboration of the future will be more complex than just chatting—verbally or electronically—with a friend. The computational power of personal computers can lend a hand here; software can support the collaboration process and help to manage its complexity. It can organize the sharing of communication, maintaining both sociability and privacy. It can personalize information access to different user perspectives and can order knowledge proposals for group negotiation.

Computer support can help us transcend the limits of individual cognition. It can facilitate the formation of small groups engaged in deep knowledge building. It can empower such groups to construct forms of group cognition that exceed what the

group members could achieve as individuals. Software functionality can present, coordinate, and preserve group discourse that contributes to, constitutes, and represents shared understandings, new meanings, and collaborative learning that is not attributable to any one person but that is achieved in group interaction.

Initial attempts to engage in the realities of computer-supported knowledge building have, however, encountered considerable technical and social barriers. The transition to this new mode of interaction is in some ways analogous to the passage from oral to literate culture, requiring difficult changes and innovations on multiple levels and over long stretches of time. But such barriers signal opportunities. By engaging in experimental attempts at computer-supported, small-group collaboration and carefully observing where activity breaks down, researchers can identify requirements for new software.

The design studies presented in this book explore innovative functionality for collaboration software. They concentrate especially on mechanisms to support *group formation*, *multiple interpretive perspectives*, and the *negotiation of group knowledge*. The various applications and research prototypes reported in the first part of this book span the divide between cooperative work and collaborative learning, helping us to recognize that contemporary knowledge workers must be lifelong learners and also that collaborative learning requires flexible divisions of labor.

The attempt to design and adopt collaboration software led to a realization that we need to understand much more clearly the social and cognitive processes involved. In fact, we need a multifaceted theory for computer-supported collaboration, incorporating empirically based analyses and concepts from many disciplines. This book, in its central part, pivots around the example of an empirical micro-analysis of small-group collaboration. In particular, it looks at *how the group constructs intersubjective knowledge that appears in the group discourse itself*, rather than theorizing about what takes place in the minds of the individual participants.

The notion that it is important to take the group, rather than the individual, as the unit of analysis ultimately requires developing, from the ground up, a new *theory of collaboration* in the book's final part. This theory departs from prevalent cognitive science, grounded as it is on mental representations of individuals. Such a theory builds on related efforts in social-cultural theory, situated cognition, and ethnomethodology, as well as their post-Kantian philosophical roots.

Collaboration as Group Cognition

This book does not aspire to the impossible task of describing all the ways that technology does or could affect working and learning. I work and I learn in innumerable ways and modes—and everyone else works and learns in additional ways,

many different from mine. Working and learning with other people mixes these ways into yet more complex varieties. Technology multiplies the possibilities even more. So this book chooses to focus on a particular form of working and learning—one that seems especially attractive to many people and may be particularly responsive to technological support but one that is also rather hard to point out and observe in the current world. It is the holy grail of cooperative knowledge work and collaborative learning—*the emergence of shared group cognition* through effective collaborative knowledge building.

The goal of collaborative knowledge building is much more specific than that of e-learning or distance education generally, where computer networks are used to communicate and distribute information from one teacher to several students who are geographically dispersed. Collaborative knowledge building stresses supporting interactions among the students themselves, with a teacher playing more of a facilitating than instructing role. Moreover, knowledge building involves the construction or further development of some kind of knowledge artifact. That is, the students are not simply socializing and exchanging their personal reactions or opinions about the subject matter but might be developing a theory, model, diagnosis, conceptual map, mathematical proof, or presentation. These activities require the exercise of high-level cognitive activities. In effective collaborative knowledge building, the group must engage in thinking together about a problem or task and produce a knowledge artifact such as a verbal problem clarification, a textual solution proposal, or a more developed theoretical inscription that integrates their different perspectives on the topic and represents a shared group result that they have negotiated.

We all know from personal experience—or think we know based on our tacit acceptance of prevalent folk theories—that individual people can think and learn on their own. It is harder to understand how a small group of people collaborating online can think and learn as a group and not just as the sum of the people in the group thinking and learning individually.

Ironically, the counterintuitive notion of group cognition turns out to be easier to study than individual learning. Whereas individual cognition is hidden in private mental processes, group cognition is necessarily publicly visible. This is because any ideas involved in a group interaction must be displayed for the members of the group to participate in the collaborative process. In this book, I try to take advantage of such displays to investigate group cognition without reducing it to an epiphenomenon of individual cognition. This does not mean that I deny that individuals have private thoughts: I simply do not rely on our commonsense intuitions and introspections about such thoughts. In the end, consideration focused on the group unit may have implications for understanding individual cognition as a socially grounded and mediated product of group cognition.

How does a group build its collective knowing? A noncognitivist approach avoids speculating on psychological processes hidden in the heads of individuals and instead looks to empirically observable group processes of interaction and discourse. The roles of individuals in the group are not ignored but are viewed as multiple interpretive perspectives that can conflict, stimulate, intertwine, and be negotiated. The spatiotemporal world in which collaborative interactions are situated is not assumed to be composed of merely physical as opposed to mental ideas but is seen as a universe filled with meaningful texts and other kinds of artifacts—human-made objects that embody shared meanings in physical, symbolic, digital, linguistic, and cultural forms.

The concern with the processes and possibilities of building group knowing has implications for the choice of themes investigated in this book. The software prototypes reported on in part I, for instance, were attempts to support the formation of teams that had the right mix for building knowledge as a group, to represent the multiple perspectives involved in developing group ideas, and to facilitate the negotiation of group knowledge that arose. Certainly, there are other important processes in online collaboration, but these are of particular concern for small-group knowledge building. Similarly, the empirical analysis in part II zooms in on the way in which the participants in an observed group of students constructed knowledge in their discourse that could not be attributed to any simple conjunction of their individual contributions. Finally, the theoretical reflections of part III try to suggest a conceptual framework that incorporates these notions of “interpretive perspectives” or “knowledge negotiation” within a coherent view of how group cognition takes place in a world of discourse, artifacts, and computer media.

Rather than centering on practical design goals for *computer-supported cooperative work* (CSCW) industrial settings or *computer-supported collaborative learning* (CSCL) classrooms, the following chapters explore foundational issues of how small groups can construct meaning at the group level. The ability of people to engage in effective group cognition in the past has been severely constrained by physical limits of the human body and brain. We can really relate to only a small number of individual people at a time or follow only one primary train of thought at a time, and most business meetings or classroom activities are structured, moderated, and delimited accordingly. Moreover, we quickly forget many of the details of what was said at such meetings. Collaboration technology has enormous potential to establish many-to-many interactions, to help us manage them, and to maintain logs of what transpired. Figuring out how to design and deploy collaboration technologies and social practices to achieve this still-distant potential is the driving force that is struggling to speak through these essays.

The structure of the book follows the broad strokes of my historical path of inquiry into computer-supported group cognition. Part I reports on several attempts

to design online technologies to support the collaborative building of knowing—that is, computer-mediated group sense making—in which I was involved. Part II shows how I responded to the need I subsequently felt to better understand phenomena of collaboration—such as group formation, perspective sharing, and knowledge negotiation through microanalysis of group interaction—in order to guide such software design. In turn, part III indicates how this led me to formulate a conceptual framework and a research methodology: a theory of collaboration, grounded in empirical practice and exploration. Although theory is typically presented as a solid foundational starting point for practice, this obfuscates its genesis as a conceptual reflection in *response* to problems of practice and their circumstances. I have tried to avoid such reification by presenting theory at the end of the book because it emerged as a result of design efforts and empirical inquiry.

The Problematic of CSCL and the Approach of This Book

This book documents my engagement with the issues of CSCL as a research field. Although I believe that much of the group-cognition approach presented is also applicable to CSCW, my own research during the decade represented here was more explicitly oriented to the issues that dominated CSCL at the time. In particular, CSCL is differentiated from related domains in the following ways:

- *Group* The focus is not on individual learning but on learning in and by small groups of students.
- *Cognition* The group activity is not one of working but of constructing new understanding and meaning within contexts of instruction and learning.
- *Computer support* The learning does not take place in isolation but with support by computer-based tools, functionality, microworlds, media, and networks.
- *Building* The concern is not with the transmission of known facts but with the construction of personally meaningful knowledge.
- *Collaborative* The interaction of participants is not competitive or accidental but involves systematic efforts to work and learn together.
- *Knowledge* The orientation is not to drill and practice of specific elementary facts or procedural skills but to discussion, debate, argumentation, and deep understanding.

The fact that these points spell out the title of this book is an indication that the book consists of an extended reflection on the defining problems of CSCL.

The history of CSCL research and theory can be schematically viewed as a gradual progression of ever-increasing critical distance from its starting point, consisting of conceptualizations of learning inherited from dominant traditions in the fields of

education and psychology. Much of the early work in CSCL started from this individualistic notion of learning and cognition. For instance, the influence of artificial intelligence (AI) on CSCL—which can be seen particularly clearly in my first three studies—often relied on computational cognitive models of individual learners. For me, at least, dramatic shifts away from this tradition came from the following sources:

- *Mediated cognition* Vygotsky’s work from the 1920s and 1930s only became available in English 50 years later, when it proposed a radically different view of cognition and learning as socially and collaboratively mediated.
- *Distributed cognition* This alternative, developed by a number of writers (including Suchman, Winograd, Pea, and Hutchins), also stressed the importance of not viewing the mind as isolated from artifacts and other people.
- *Situated learning* Lave’s work applied the situated perspective to learning, showing how learning can be viewed as a community process.
- *Knowledge building* Scardamalia and Bereiter developed the notion of community learning with a model of collaborative knowledge building in computer-supported classrooms.
- *Meaning making* Koschmann argued for reconceptualizing knowledge building as meaning making, drawing on theories of conversation analysis and ethnomethodology.
- *Group cognition* This book arrives at a theory of group cognition by pushing this progression a bit further with the help of a series of software-implementation studies, empirical analyses of interaction, and theoretical reflections on knowledge building.

The notion of group cognition emerged out of the trajectory of the research that is documented in this volume. The software studies in the early chapters attempted to provide support for collaborative knowledge building. They assumed that collaborative knowledge building consisted primarily of forming a group, facilitating interaction among the multiple personal perspectives brought together, and then encouraging the negotiation of shared knowledge. When the classroom use of my software resulted in disappointing levels of knowledge building, I tried to investigate in more detail how knowledge building occurs in actual instances of collaborative learning.

The explorative essays in the middle of the book prepare the way for that analysis and then carry out a microanalysis of one case. The fundamental discovery made in that analysis was that, in small-group collaboration, *meaning is created across the utterances of different people*. That is, the meaning that is created is not a cog-

nitive property of individual minds but a characteristic of the group dialogue. This is a striking result of looking closely at small-group discussions; it is not so visible in monologues (although retrospectively these can be seen as internalized discourses of multiple voices), in dialogues (where the utterances each appear to reflect the ideas of one or the other member of the dyad), or in large communities (where the joint meaning becomes fully anonymous). I call this result of collaborative knowledge building *group cognition*.

For me, this discovery—already implied in certain social science methodologies like conversation analysis—led to a conception of group cognition as central to understanding collaboration and consequently required a rethinking of the entire theoretical framework of CSCL: collaboration, knowledge, meaning, theory building, research methodology, design of support. The paradigm shift from individual cognition to group cognition is challenging—even for people who think they already accept the paradigms of mediated, distributed, and situated cognition. For this reason, the essays in the last part of the book not only outline what I feel is necessary for an appropriate theory but also provide a number of reflections on the perspective of group cognition itself. While the concept of group cognition that I develop is closely related to findings from situated cognition, dialogic theory, symbolic interactionism, ethnomethodology, and social psychology, I think that my focus on small-group collaboration casts it in a distinctive light particularly relevant to CSCL. Most important, I try to explore the core phenomenon in more detail than other writers, who tend to leave some of the most intriguing aspects as mysteries.

Accomplishing this exposition on group cognition requires spelling out a number of interrelated points, each complex in itself. A single conference or journal paper can enunciate only one major point. This book is my attempt to bring the whole argument together. I have organized the steps in this argument into three major book parts:

Part I, *Design of Computer Support for Collaboration*, presents eight studies of technology design. The first three apply various AI approaches (abbreviated as DODE, LSA, CBR) to typical CSCL or CSCW applications, attempting to harness the power of advanced software techniques to support knowledge building. The next two shift the notion of computer support from AI to providing collaboration media. The final three try to combine these notions of computer support by creating computational support for core collaboration functions in the computational medium. The chapters discuss how to

1. Support teacher collaboration for constructivist curriculum development (written in 1995),
2. Support student learning of text production in summarization (1999),

3. Support the formation of groups of people who can work effectively together (1996),
4. Define the notion of personal interpretive perspectives of group members (1993),
5. Define the role of computational media for collaborative interactions (2000),
6. Support group and personal perspectives (2001),
7. Support group work in collaborative classrooms (2002), and
8. Support the negotiation of shared knowledge by small groups (2002).

Part II, *Analysis of Collaborative Knowledge Building*, consists of five essays related to research methodology for studying small-group interaction. First, a process model of knowledge building shows how utterances from multiple perspectives may be negotiated to produce shared knowledge. Second, methodological considerations argue that the most important aspects of collaboration are systematically obscured by the approach taken by many leading CSCL studies. A solution is then proposed that integrates knowledge building and merged perspectives with artifacts from distributed cognition theory and the close interpretation of utterances from conversation analysis. This solution is applied to an empirical case of collaboration. This case reveals how group cognition creates shared meaning through the thick interdependencies of everyone's utterances. It also shows how the group builds knowledge about meaning in the world. In particular, these chapters provide

9. A process model of collaborative knowledge building, incorporating perspectives and negotiation (2000),
10. A critique of CSCL research methodologies that obscure the collaborative phenomena (2001),
11. A theoretical framework for empirical analysis of collaboration (2001),
12. Analysis of five students who are building knowledge about a computer simulation (2001), and
13. Analysis of the shared meaning that they built and its relation to the design of the software artifact (2004).

Part III, *Theory of Group Cognition*, includes eight chapters that reflect on the discovery of group meaning in chapter 12 and its further analysis in chapter 13. As preliminary context, previous theories of communication are reviewed to see how they can be useful, particularly in contexts of computer support. Then a broad-reaching attempt is made to sketch an outline of a social theory of collaborative knowledge building based on the discovery of group cognition. A number of specific issues are taken up from this, including the distinction between meaning making

at the group level versus interpretation at the individual level and a critique of the popular notion of common ground. Chapter 18 develops the alternative research methodology hinted at in chapter 10. Chapters 19 and 20 address philosophical possibilities for group cognition, and the final chapter complements chapter 12 with an initial analysis of computer-mediated group cognition, as an indication of the kind of further empirical work needed. The individual chapters of this final part offer

14. A review of traditional theories of communication (2003),
15. A sketch of a theory of building collaborative knowing (2003),
16. An analysis of the relationship of group meaning and individual interpretation (2003),
17. An investigation of group meaning as common ground versus as group cognition (2004),
18. A methodology for making group cognition visible to researchers (2004),
19. Consideration of the question, “Can groups think?” in parallel to the AI question, “Can computers think?” (2004),
20. Exploration of philosophical directions for group-cognition theory (2004), and
21. A wrap-up of the book and an indication of future work (2004).

The discussions in this book are preliminary studies of a science of computer-supported collaboration that is methodologically centered on the group as the primary unit of analysis. From different angles, the individual chapters explore how meanings are constituted, shared, negotiated, preserved, learned, and interpreted socially by small groups within communities. The ideas these essays present themselves emerged out of specific group collaborations.

Situated Concepts

The studies of this book are revised forms of individual papers that were undertaken during the decade between my dissertation at the University of Colorado and my research at Drexel University and were published on various specific occasions. In bringing them together, I have tried to retain the different voices and perspectives that they expressed in their original situations. They look at issues of online collaboration from different vantage points, and I wanted to retain this diversity as a sort of collaboration of me with myself—a collection of selves that I had internalized under the influences of many people, projects, texts, and circumstances. The format of the book thereby reflects the theory it espouses: that knowledge emerges from situated activities involving concrete social interactions and settings and that

such knowledge can be encapsulated in vocabularies and texts that are colored by the circumstances of their origins.

Thus, the main chapters of this book are self-contained studies. They are reproduced here as historical artifacts. The surrounding apparatus—this overview, the part introductions, the chapter lead-ins, and the final chapters—has been added to make explicit the gradual emergence of the theme of group cognition. When I started to assemble the original essays, it soon became apparent that the whole collection could be significantly more than the sum of its parts, and I wanted to bring out this interplay of notions and the implications of the overall configuration. The meaning of central concepts, like *group cognition*, are not simply defined; they evolve from chapter to chapter in the hope that they will continue to grow productively in the future.

Concepts can no longer be treated as fixed, self-contained, eternal, universal, and rational, for they reflect a radically historical world. The modern age of the last several centuries may have questioned the existence of God more than the medieval age, but it still maintained an unquestioned faith in a god's-eye view of reality. For Descartes and his successors, an objective physical world was knowable in terms of a series of facts that were expressible in clear and distinct propositions using terms defined by necessary and sufficient conditions. While individuals often seemed to act in eccentric ways, one could still hope to understand human behavior in general in rational terms.

The twentieth century changed all that. Space and time could henceforth be measured only relative to a particular observer; position and velocity of a particle were in principle indeterminate; observation affected what was observed; relatively simple mathematical systems were logically incomplete; people turned out to be poor judges of their subconscious motivations and unable to articulate their largely tacit knowledge; rationality frequently verged on rationalization; revolutions in scientific paradigms transformed what it meant in the affected science for something to be a fact, a concept, or evidence; theories were no longer seen as absolute foundations but as conceptual frameworks that evolved with the inquiry; and knowledge (at least in most of the interesting cases) ended up being an open-ended social process of interpretation.

Certainly, there are still empirical facts and correct answers to many classes of questions. As long as one is working within the standard system of arithmetic, computations have objective answers—by definition of the operations. Some propositions in natural language are also true, like “This sentence is declarative.” But others are controversial, such as “Knowledge is socially mediated,” and some are even paradoxical: “This sentence is false.”

Sciences provide principles and methodologies for judging the validity of propositions within their domain. Statements of personal opinion or individual observa-

tion must proceed through processes of peer review, critique, evaluation, argumentation, negotiation, refutation, and so on to be accepted within a scientific community; that is, to evolve into knowledge. These required processes may involve empirical testing, substantiation, or evidence as defined in accord with standards of the field and its community. Of course, the standards themselves may be subject to interpretation, negotiation, or periodic modification.

Permeating this book is the understanding of knowledge, truth, and reality as products of social labor and human interpretation rather than as simply given independently of any history or context. *Interpretation* is central. The foundational essay of part I (chapter 4) discusses how it is possible to design software for groups (groupware) to support the situated interpretation that is integral to working and learning. Interpretation plays the key analytic role in the book, with the analysis of collaboration that forms the heart of part II (chapter 12) presenting an interpretation of a moment of interaction. And in part III (particularly chapter 16), the concepts of interpretation and meaning are seen as intertwined at the phenomenological core of an analysis of group cognition. Throughout the book, the recurrent themes of multiple interpretive perspectives and of the negotiation of shared meanings reveal the centrality of the interpretive approach.

There is a philosophy of interpretation, known since Aristotle as *hermeneutics*. Hans-Georg Gadamer (1988) formulated a contemporary version of philosophical hermeneutics, based largely on ideas proposed by his teacher, Martin Heidegger (1996). A key principle of this hermeneutics is that the meaning of a term should be interpreted based on the history of its effects in the world. Religious, political, and philosophical concepts, for instance, have gradually evolved their meanings as they have interacted with world history and been translated from culture to culture. Words like *being*, *truth*, *knowledge*, *learning*, and *thought* have intricate histories that are encapsulated in their meaning but that are hard to articulate. Rigorous interpretation of textual sources can begin to uncover the layers of meaning that have crystallized and become sedimented in these largely taken-for-granted words.

If we now view meaning making and the production of knowledge as processes of interpretive social construction within communities, then the question arises of whether such fundamental processes can be facilitated by communication and computational technologies. Can technology help groups to build knowledge? Can computer networks bring people together in global knowledge-building communities and support the interaction of their ideas in ways that help to transform the opinions of individuals into the knowledge of groups?

As an inquiry into such themes, this book eschews an artificially systematic logic of presentation and, rather, gathers together textual artifacts that view concrete investigations from a variety of perspectives and situations. My efforts to build software systems were not applications of theory in either the sense of foundational

principles or predictive laws. Rather, the experience gained in the practical efforts of part I motivated more fundamental empirical research on computer-mediated collaboration in part II, which in turn led to the theoretical reflections of part III that attempt to develop ways of interpreting, conceptualizing, and discussing the experience. The theory part of this book was written to develop themes that emerged from the juxtaposition of the earlier, empirically grounded studies.

The original versions of the chapters were socially and historically situated. Concepts they developed while expressing their thoughts were, in turn, situated in the contexts of those publications. In being collected into the present book, these papers have been only lightly edited to reduce redundancies and to identify cross-references. Consistency of terminology across chapters has not been enforced as much as it might be to allow configurations of alternative terminologies to bring rich complexes of connotations to bear on the phenomena investigated.

These studies strive to be *essays* in the postmodern sense described by Theodor Adorno (1984, p. 160):

In the essay, concepts do not build a continuum of operations, thought does not advance in a single direction, rather the aspects of the argument interweave as in a carpet. The fruitfulness of the thoughts depends on the density of this texture. Actually, the thinker does not think, but rather transforms himself into an arena of intellectual experience, without simplifying it. . . . All of its concepts are presentable in such a way that they support one another, that each one articulates itself according to the configuration that it forms with the others.

In Adorno's book *Prisms* (1967), essays on specific authors and composers provide separate glimpses of art and artists, but there is no development of a general aesthetic theory that illuminates them all. Adorno's influential approach to cultural criticism emerged from the book as a whole, implicit in the configuration of concrete studies but nowhere in the book articulated in propositions or principles. His analytic paradigm—which rejected the fashionable focus on biographical details of individual geniuses or eccentric artists in favor of reflection on social mediations made visible in the artworks or artifacts themselves—was too incommensurable with prevailing habits of thought to persuade an audience without providing a series of experiences that might gradually shift the reader's perspective. The metaphor of prisms—that white light is an emergent property of the intertwining of its constituent wavelengths—is one of bringing a view into the light by splitting the illumination itself into a spectrum of distinct rays.

The view of collaboration that is expressed in this book itself emerged gradually, in a manner similar to the way that *Prisms* divulged its theories, as I intuitively pursued an inquiry into groupware design, communication analysis, and social philosophy. While I have made some connections explicit, I also hope that the central meanings will emerge for each reader through his or her own interpretive interests.

In keeping with hermeneutic principles, I do not believe that my understanding of the connotations and interconnections of this text is an ultimate one; certainly, it is not a complete one, the only valid one, or the one most relevant to a particular reader. To publish is to contribute to a larger discourse, to expose one's words to unanticipated viewpoints. Words are always open to different interpretations.

The chronology of the studies has generally been roughly maintained within each of the book's parts, for they document a path of discovery, with earlier essays anticipating what was later elaborated. The goal in assembling this collection has been to provide readers with an intellectual experience open-ended enough that they can collaborate in making sense of the enterprise as a whole—to open up “an arena of intellectual experience” without distorting or excessively delimiting it so that it can be shared and interpreted from diverse perspectives.

The essays were written from my own particular and evolving perspective. They are linguistic artifacts that were central to the intellectual development of that perspective and should be read as situated within that gradually developing interpretation. It may help the reader to understand this book if some of the small groups that incubated its ideas are named.

Collaborating with Groups

Although most of the original papers were published under my name, they are without exception collaborative products, artifacts of academic group cognition. Acknowledgments in the notes section at the end of the book indicate the most immediate intellectual debts. Due to collaboration technologies like the Web and e-mail, our ideas are ineluctably the result of global knowledge building. Considered individually, there is little in the way of software features, research methodology, or theoretical concept that is completely original here. Rather, available ideas have been assembled as tools or intellectual resources for making sense of collaboration as a process of constituting group knowing. If anything is original, it is the mix and the twist of perspectives. Rather than wanting to claim that any particular insight or concept in this book is absolutely new, I would like to think that I have pushed rather hard on some of the ideas that are important to CSCL and brought unique considerations to bear. In knowledge building, the configuration of existing ideas and the intermingling of a spectrum of perspectives on those ideas count.

In particular, the ideas presented here have been developed through the work of certain knowledge-building groups or communities:

- The very notion of knowledge-building communities was proposed by Scardamalia and Bereiter and the Computer-Supported International Learning

Environment (CSILE) research group in Toronto. They pioneered CSCL, working on pedagogical theory, system design, and evaluation of computer-supported classroom practices.

- They cited the work of Lave and Wenger on situated learning, a distillation of ideas brewing in an active intellectual community in the San Francisco Bay area that had a formative impact on CSCW in the 1970s.
- The sociocultural theory elaborated there, in turn, had its roots in Vygotsky and his circle, which rose out of the Russian revolution. The activity theory that grew out of that group's thinking still exerts important influences in the CSCW and CSCL communities.

The personal experience behind this book is perhaps most strongly associated with:

- McCall, Fischer, and the Center for LifeLong Learning and Design in Colorado, where I studied, collaborated, and worked on Hermes and CIE in the early 1990s (see chapters 4 and 5);
- The Computers and Society research group led by Herrmann at the University of Dortmund (now at Bochum), which collaborated on WebGuide and negotiation support (chapters 6 and 9);
- Owen Research, Inc., where TCA and the Crew software for NASA were developed (chapters 1 and 3);
- The Institute for Cognitive Science at Boulder, where State the Essence was created (chapter 2);
- The Innovative Technology for Collaborative Learning and Knowledge Building (ITCOLE) project in the European Union (2001–2002), in which I designed BSCL and participated as a visiting scientist in the CSCW group at Fraunhofer-FIT (chapters 7 and 8);
- The research community surrounding the conferences on computer support for collaborative learning, where I was program chair in 2002 (chapter 11); and
- The Virtual Math Teams (VMT) project that colleagues and I launched at Drexel University in 2003 (chapter 21).

But today knowledge building is a global enterprise, and most of the foundational concepts—like knowledge, learning, and meaning—have been forged in the millennia-long discourse of Western philosophy, whose history is reviewed periodically in the following chapters.

Technology as Mediation

When I launched into software development with a fresh degree in artificial intelligence, I worked eagerly at building cognitive aids—if not directly machine cognition—into my systems, developing rather complicated algorithms using search procedures, semantic representations, case-based reasoning, fuzzy logic, and an involved system of hypermedia perspectives. These mechanisms were generally intended to enhance the cognitive abilities of individual system users. When I struggled to get my students to use some of these systems for their work in class, I became increasingly aware of the many barriers to the adoption of such software. In reflecting on this, I began to conceptualize my systems as artifacts that mediated the work of users. It became clear that the hard part of software design was dealing with its social aspects. I switched my emphasis to creating software that would promote group interaction by providing a useful medium for interaction. This led me to study collaboration itself and to view knowledge building as a group effort.

As I became more interested in software as mediator, I organized a seminar with colleagues and graduate students from different fields on computer mediation of collaborative learning. I used the software discussed in chapter 6 and began the analysis of the moment of collaboration that over the years evolved into chapter 12. We tried to deconstruct the term *mediation*, as used in CSCL, by uncovering the history of the term's effects that are sedimented in the word's usage today. We started with its contemporary use in Jean Lave and Etienne Wenger's *Situated Learning* (1991, p. 50): "Briefly, a theory of social practice emphasizes the relational interdependency of agent and world, activity, meaning, cognition, learning and knowing. . . . Knowledge of the socially constituted world is socially mediated and open ended."

This theory of social practice can be traced back to Lev Vygotsky. Vygotsky described what is distinctive to human cognition, psychological processes that are not simply biological abilities, as *mediated cognition*. He analyzed how both signs (words, gestures) and tools (instruments) act as artifacts that mediate human thought and behavior—and he left the way open for other forms of mediation: "A host of other mediated activities might be named; cognitive activity is not limited to the use of tools or signs" (Vygotsky, 1978, p. 55).

Vygotsky attributes the concept of indirect or mediated activity to Hegel and Marx. Where Hegel loved to analyze how two phenomena constitute each other dialectically—such as the master and slave, each of whose identity arises through their relationship to each other—Marx always showed how the relationships arose in concrete socioeconomic history, such as the rise of conflict between the capitalist class and the working class with the establishment of commodity exchange and

wage labor. The minds, identities, and social relations of individuals are mediated and formed by the primary factors of the contexts in which they are situated.

In this book, mediation plays a central role in group cognition, taken as an emergent phenomenon of small-group collaboration. The computer support of collaboration is analyzed as a mediating technology whose design and use forms and transforms the nature of the interactions and their products.

Mediation is a complex and unfamiliar term. In popular and legal usage, it might refer to the intervention of a third party to resolve a dispute between two people. In philosophy, it is related to *media*, *middle*, and *intermediate*. So in CSCL or CSCW, we can say that a software environment provides a *medium* for collaboration or that it plays an *intermediate* role in the *midst* of the collaborators. The contact between the collaborators is not direct or *im-mediate* but is *mediated* by the software. Recognizing that when human interaction takes place through a technological medium the technical characteristics influence—or *mediate*—the nature of the interaction, we can inquire into the effects of various media on collaboration. For a given task, for instance, should people use a text-based, asynchronous medium? How does this choice both facilitate and constrain their interaction? If the software intervenes between collaborating people, how should it represent them to each other to promote social bonding and understanding of each other's work?

The classic analyses of mediation will reappear in the theoretical part of the book. The term *mediation*—perhaps even more than other key terms in this book—takes on a variety of interrelated meanings and roles. These emerge gradually as the book unfolds; they are both refined and enriched—mediated—by relations with other technical terms. The point for now is to start to think of group-collaboration software as artifacts that mediate the cognition of their individual users and support the group cognition of their user community.

Mediation by Small Groups

Small groups are the engines of knowledge building. The knowing that groups build up in manifold forms is what becomes internalized by their members as individual learning and externalized in their communities as certifiable knowledge. At least, that is a central premise of this book.

The last several chapters of this book take various approaches to exploring the concept of group cognition because this concept involves such a difficult, counter-intuitive way of thinking for many people. This is because cognition is often assumed to be associated with psychological processes contained in individual minds.

The usual story, at least in Western culture of the past three hundred years, goes something like this: an individual experiences reality through his senses (the para-



Figure 1

Auguste Rodin, *The Thinker*, 1881, bronze, Rodin Museum, Philadelphia, PA. Photo: G. Stahl, 2004.

digmatic rational thinker in this tradition is often assumed to be male). He thinks about his experience in his mind; *cognition*, stemming from the Latin *cogito* for “I think,” refers to mental activities that take place in the individual thinker’s head (see figure 1). He may articulate a mental thought by putting it into language, stating it as a linguistic proposition whose truth value is a function of the proposition’s correspondence with a state of affairs in the world. Language, in this view, is a medium for transferring meanings from one mind to another by representing reality. The

recipient of a stated proposition understands its meaning based on his own sense experience as well as his rather unproblematic understanding of the meanings of language.

The story based on the mediation of group cognition is rather different: here, language is an infinitely generative system of symbolic artifacts that encapsulate and embody the cultural experiences of a community. Language is a social product of the interaction of groups—not primarily of individuals—acting in the world in culturally mediated ways. Individuals who are socialized into the community learn to speak and understand language as part of their learning to participate in that community. In the process, they internalize the use of language as silent self-talk, internal dialogue, rehearsed talk, narratives of rational accountability, senses of morality, conflicted dream lives, habits, personal identities, and their tacit background knowledge largely preserved in language understanding. In this story, cognition initially takes place primarily in group processes of interpersonal interaction, which include parent and child, friend and friend, husband and wife, teacher and student, boss and employee, extended family, social network, gang, tribe, neighborhood, community of practice, and so on. The products of cognition exist in discourse, symbolic representations, meaningful gestures, patterns of behavior; they persist in texts and other inscriptions, in physical artifacts, in cultural standards, and in the memories of individual minds. Individual cognition emerges as a secondary effect, although it later seems to acquire a dominant role in our introspective narratives.

Most people have trouble accepting the group-based story at first and viewing collaborative phenomena in these terms. Therefore, the group emphasis emerges gradually in this book rather than being assumed from the start. Indeed, that is what happened during my decade-long inquiry that is documented in these studies.

Although one can see many examples of the decisive role of small groups in the CSCW and CSCL literature, their pivotal function is rarely explicitly acknowledged and reflected on. For instance, the two prevailing paradigms of learning in CSCL—which are referred to in chapter 17 as the acquisition metaphor and the participation metaphor—focus on the individual and the community, respectively, not on the intermediate small group. In the former paradigm, learning consists in the acquisition of knowledge by an individual; for instance, a student acquires facts from a teacher's lesson. In the later, learning consists in knowledgeable participation in a community of practice; for instance, an apprentice becomes a more skilled practitioner of a trade. But if one looks closely at the examples typically given to illustrate each paradigm, one sees that there is usually a small group at work in the specific learning situation. In a healthy classroom there are likely to be cliques of students learning together in subtle ways, even if the lesson is not organized as collaborative learning with formal group work. Their group practices may or may not

be structured in ways that support individual participants to learn as the group builds knowledge. In apprenticeship training, a master is likely to work with a few apprentices, and they work together in various ways as a small group; it is not as though all the apprentice tailors or carpenters or architects in a city are being trained together. The community of practice functions through an effective division into small working groups.

Some theories, like activity theory, insist on viewing learning at both the individual and the community levels. Although their examples again typically feature small groups, the general theory highlights the individual and the large community but has no theoretical representation of the critical small groups in which the individuals carry on their concrete interactions and into which the community is hierarchically structured (see chapter 21).

My own experiences during the studies reported here and in my apprenticeships in philosophy and computer science that preceded them impressed on me the importance of working groups, reading circles, and informal professional discussion occasions for the genesis of new ideas and insights. The same can be seen on a world-historical scale. Quantum jumps in human knowledge building emerge from centers of group interaction: the Bauhaus designers at Weimar, the postimpressionist artists in Paris salons, the Vienna Circle, the Frankfurt School—in the past, these communities were necessarily geographic locations where people could come together in small groups at the same time and place.

The obvious question once we recognize the catalytic role of small groups in knowledge building is whether we can design computer-supported environments to create effective groups across time and space. Based on my experiences, documented in part I, I came to the conclusion that to achieve this goal we need a degree of understanding of small-group cognition that does not currently exist. To design effective media, we need to develop a theory of mediated collaboration through a design-based research agenda of analysis of small-group cognition. Most theories of knowledge building in working and learning have focused primarily on the two extreme scales: the individual unit of analysis as the acquirer of knowledge and the community unit of analysis as the context within which participation takes place. We now need to focus on the intermediate scale: the small-group unit of analysis as the discourse in which knowledge actually emerges.

The size of groups can vary enormously. This book tends to focus on small groups of a few people (say, three to five) meeting for short periods. Given the seeming importance of this scale, it is surprising how little research on computer-supported collaboration has focused *methodologically* on units of this size. Traditional approaches to learning—even to collaborative learning in small groups—measure effects on individuals. More recent writings talk about whole communities of

practice. Most of the relatively few studies of collaboration that do talk of groups look at dyads, where interactions are easier to describe but qualitatively different from those in somewhat larger groups. Even in triads, interactions are more complex, and it is less tempting to attribute emergent ideas to individual members than in dyads.

The emphasis on the group as unit of analysis is definitive of this book. It is not just a matter of claiming that it is time to focus software development on groupware. It is also a methodological rejection of individualism as a focus of empirical analysis and cognitive theory. The book argues that software should support cooperative work and collaborative learning, should be assessed at the group level, and should be designed to foster group cognition.

This book provides different perspectives on the concept of group cognition, but the concept of group cognition *as discourse* is not fully or systematically worked out in detail. Neither are the complex layers of mediation presented, by which interactions at the small-group unit of analysis mediate between individuals and social structures. This is because it is premature to attempt this; much empirical analysis is needed first. The conclusions of this book simply try to prepare the way for future studies of group cognition.

The Promise of Collaborating with Technology

Online workgroups are becoming increasingly popular, freeing learners and workers from the traditional constraints of time and place for schooling and employment. Commercial software offers basic mechanisms and media to support collaboration. However, we are still far from understanding how to work with technology to support collaboration in practice. Having borrowed technologies, research methodologies, and theories from allied fields, it may now be time for the sciences of collaboration to forge their own tools and approaches, honed to the specifics of the field.

This book tries to explore how to create a science of collaboration support grounded in a fine-grained understanding of how people act, work, learn, and think together. It approaches this by focusing the discussion of software design, interaction analysis, and conceptual frameworks on central, paradigmatic phenomena of small-group collaboration, such as multiple interpretive perspectives, intersubjective meaning making, and knowledge building at the group unit of analysis.

The view of group cognition that emerges from the following essays is one worth working hard to support with technology. Group cognition is presented in stronger terms than previous descriptions of distributed cognition. Here it is argued that high-level thinking and other cognitive activities take place in group discourse and that

these are most appropriately analyzed at the small-group unit of analysis. The focus on mediation of group cognition is presented more explicitly than elsewhere, suggesting implications for theory, methodology, design, and future research generally.

Technology in social contexts can take many paths of development in the near future. Globally networked computers provide a promise of a future of worldwide collaboration founded on small-group interactions. Reaching such a future will require overcoming the ideologies of individualism in system design, empirical methodology, and collaboration theory, as well as in everyday practice.

This is a tall order. Today, many people react against the ideals of collaboration and the concept of group cognition based on unfortunate personal experiences, the inadequacies of current technologies, and deeply ingrained senses of competition. Although so much working, learning, and knowledge building takes place through teamwork these days, goals, conceptualizations, and reward structures are still oriented toward individual achievement. Collaboration is often feared as something that might detract from individual accomplishments, rather than valued as something that could facilitate a variety of positive outcomes for everyone. The specter of group-think—where crowd mentality overwhelms individual rationality—is used as an argument against collaboration rather than as a motivation for understanding better how to support healthy collaboration.

We need to continue designing software functionality and associated social practices; continue analyzing the social and cognitive processes that take place during successful collaboration; and continue theorizing about the nature of collaborative learning, working, and acting with technology. The studies in this book are attempts to do just that. They are not intended to provide final answers or to define recipes for designing software or conducting research. They do not claim to confirm the hypotheses, propose the theories, or formulate the methodologies they call for. Rather, they aim to open up a suggestive view of these bewildering realms of inquiry. I hope that by stimulating group efforts to investigate proposed approaches to design, analysis, and theory, they can contribute in some modest measure to our future success in understanding, supporting, and engaging in effective group cognition.

