



CHANGING MINDS

Computers, Learning, and Literacy

ANDREA A. DI SESSA

Changing Minds

Changing Minds

Computers, Learning, and Literacy

Andrea A. diSessa

A Bradford Book
The MIT Press
Cambridge, Massachusetts
London, England

© 2000 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.

This book was set in Sabon by Achorn Graphic Services, Inc. and was printed and bound in the United States of America.

Library of Congress Cataloging-in-Publication Data

diSessa, Andrea A.

Changing minds: computers, learning, and literacy / Andrea A. diSessa.

p. cm.

Includes bibliographical references and index.

ISBN 0-262-04180-4 (hc: alk. paper)

1. Education—Data processing. 2. Learning, Psychology of. 3. Literacy. I. Title.

LB1028.43.D57 2000

370'.285 21—dc21

99-040215

To Melinda

Contents

Preface	ix
Acknowledgments	xvii
1 Computational Media and New Literacies—The Very Idea	1
2 How It Might Be	29
3 Snapshots: A Day in the Life	45
4 Foundations of Knowledge and Learning	65
5 Intuition and Activity Elaborated	89
6 Explaining Things, Explainable Things	109
7 Designing Computer Systems for People	131
8 More Snapshots: Kids Are Smart	165
9 Stepping Back, Looking Forward	209
Notes and Resources	249
Index	267

Preface

Can computers convey to humans a new increment of intellectual power that rivals what conventional literacy has given us? Can education—science education in particular—be transformed by the computer’s presence so that children learn much more, learn it earlier and more easily, and, fundamentally, learn it with a pleasure and commitment that only a privileged few now feel toward school learning? For me, these questions unfold a scope and depth that is frightening and wonderful. The questions are frightening because, in order to answer them with confidence, we need to understand much more than we do today about learning, about how materials (such as text and computers) can make us smarter and help us learn, and about how cultures appropriate a competence base such as a literacy. The questions are wonderful for almost the same reason: trying to answer them brings us to profound and almost timeless questions about people and how they may think and act. The questions are wonderful also because, to my best estimation, both should be answered in the affirmative. Computers *can* make us smarter, if not wiser, and *can* revolutionize education. If this proves correct, civilization has some truly excellent days ahead of it—provided we recognize the possibilities and act in accordance with them.

I can’t answer these questions definitively, but I can show why I believe the best data and theory available suggest immense promise for computers and human potential. Twenty-five years of work have only increased my commitment to the view that computers can be the basis of a new literacy and that great improvements in science education can follow. Part of my commitment comes from practical and theoretical successes at understanding the issues and possibilities; part of it comes from a better

understanding of critics' reasoning, including why it is they think the way they do! Given the promise, I believe that it is irresponsible not to pursue computational literacies, even if we are still uncertain of the outcome, and I would be personally irresponsible if I did not trouble to put the case forward.

This book is not likely to be an easy read, despite the fact that I worked hard to make it accessible to anyone with a good high school science background or, at minimum, to people who know enough to enter responsibly into discussions of science education. Part of the difficulty is the depth of the issues we must face. There are no easy answers to questions such as "How might people think?"

Another part of the difficulty is breadth. Understanding a literacy involves understanding thinking, understanding how material resources can help us think better, and understanding what it means socially for a culture to become literate. I expect almost all readers will be differentially interested in and prepared to engage this breadth, so I have been careful to provide clues and guideposts for readers to pick and choose what suits them. Despite the challenge of breadth, one of the most exciting things about writing this book has been the chance to put together in one place the many aspects of my own work that bear on these issues. I would be disappointed if no one read it all.

Finally, this book will be difficult because many of the ideas cut against the grain of contemporary thought about learning, about education, and about what good computer software looks like. Coming to understand the depth and influence of this misalignment has been exciting and enlightening, if not always encouraging.

Allow me to forestall a few misunderstandings about the intent of this book. No matter what I do or say, many readers will interpret this book as being about a particular piece of software—Boxer, the system my research group and I are using to experiment with computational literacies. I have been as careful as I can manage to frame the issues at the appropriate level of generality. The prime foci of this book are computational media and new literacies, not Boxer per se. On the other hand, concreteness demands I explain how we worked with students and what they did. Furthermore, Boxer is almost certainly the best-developed exemplar of software intended to support a literacy model of the use of computers

in education. Despite the fact that I would love to escape the crush of responsibility for developing and implementing a complex software system, we have had no reasonable alternative in order to undertake the experiments that we felt were needed. In addition, it would be foolish not to show some of the cleverness, I hope, we've installed in Boxer to others who might choose to pursue similar work.

Critics complain that I am unrelentingly positive in my portrayals of students' work. They want me also to explain failures and disappointments, claiming only that will convey a sense of honesty and impartiality. There are practical, scientific, and "political" reasons why I don't include many examples of students' or teachers' or our own failures. The practical reasons are simplest. I have underscored the complexity of the arguments in this book, even as things stand. If I were to qualify every statement and explain the basis and limits of every judgment, the book would lose any semblance of accessibility.

Scientifically, I have tried hard to maintain a sense of balance. My disclaimers of uncertainty already (appropriately) put this book far outside the regime of unqualified confidence that pervades general public discussion of computers and education. If I want to express my convictions and make any contribution at all to public discussion, I cannot lock these ideas in the intricate chambers of pure scientific discourse. In addition, many of our project's failures and difficulties are boring. We could not afford sufficiently fast machines and the programming necessary for a truly smooth and professional interface in our early experiments. How illuminating is it to know this caused difficulty? Many others have demonstrated beyond question that teachers need time and support to teach in new ways. Our specific difficulties in this regard are, for the most part, completely unremarkable. Finally, I have reserved my most somber and uncertain tones for those subjects that I feel most deserve them—social and cultural change, not the potential of students and teachers.

Educational research entails positioning oneself with respect to contemporary landmark points of view. One of the fault lines in current research is how we view teachers and students. On one side, teachers need strict accountabilities, and students won't ever learn unless we discipline their unruly thought patterns and unacademic interests. I have chosen to remind people, as often as necessary, that many if not most of the faults

attributed to teachers and students lie in the inability or unwillingness of scientists or citizens who decide public priorities to support them sufficiently. I want to show not that children “naturally” learn anything put in front of them or that all teachers are brilliant, but that there are solid scientific and empirical reasons to believe that much more can be achieved in learning than most expect. Against the sea of demonstrated failure, a few successes deserve special attention for what they show about the directions to pursue for improvement.

Guide to the Book

Chapter 1 begins by introducing the core claim of the book: computers can be the basis for an empowering new literacy (or family of literacies); hence, they can change the way people think and learn. I also introduce three fundamental views on literacy that organize discussion in the remainder of the book. The cognitive view asks the question, “How can we see the advantages that competence with an external representational system—such as text, algebra, graphing, or computer-implemented systems—conveys to individuals?” The material view of literacy puts the emphasis on the external forms of literacy, what particular characteristics they have, and how they combine with our thoughts and actions to create particular new capabilities. The social view of literacy inquires how literate communities surpass the accomplishment allotted to individuals, no matter how brilliant. In addition, the social view holds the key to understanding whether, when, and how computational literacy might actually take hold. The subtle concept of *social niches* is introduced to explore and explain the viability of some (but not other) forms of literacy and to support an appropriately anarchic view of what good literacy does for people. I use examples from everyday life and from the history of science and education to illustrate these ideas.

Chapter 2 introduces a concrete and realistic image of how new computational representations can change the landscape of learning important scientific ideas. I argue that programming is incredibly apt for “speaking” about the mathematics and physics of motion. In addition, programming can make learning about motion far more motivating than any other available means. The chapter ends by speculating on the properties and

plausibility of changing schools into tool-rich communities, where teachers do far more than consume what others feel is best for them to use in the way of educational software.

Chapter 3 is a set of accounts of children and teachers who have experimented with us concerning learning with computational media. The intent is to enliven the abstractions of previous chapters with real-world happenings that illustrate important ideas. I hope to give readers a feeling for the data that convinces me that radical improvements in learning are possible.

Chapters 4 and 5 are the cognitive and philosophical core of the book. They introduce the learning theory that explains why computers can be such powerful catalysts of change in education. I explain how intuitive knowledge—a much ignored and maligned component of human competence—is actually the platform on which students build scientific understanding. Computer-based representations support and cultivate this important kind of knowledge especially well.

In addition, chapters 4 and 5 discuss how critical it is to view human beings as engaged actors in the world and not simply repositories of instrumental knowledge. People choose to do some things and not others. They have a sense of self that entails ownership over some activities, but alienation and disconnection from others. I introduce a sketch of a future theory of activities that can help us understand and design more human trajectories of competence. I introduce the goal of *committed learning*, where individuals feel a personal connection to the form of extended action in which they participate.

Chapters 6 and 7 are the most technical in the book. They deal with the material forms for possible new literacies. Most particularly, I undertake to explain how to design computer systems that are both comprehensible and powerful. To this end, I introduce two views of computer systems. The *structural view* explains the logic of the system in its own terms. Structurally simple and powerful systems afford long trajectories of learning and use. On the other hand, a *functional view* of systems explains how they can feel familiar and support fluent everyday activity. The functional view connects computer systems most directly to what people care about. I use imaginary designers of incomprehensible bicycles

and the foibles of real calculators to explore the intricate interweaving of function and structure that is the hallmark of the best design.

Chapters 6 and 7 contain two somewhat hidden elements. First, I argue for—and try to show how it is possible to create—two-way literacies, where everyone is a creator as well as a consumer of new material forms. Reading without writing would be an absurdly limited form of literacy in modern society, but far too many people ignore the equivalent of writing for computers. In addition, chapter 7 contains the only technical description of the Boxer software in the book (“A Structure for Cyberspace”). Even readers not interested in principles of design, which Boxer is used to illustrate, will benefit from an up close look at a computational medium. Understanding Boxer will help in understanding the examples given elsewhere in the book, and in understanding their general implications. The description of Boxer also explains how the process of developing learning materials may change with the advent of two-way computational media. In contrast to the dominant paradigm of top-down development by experts, the image of more organic, extended, and open forms of educational materials development, including teachers and students, is a distinctive emblem of what we have been trying to achieve with Boxer.

Chapter 8 returns to accounts of real children and teachers in action. This time, however, we have the benefit of a lot of theoretical preparation for seeing the meaning of new and different things that are made possible by computational media. In particular, I introduce the exotic but wonderful idea that computers can support not just a new fixed form (such as text) for new intelligence, but constant innovation in representational forms, even by teachers and students. If I hadn’t seen it with my own eyes, I wouldn’t believe it myself! The chapter again highlights the role of the computer in engaging intuitive knowledge and how students can reach surprising levels of accomplishment in programming because of the characteristics of computational media.

Chapter 9 complements the focus on cognitive issues in chapters 4 and 5, and the focus on issues of material form in chapters 6 and 7. We return to social and cultural issues to examine in some detail the resonance or antiresonance of computational literacies with current common sense about computer systems and how people learn. As part of this study, I

look at the World Wide Web with respect to how it builds toward the hopes for improved learning presented in this book. Then, I look at the intuitions and implicit assumptions about computers and learning of several important communities, including educators, technologists, and research funders. This study leads to a renewed appreciation for the scale and difficulty of change in the direction of the most optimistic hopes expressed in this book. If the possible rewards of new literacies are opulent, the challenges are commensurate.

Highlights for Different Audiences

A Minimal Read

For those who want a sense of this book's ideas with minimal effort, I suggest the following sections.

Chapter 1: Introduction; A Cognitive View of Material Intelligence (optional)

Chapter 2: Beyond Algebra

Chapter 3: Fish and Water; Ownership by Individuals and Communities; Materials That Live and Grow (optional)

General Audience

The general audience interested in education will need the outline of the book's project in chapter 1, although the theoretical ideas concerning literacy may be less important. The examples of student and teacher work in chapter 3 and, to a lesser extent, the more exotic examples in chapter 8 should be important to this audience. The discussion of learning in chapters 4 and 5 should be accessible to the more dedicated reader; probably the most important new ideas are introduced in "The Structure of Activities and 'Committed Learning'" in chapter 4 and elaborated in the second half of chapter 5. Although chapters 6 and 7 introduce some technical issues, the discussion of Boxer in chapter 7 ("A Structure for Cyberspace") should be accessible and helpful to all.

Technology Designers

The "materials" chapters will be of most interest. Ideas concerning structural and functional understanding of systems (chapters 6 and 7) and concerning alternative social patterns of materials development (the end

of chapter 2 and last third of chapter 7) should be relevant. The discussions of children's and teachers' programming in the chapters describing student and teacher work (chapters 3 and 8) may be of interest.

Theorists of Literacy

The discussion of literacy in chapter 1 should be relevant. In particular, putting social viability in a central position with respect to defining literacy may be a surprising move. The idea of social niches inevitably leads to an expectation of diversity and multifaceted influence for literacy, which is quite different from most general accounts. Discussions of technical and scientific subliteracies, such as algebra and programming, might be of interest to contrast with more conventionally narrow definitions of literacy having to do exclusively with written language. The discussion of meta-representational competence and how computational literacies may fundamentally alter which people create new representational forms should resonate with some accounts of the power of literacy and contrast with others.

Socioculturally Oriented Readers

This group will find chapters 1 and 9 most relevant. "Activity theorists" may be interested to consider the approach taken to understanding activities in chapters 5 and 6. My general approach to understanding the power and influence of artifacts (scattered throughout the book) may be of interest.

Students of Learning and Conceptual Change

Chapters 4 and 5 are most relevant. The discussion of intuitive knowledge may be old hat, but the move toward enfolding activity into epistemology should be provocative.

Acknowledgments

Proper acknowledgments for the work reported in this book—which included a huge software project and the difficult pragmatics of doing truly unusual things in schools—probably deserves another book. At the risk of slighting many contributions, I'll attempt a miniature version. My grateful thanks in no way implies funders' or individuals' endorsement of the ideas in this book.

The most important proximal cause of this book was a year as a Fellow at the Center for Advanced Study in the Behavioral Sciences, supported by a fellowship from the Spencer Foundation (grant no. 19940032). The Spencer Foundation seems to have provided support for my work at precisely the most critical times. In particular, it also supported my work on intuitive knowledge (grant no. B1393) reported in chapters 4 and 5.

The development of Unix Boxer and most of the learning experiments reported here were supported by a pair of grants from the National Science Foundation (NSF), from a program directed by Dr. Andrew Molnar (grant nos. MDR 86-42177 and MDR 88-50363). A more recent grant allowed us to investigate students' metarepresentational competence (grant no. RED 95-53902), reported in chapter 8.

Some additional support for the development of Boxer came in the way of generous equipment donations from Sun Microsystems and a grant from Apple Computer that allowed us to hire a computer science student. The Graduate School of Education of the University of California, Berkeley, provided some direct support for my project and also indirect support in the way of computational infrastructure.

Seymour Papert got me started in computers and education, and inspired many of the early ideas that eventually led to this book. Alan Kay's

work helped motivate the move to a serious consideration of a literacy perspective.

Ed Lay has been the chief designer and implementor of Boxer for nearly fifteen years. Boxer was conceptualized under the auspices of the MIT Laboratory for Computer Science under a grant from the Defense Advanced Research Projects Agency (DARPA), Michael Dertouzos, principal investigator. Hal Abelson codirected that work and, critically, sketched the single most difficult algorithm, which made Boxer efficient enough to be practical. Several talented programmers made contributions: Gregor Kiczalis wrote the first version of the editor; Leigh Klotz wrote the core of the Boxer interpreter; Jeremy Roschelle wrote the first real Boxer graphics system; and Michael Travers, introduced to us by Alan Kay, did the initial port of Boxer to the Macintosh platform.

The two teachers who have contributed most to the project are Henri Picciotto and Tina Kolpokowski. Henri has defined “teacher-leader” for me in terms of innovation with computational media. His constant feedback has made Boxer much better. Tina guided our first Boxer course, and her courage and brilliant teaching have continued to inspire.

Several people were instrumental in our school trials. Don Ploger provided many kinds of help and innovation during our two initial NSF grants, and he is still showing how to bring computational media into the real world. Many graduate students made important intellectual and practical contributions. Those whose influence has shown most in this book are Bruce Sherin and David Hammer. I can’t omit all the young students whom we tried to help learn and who, in return, taught us so much.

The most helpful feedback on the text of this book was provided by Michael Eisenberg, Richard Noss, and Melinda diSessa. Graduate student members of my research group (including Rafael Granados, Rodrigo Madanes, and Andy Elby) influenced content and style. Two anonymous reviewers of the manuscript also had a significant influence. Kathleen Much of the Center for Advanced Study in the Behavioral Sciences helped improve the writing of many of the chapters. Maurice Anker developed the software extension that created the high-resolution Boxer figures used in this book.

Critics of the ideas presented in this book made important contributions, not all of which were intended. You know who you are.

Although I list them last, my family is first in my heart. Melinda provided advice, help, support, and encouragement, without which this book would have been impossible. My sons, Kurt and Nicholas, are two of those “smart kids” from whom I have learned so much about thought, action, and human potential. I love them, too.

Changing Minds

Computational Media and New Literacies—The Very Idea

Literacy in the conventional sense of being able to read and write is both highly valued and commonplace in contemporary society. Although almost everything else—especially values—seems to be in dispute, no one questions the importance of reading and writing as foundational skills. Of course, there is plenty of disagreement about exactly what constitutes literacy and how we should go about bringing up children to become literate. Still, not even the most extremist politicians can expect to win converts by cheering the latest study that shows college students can neither string two sentences together coherently nor read a map.

Because the social value of literacy is so important to this book, it is worth taking a few moments to evoke a more lively sense of the multiple roles literacy plays in our lives. Everyday life is a good place to start. When I get up in the morning, I usually find time to look at the newspaper. I glance through international events, partly just to keep up, partly because I have a few special interests stemming from overseas friends and personal associations from travel. I am not very fond of national politics, but it is interesting to see who is trying to do away with the U.S. Department of Education this year and whether National Science Foundation funding for social sciences will really go away.

I usually look in the business section mainly because that is the most likely place to find technology news, but also because I hope to find useful information that will help me save for retirement and pay for my sons' college education. Sometimes I'll find a good recipe and other times a piece of medical or health information of use to my family.

My interests in newspaper news are partly personal, organized by my own orientations and multiple group memberships, and partly profes-

sional. I keep up with some aspects of my work that don't get covered in professional journals (such as what features one gets in an inexpensive home computer these days), and I "accidentally" become a better-informed citizen and voter. For all of this, I lead a bit richer, probably slightly better, and more meaningful life. Many people buy newspapers, and I'm sure they have similar experiences.

Mail time is another bit of everyday life that reminds us how deeply literacy pervades our lives, frequently without our notice: letters from offspring or parents (we'd better write back), magazines, solicitations that every once in a while get noticed and acted on, forms to fill out (taxes!), sometimes with daunting written instructions (taxes!) . . .

Work gives us another perspective on literacy. As an academic, I have a special relation to literacy. It would not be a bad approximation to say my professional life *is* reading and writing. This book, for example, may be the single best representation of at least fifteen years' work on computational media, and it is likely to be only a small percentage of my career writing output. I'm writing now at home in front of a wall of books eight feet high and twenty feet wide; perhaps half of them are professional books. My professional dependence on literacy may be easy to dismiss as atypical in society—and surely it is atypical—but I am not too modest to claim that academia makes significant contributions, particularly in educating the young and in pursuing new knowledge outside of narrow special interests that measure new accomplishments only by dollars or by political or social power. There are many other "niche players" in society for whom literacy is nearly as important as in the lives of academics. Science and high technology are critically literate pursuits. I am certainly glad my personal doctor reads and that some doctors can write well enough to convey new ideas and practices effectively. In a wider scope, business and bureaucracies run on information, reports, memos, spreadsheets, concept papers, and so on.

A third perspective on literacy may be the most obvious and most important. Literacy is infrastructural and absolutely essential to education, to creating people who are knowledgeable and competent. *Infrastructural* means that literacy is not just a result of the educational process, but a driving force within it. Every class has textbooks, not only English class or other overtly literacy-oriented classes. If you can't read well enough

or don't have basic mathematical literacy, you can't profit from history, science, or mathematics textbooks. Education has producers as well as consumers. Teachers, too, read to learn more and improve their practice. Someone has to write textbooks. Most teachers, especially the best, also write to help students—notes, handouts, evaluations—even if they are not writing to and for fellow teachers.

Enter the computer, a “once in several centuries” innovation, as Herbert Simon put it. Computers are incontestably transforming our civilization. Comparisons of our current information revolution to the Industrial Revolution are commonplace and apt. Almost no corner of society is untouched by computers. Most dramatically, science and business are not remotely the same practices they were twenty years ago because of the widespread influence of computers.

Education and schooling are, as yet, an ambiguous case. Few can or should claim that computers have influenced the cultural practices of school the way they have other aspects of society, such as science and business. Just look at texts, tests, and assignments from core subjects. They have changed little so far. Numbers tell a more optimistic but still muted story of penetration. In 1995, K–12 schools in the United States had about three computers per “average” thirty-student classroom. A decent informal benchmark I use is one computer per three students before core practices can be radically changed. This is the ratio at which students can be working full-time, three to a machine, a number that I know from personal experience can work very well, or each student can work alone one-third of the time, well above the threshold for infrastructural influence. One computer per ten students seems some distance from one per three, but consider that schools have been adding regularly to their stock of computers by about one-half computer per classroom per year. At that rate, average schools can easily meet my benchmark in a decade and a half. More than 10 percent of the high schools in the country are *already* above the threshold benchmark.

I fully expect the rate of computer acquisition to accelerate. That one-half computer per classroom is a fraction of what school districts spend per pupil, let alone per classroom, each year. Add the facts that in, say, ten years, computers will be easily ten times more powerful (thirty is a more responsible scientific estimate), that they will cost less, and that

there will be vastly more good learning materials available, and I see inevitability. Despite amazing entrenchment, general conservatism, small budgets, and low status, schools will soon enough be computer-rich communities, unless our society is suicidally reluctant to share the future with its young.

Assuring ourselves that schools will have enough computers to do something interesting is a long way from assuring ourselves that something good—much less the very best we can manage—will happen. That is precisely what this book is about. What is the very best thing that can happen with computer use in education? What might learning actually be like then? How can you assure yourself that any vision is plausible and attainable? What sort of software must be created, and what are the signposts to guide us on the way to realizing “the best”?

I’ve already set the standard and implicitly suggested the key:

Computers can be the technical foundation of a new and dramatically enhanced literacy, which will act in many ways like current literacy and which will have penetration and depth of influence comparable to what we have already experienced in coming to achieve a mass, text-based literacy.

Clearly, I have a lot of explaining to do. This is not a very popular image of what may happen with computers in education. For that matter, it is not a very unpopular image either in the sense of having substantial opposition with objections that are deeply felt or well thought out. Instead, I find that most people have difficulty imagining what a *computational literacy*, as I propose to call it, may mean, or they dismiss it as easy and perhaps as already attained, or they find it immediately implausible, almost a contradiction in terms, so that it warrants little thought.

I need to identify and reject an unfortunate cultural artifact that can easily get in the way of thinking seriously about relevant issues. *Computer literacy* is a term that has been around since the early days of computers. It means something like being able to turn a computer on, insert a CD, and have enough keyboarding and mouse skills to make a few interesting things happen in a few standard applications. Computational literacy is different. In the first instance, the scale of achievement involved in com-

puter literacy is microscopic compared to what I am talking about. It is as if being able to decode, haltingly, a few “typical” words could count as textual literacy.

If a true computational literacy comes to exist, it will be infrastructural in the same way current literacy is in current schools. Students will be learning and using it constantly through their schooling careers and beyond in diverse scientific, humanistic, and expressive pursuits. Outside of schools, a computational literacy will allow civilization to think and do things that will be new to us in the same way that the modern literate society would be almost incomprehensible to preliterate cultures. Clearly, by computational literacy I do not mean a casual familiarity with a machine that computes. In retrospect, I find it remarkable that society has allowed such a shameful debasing of the term *literacy* in its conventional use in connection with computers; perhaps like fish in the ocean, we just don’t see our huge and pervasive dependence on it.

I find that substituting the phrase *material intelligence* for *literacy* is a helpful ploy. People instinctively understand intelligence as essential to our human nature and capacity to achieve. Material intelligence, then, is an addition to “purely mental” intelligence. We can achieve it in the presence of appropriate materials, such as pen and paper, print, or computers. This image is natural if we think of the mind as a remarkable and complex machine, but one that can be enhanced by allowing appropriate external extensions to the mechanism, extensions that wind up improving our abilities to represent the world, to remember and reason about it. The material intelligence—literacy—I am referring to is not artificial intelligence in the sense of placing our own intelligence or knowledge, or some enhanced version of it, into a machine. Instead, it is an intelligence achieved cooperatively with external materials.

In the remainder of this introductory chapter, I have one overarching goal. I want to examine traditional literacy in some detail, including both micro- and macrocomponents. The microfocus shows a little about how traditional literacy actually works in episodes of thinking with a materially enhanced intelligence. The macrofocus introduces some large-scale and irreducibly social considerations that determine whether a new literacy is achievable and how. Much of the rest of the book builds on these views of conventional literacy, extrapolating them to consider what

exactly a computational literacy might mean, what it might accomplish for us, whether it is plausible, and how we can act to bring it about.

Three Pillars of Literacy

Before getting down to details, we might find it useful to set a rough framework for thinking about the many features and aspects of literacy. I think of literacy as built on three foundational pillars. First, there is the material pillar. That is, literacy involves external, materially based signs, symbols, depictions, or representations. This last set of terms, as well as others, holds an essential magic of literacy: we can install some aspects of our thinking in stable, reproducible, manipulable, and transportable physical form. These external forms become in a very real sense part of our thinking, remembering, and communicating. In concert with our minds, they let us act as if we could bring little surrogates of distant, awkwardly scaled (too big or too small), or difficult to “touch” aspects of the real world to our desktop and manipulate them at will. We can read a map, check our finances, write our itinerary, and plan an automobile trip across the United States. Even more, we can create and explore possible worlds of fantasy or reality (as in a scientific exploration) with a richness, complexity, care, and detail far transcending what we may do with the unaided mind.

The material bases for literacy are far from arbitrary, but are organized into intricately structured subsystems with particular rules of operation, basic symbol sets, patterns of combination, conventions, and means of interpretation. These subsystems all have a particular character, power, and reach, and they also have limits in what they allow us to think about. Associated with them are particular modes of mediated thought and connections to other subsystems. Written language, the prototype of literacy, has an alphabet, a lexicon, a grammar, and a syntax, and above these technical levels are conventions of written discourse, genres, and styles, and so on. Written language is expansive in what may be thought through it, it is variable in its level of precision—we can use it carefully or casually, from a jotted note to a formal proof—and it is generally a wonderful complement to other subsystems, for example, as annotation over the graphical-geometric component of maps.

Other subsystems have a different character. Arithmetic, for example, is much narrower in what you may write about with it. You can't write much good poetry or philosophy in numbers. But what it does allow us to think about, it does with great precision. We can draw inferences (calculate) using arithmetic either perfectly or with as much precision as we care to spend time to achieve. The power of arithmetic is tightly connected with other components of human intellect. For example, scientific understanding frequently is what liberates arithmetic as a useful tool; an engineer can calculate how big a beam is needed in a building because we understand scientifically how size, shape, and material relate to strength. Other important mathematical subsystems—algebra, calculus, graph drawing and interpreting, and so on—also have their own character. Each has its own structure, expressive range, associated modes of thought, and “intellectual allies.”

The material pillar of literacy has two immensely important features: the material subsystems of literacy are technologically dependent, and they are designed. It is not at all incidental to contemporary literacy that paper and pencils are cheap, relatively easy to use, and portable. Think back to quills and parchment, or even to cuneiform impressions or rock painting or carving, and consider what you have done today with letters that would have been impossibly awkward without modern, cheap, portable implements. Think what difference the printing press made in creating a widespread, popular, and useful literacy.

Coming directly to the heart of this book, computer technology offers a dazzling range of inscription forms (spreadsheets, electronically processed images and pictures, hypertext, etc.), of reactive and interactive patterns (think of game interfaces—from text typed in and new text returned in reaction, to intense, real-time reflex interaction, to contemplative browsing of a visually based interactive mystery story), of storage and transmission modes (CDs to worldwide networking), and of autonomous actions (simulations, calculation). With all these new forms and more to come, it seems inconceivable our current material literacy basis could remain unaffected.

I noted also that all these inscription forms, both the historical ones and those in current and future development, have been designed—either in acts of inspiration (e.g., the invention of zero or the pulldown menu)

or slowly over generations by an accumulation of little ideas and societal trial and error. We have much to gain by thinking carefully about what the whole game of literacy is and about what we can do with computers that can either hasten or undermine new possibilities.

The second pillar of literacy is mental or cognitive. Clearly the material basis of literacy stands only in conjunction with what we think and do with our minds in the presence of inscriptions. A book is only a poor stepping stool to a nonreader. Material intelligence does not reside in either the mind or the materials alone. Indeed, the coupling of external and internal activity is intricate and critical.

This mutual dependence has both constraining and liberating aspects. Our minds have some characteristics that are fixed by our evolutionary state. Nobody can see and remember a thousand items presented in a flash or draw certain kinds of inferences as quickly and precisely as a computer. On the positive side, our ability to talk and comprehend oral language is at least partly physiologically specific, and without this physical equipment, written literacy would also probably be impossible. Similarly, I believe that new computer literacies will build on and extend humans' impressive spatial and dynamic interactive capabilities far more than conventional literacy does. I have much more to say about these issues later, mainly in chapters 4, 5, and 8.

New computational inscription systems should therefore build on strengths in human mental capacities, and they must also recognize our limitations. Intelligence is a complex and textured thing. We know little enough about it in detail, and we will certainly be surprised by its nature when it is materially enhanced in quite unfamiliar ways. The simultaneous tracking of our understanding of intelligence and knowledge along with materially enhanced versions of them is, for me, among the most scientifically interesting issues of our times. It may be among the most practically relevant issues for the survival and prospering of our civilization.

The third pillar of literacy is social, the basis in community for enhanced literacies. Although one may imagine that an individual could benefit in private from a new or different material intelligence, literacy in the sense investigated in this book is unambiguously and deeply social. Let's take a look at the boundary between the social and the individual to get a feeling for the issues.

Sir Isaac Newton (1642–1727) is generally credited with inventing the calculus as part of building the intellectual infrastructure for his own accomplishments in understanding mechanics, the science of force and motion. His feat was one of those rare but especially impressive events in the history of science when a new material intelligence emerged out of the specific needs of an investigation; that new intelligence clearly contributed to Newton's ability to state and validate his new scientific accomplishments.

Fundamentally, the calculus is a way of writing down and drawing inferences about (i.e., calculating) various aspects of changing quantities. Newton wanted to reason about instantaneous properties of motion that were difficult to capture using prior conceptions and representations. A planet traveling around the sun is constantly changing its speed. Averages and constant speed situations, which were handled adequately by prior techniques, simply weren't up to dealing with facts about instants in a constantly and nonuniformly changing situation. The calculus allowed Newton to capture relations in those instants. Thinking about laws of nature that work in instants and at points in space has turned out to be one of the most fundamental and enduring moves of all time in physics. Nature's causality is local: there is no such thing as "action at a distance" (or "at a later time") in modern physics.

Newton's calculus sounds like a case of a new material intelligence emerging in the hands of an individual, which enabled and in part constituted a fundamental advance for all of science, but the details of the story betray important social components. In the first instance, Newton's accomplishment was clearly not developed on a blank slate. He borrowed and extended techniques, even graphical techniques, that had been around certainly since Galileo (1564–1642), fifty years earlier. (Galileo, in turn, cribbed many of these from his predecessors.) Newton himself said, "If I have seen farther than most, it is because I stood on the shoulders of giants," and this was as true for the calculus as for his laws of physics.

Neither was the development of calculus finished in Newton's work. G. W. Leibniz (1646–1716), most believe, independently developed the calculus at about the same time. Indeed, the notational form mainly in use today is Leibniz's, not Newton's. Although I can't prove it, I believe the reasons for this fact are in important measure pedagogical. Leibniz's

notation is easier to learn; it is powerfully heuristic in suggesting useful techniques and ways of thinking about change; and it even makes obvious certain important theorems. For example, in Leibniz's notation the rate of change of a quantity, x , given a small change in another, t , looks like just what it is, a ratio, dx/dt . (The d in dx and dt stands for a change, or "delta," in the quantity.) Newton's notation is opaque, \dot{x} . In Leibniz's notation, the "change of variable theorem,"

$$\frac{dz}{dy} \frac{dy}{dx} = \frac{dz}{dx}$$

looks obvious, even if it is not. "Cancel the dy s" appears to prove the theorem. Newton is not so helpful. His notation dealt easily only with changes in time, which he called "fluxions," so he pretty much had to state this theorem for the case that y is time, and he had to do it in words that hide the real generality of the theorem. Newton's statement of the theorem used the term *velocity* to describe dz/dy and dx/dy , whereas Leibniz's notation makes change over time only a special case: "The moments [spatial rate of change] of flowing quantities are as [the ratio of] the velocities [time rate of change] of their flowing or increasing."

The morals of Leibniz's contributions are both obvious and subtle. Obviously, once again, science is a strongly cumulative social enterprise. To reach their full potential, contributions must be both shared and extended by others. I would extend this precept beyond the bounds of professional science: incremental material intelligence in the hands of a genius, or even in the hands of a scientific or technological elite, is small in comparison to the huge possibilities of popular new literacies.

The second moral from Leibniz is more subtle, but it explains why I spend two chapters on the material basis of computational media. The inscribed form of thought is critically important. I've suggested that Leibniz has helped generations of scientists and mathematicians in training, even if his purely conceptual accomplishment was entirely redundant with Newton's. I can highlight this claim with a somewhat speculative thought experiment. The fact is that calculus has become absolutely infrastructural in the educational process of scientists, engineers, and a broad range of other technical professions. All learners in these categories are funneled through freshman calculus, if they did not already study calculus

in high school. Further learning is dependent on this prerequisite. Upper division textbooks, for example, assume it in their exposition.

This move to infrastructural status for calculus was not easy. It took more than two centuries! In the twentieth century, a few bold universities decided it was possible and useful to teach calculus in the early and universal (that is, for all technical students) infrastructural mode. It succeeded, more or less, and gradually more schools jumped on the bandwagon. They had the advantage of knowing that teaching calculus this way was possible, and they could capitalize on the know-how of the early innovators. In the meantime, other professors and textbook writers for other classes began to take the teaching of calculus for granted. They became dependent on it. Calculus came to be infrastructural.

Focus on two critical phases. First, suppose calculus was just 10 percent more difficult to learn. Would those early innovators have had the courage to guess it might succeed? Similarly, at the second phase, if 10 percent fewer students “got it,” would these innovators have declared success, and would others have followed and had enough success for the whole project to succeed? Finally, might Leibniz’s notation have made that small difference by which the snowball of calculus got over the crest to start the eventual avalanche of infrastructural adoption?

I am not interested in verifying any particular account of these events. The general principles are clear. The emergence of a material intelligence as a literacy, as infrastructural, depends on complex social forces of innovation, adoption, and interdependence, even if (as I have argued is generally false) it originated with an individual or a small group. Furthermore, under some circumstances at least, small differences in learnability can make huge differences in eventual impact.

Here are the implications of this history of calculus with respect to the broader aims of this book. We may now have sufficiently learnable and powerful computational inscription systems to have dramatic literacy implications. For example, learning some important parts of mathematics and science may be transformed from a pleasurable success for a few but a painful failure for most to an infrastructural assumption for our whole society, and this transformation depends in an essential way on details of material form and on social forces, which it behooves us to understand.

A Cognitive View of Material Intelligence

My goal for this section is to illustrate and explicate some of the details of how material intelligence works to enhance the power of individual human beings. I have chosen to look at a small part of the works of Galileo for several reasons. The first is a version of the invisibility of water to fish. I want to take us a little away from our familiar everyday world of literacy so that some things we otherwise take for granted may stand out.

The second reason to consider Galileo is that doing so illustrates a somewhat technical and scientific component of literacy. Making mathematics and science easier and more interesting to learn was my first motivation for thinking about computers, and it is still my primary concern. I firmly believe computers will also have revolutionary literacy effects in art and the humanities generally, but this book is rich and complex enough dealing with mathematics and science. As a bonus, this little story leads directly into my own experiences in using computers to teach children about motion.

The last reason to look at Galileo and to return to the early part of the seventeenth century is to remind us, by contrast with what exists today, that literacy is created. What we had is not what we have, and without the slightest doubt it is not what we will have. The process of literacy creation happens on the scale of decades, if not centuries, even for some relatively small components of literacy. If we want to think about new literacies—and I think we must, given their importance—we must also free ourselves to think about the coming decades, not just next year.

Let me start this little parable of Galileo and literacy as it first appeared to me—as a puzzle. Just at the beginning of his treatment of motion in *Dialogues Concerning Two New Sciences*, at the outset of what is generally regarded as his greatest accomplishment, Galileo defines uniform motion, motion with a constant speed. The section that follows this definition consists of six theorems about uniform motion and their proofs. Below, I reproduce those theorems. Despite the unfamiliarity of the language, I urge you to try to follow along and think what, in essence, Galileo is getting at in these theorems and how we would express it in modern terms.

THEOREM 1 If a moving particle, carried uniformly at constant speed, traverses two distances, then the time intervals required are to each other in the ratio of these distances.

THEOREM 2 If a moving particle traverses two distances in equal intervals of time, these distances will bear to each other the same ratio as their speeds. And conversely, if the distances are as the speeds, then the times are equal.

THEOREM 3 In the case of unequal speeds, the time intervals required to traverse a given space are to each other inversely as the speeds.

THEOREM 4 If two particles are carried with uniform motion, but each with a different speed, then the distances covered by them during unequal intervals of time bear to each other the compound ratio of the speeds and time intervals.

THEOREM 5 If two particles are moved at a uniform rate, but with unequal speeds, through unequal distances, then the ratio of the time intervals occupied will be the products of the distances by the inverse ratio of the speeds.

THEOREM 6 If two particles are carried at a uniform rate, the ratio of their speeds will be the product of the ratio of the distances traversed by the inverse ratio of the time intervals occupied.

A modern reader (after struggling past the language of ratios and inverse ratios) must surely get the impression that here there is much ado about very little. It seems like a pretentious and grandly overdone set of variations on the theme of “distance equals rate times time.” To make matters worse, the proofs of these theorems given by Galileo are hardly trivial, averaging almost a page of text. The first proof, indeed, is difficult enough that it took me about a half-dozen readings before I understood how it worked. (See the boxed text.)

In fact this *is* a set of variations on distance equals rate times time. Allow me to make this abundantly clear. Each of these theorems is about two motions, so we can write “distance equals rate times time” for each. Subscripts specify which motion the distance (d), rate (r), and time interval (t) belong to.

$$d_1 = r_1 t_1$$

$$d_2 = r_2 t_2$$

In these terms, we can state and prove each of Galileo's theorems. Because Galileo uses ratios, first we divide equals by equals (the left and right sides of the equations above, respectively) and achieve:

$$\frac{d_1}{d_2} = \frac{r_1 t_1}{r_2 t_2}$$

THEOREM 1 In the case $r_1 = r_2$, the r terms cancel, leaving $d_1/d_2 = t_1/t_2$.

THEOREM 2 In the case $t_1 = t_2$, the t terms cancel, leaving $d_1/d_2 = r_1/r_2$. Conversely, if $d_1/d_2 = r_1/r_2$ then $t_1/t_2 = 1$ or $t_1 = t_2$.

THEOREM 3 In the case of $d_1 = d_2$, the d terms cancel, leaving $(r_1/r_2)(t_1/t_2) = 1$, or $t_1/t_2 = r_2/r_1$.

THEOREM 4 This is precisely our little ratio lemma, $d_1/d_2 = (r_1/r_2)(t_1/t_2)$.

THEOREM 5 Solve the equation above for t_1/t_2 ; $t_1/t_2 = (d_1/d_2)(r_2/r_1)$.

THEOREM 6 Solve for r_1/r_2 ; $r_1/r_2 = (d_1/d_2)(t_2/t_1)$.

For direct contrast, I reproduce Galileo's proof of theorem 1, which is *one-sixth* of the job we did with algebra, in box 1.

So now we've redone a significant piece of work by one of the great geniuses of Western science, with amazing ease. Solving problems is always easier after the first time around, but the difference here is almost mindboggling. What we did would constitute only an exercise for a ninth-grade mathematics student.

That, in fact, is the key. Galileo never had ninth-grade mathematics; he didn't know algebra! There is not a single "=" in all of Galileo's writing.


The fault is not with Galileo or with the education provided by his parents or with the schooling of the times. Algebra simply did not exist at that time. To be more precise, although solving for unknowns that participated in given relations with other numbers had been practiced for at least half a millennium, the modern notational system that allows writing equations as we know them—and also the easy manipulations to solve them—did not exist. Fifty years after Galileo's main work, René Descartes (1596–1650) would have a really good start on modern algebra. Later, by the end of the seventeenth century, algebra had stabilized to roughly the modern notation and manipulative practices, although it would be the twentieth century before algebra became a part of widespread technical literacy.

Box 1

If a moving particle, carried uniformly at a constant speed, traverses two distances the time-intervals required are to each other in the ratio of these distances.

Let a particle move uniformly with constant speed through two distances AB, BC, and let the time required to traverse AB be represented by DE; the time required to traverse BC, by EF; then I say that the distance AB is to the distance BC as the time DE is to the time EF.

Let the distances and times be extended on both sides towards G, H and I, K; let AG be divided into any number whatever of spaces each equal to AB, and in like manner lay off in DI exactly the same number of time-intervals each equal to DE. Again lay off in CH any number whatever of distances each equal to BC; and in FK exactly the same number of time-intervals each equal to EF; then will the distance BG and the time EI be equal and arbitrary multiples of the distance BA and the time ED; and likewise the distance HB and the time KE are equal and arbitrary multiples of the distance CB and the time FE.



And since DE is the time required to traverse AB, the whole time EI will be required for the whole distance BG, and when the motion is uniform there will be in EI as many time-intervals each equal to DE as there are distances in BG each equal to BA; and likewise it follows that KE represents the time required to traverse HB.

Since, however, the motion is uniform, it follows that if the distance GB is equal to the distance BH, then must also the time IE be equal to the time EK; and if GB is greater than BH, then also IE will be greater than EK; and if less, less. There are then four quantities, the first AB, the second BC, the third DE, and the fourth EF; the time IE and the distance GB are arbitrary multiples of the first and the third, namely of the distance AB and the time DE.

But it has been proved that *both* of these latter quantities are either equal to, greater than, or less than the time EK and the space BH, which are arbitrary multiples of the second and the fourth. Therefore, the first is to the second, namely the distance AB is to the distance BC, as the third is to the fourth, namely the time DE is to the time EF.

Q.E.D.

(From Galileo, *Dialogues Concerning Two New Sciences*. Translated by H. Crew and A. de Salvio [Northwestern University, 1939], pp. 155–156.)

In net, an average ninth-grade mathematics student plus a particular inscription system yields a material intelligence that surpasses Galileo's intelligence, at least in this domain of writing and "reasoning about" simple quantitative relationships.

We can learn more about the power that material intelligence conveys to individuals by thinking more about this example. Notice first that the equations are shorter, more concise than Galileo's natural language. Compactness has many advantages, besides saving paper. It usually results in statements that are easier to remember. Every mathematically literate person remembers, probably literally and iconically, $d = rt$, and possibly $E = mc^2$. Some even remember the solution to the quadratic equation,

$$x = \frac{-b \pm \sqrt{b^2 - 4ac}}{2a}$$

(I didn't have to look it up!). Galileo's sentences, as well written as they are, are less compact and less memorable. Our memories are better with good external inscriptions, even if we do not use the material form as memory by rereading what we wrote a while ago. Literacies leave traces of themselves in autonomous thinking, making us smarter even when we're not in the presence of the material form.

Inscription systems and associated subliteracies are a little like miniature languages in that they select a certain kind of thing to talk about and certain things to say about them. They have a certain vocabulary, one might say. Thus, each system is apt for some things and less apt for others. Every good new system enlarges the set of ways we can think about the world. If we happen to have in hand a system that is apt for learning or inquiring into a new area, we make progress quickly. If it turns out that a fairly easy inscription system enlightens a new area, then we can teach the inscription system first, and students will learn the area much more easily than those who had to work without or had to invent the system. This is a general version of where we came in: any high school student who knows algebra and Descartes's analytic geometry can learn all of Galileo's accomplishments concerning motion in very short order.

Part of expressing the right things is picking the right level of abstraction. For example, Galileo sometimes talks about two motions of one particle and sometimes about two distinct particles. These details are ir-

relevant, however; the algebraically expressed relations apply to any pair of motions. An even higher level of abstraction than that of equations turns out to be worse than one that is too detailed. To say that distance, rate, and time “are related” misses important, relevant details.

Algebra has been so spectacularly successful at picking a good level of abstraction and displaying the right kind of relations that in some parts of science, one may understandably, but incorrectly, view progress as a march from one equation to the next, from Newton ($F = ma$) to Maxwell

$$(\nabla \times \mathbf{E} = -\frac{1}{c} \frac{\partial \mathbf{B}}{\partial t}, \text{ etc.})$$

to relativity (Einstein’s $E = mc^2$) to quantum mechanics (Schroedinger’s equation,

$$i\hbar \frac{\partial}{\partial t} \Psi = -\frac{\hbar^2}{2m} \nabla^2 \Psi).$$

Yet we must not forget: (1) Algebra did not always exist; it was invented, just as other systems have been and will be developed, especially with the advent of computers. (2) Algebra is not apt for all areas of science; it has not been nearly as important in biology as in physics, and I am quite sure it will never be so central in cognitive science.

Coming to see or hypothesize patterns—discovery—is an important mental act that can be aided by literacies, especially those based on simple, systematic representational systems. Systematic representational systems aid discovery because they convert abstract “intellectual” patterns into spatial, visible ones. James Maxwell discovered an important electromagnetic phenomenon essentially because a missing term broke a nice pattern in a set of equations. There’s a miniature example here in Galileo’s six theorems. Why are there six? Might there be more, or perhaps fewer would do? Probably only the most diligent and perceptive reader noticed the pattern in Galileo’s discourse, but, at least in retrospect, the algebraic form makes it evident: Start with the ratio form of $d = rt$. The first three relationships eliminate in succession each of the three basic quantities—rate, distance, and time—by declaring a ratio equal to one; the second three express the full relationship, solving one at a time for distance, rate, and time ratios. We note also, therefore, that the first three relations are special cases of the last three.

Casting a wider net, graphs are another obvious case where a written literacy makes pattern detection easier. A graph that swoops up shows us instantly that a quantity is increasing faster and faster.

I have been listing ways in which written inscription systems can make us smarter, illustrating some details of material intelligence using algebra as an example. Inscription systems and associated subliteracies can effectively improve our memories, even without our rereading what we wrote. They may be well adapted to saying clearly, precisely, and compactly the particular things that need to be said in a particular field of study. They may also extend our abilities to detect patterns and make discoveries. The last example I want to deal with here is at the heart of intelligence: reasoning—the ability to draw inferences.

Look again at what I did with Galileo's six theorems and think how those results would appear in the minds of modern, algebra-literate knowers. I do not think an investigation is necessary. We know that $d = rt$ is a part of our current mathematical and scientific cultures, but Galileo's "six laws of uniform motion" are not. We can certainly tell what they are about, but they are not cornerstone pieces of our basic understanding. The reason is fairly obvious. The modern algebraic form is simply much better adapted to exactly what needs to be said. It produces a compact, precise, memorable statement at exactly the right level of abstraction. But what about Galileo's theorems? Have we lost them? Hardly. These theorems are so easy to derive algebraically that a student could easily manage the task, and a scientist would do it so effortlessly that no one would think to consider it a new result.

Think about it this way. Theorems are necessarily true given the axioms and definitions out of which they flow, so why do we bother writing the theorems at all? We do so simply because reasoning is sometimes expensive, and we just can't afford to reason from basic principles each time. We struggle to derive a result once, then essentially memorize it so that we have it quickly available whenever we need it. If reasoning suddenly becomes inexpensive, we can keep just the definitions and axioms in our own minds and derive particular theorems at need. In this case, reasoning became inexpensive because of a new modicum of material intelligence. We can quickly and easily see many implications of an algebraic expression by pushing symbols around. Galileo's intellectual

terrain had six small hills. The algebraically enhanced version is one tall, powerful mountain of a result that covers the whole area of those six hills and more besides, using the glue of algebraic reasoning to hold it all together.

I find this a provocative image. Not only can new inscription systems and literacies ease learning, as algebra simplified the proofs of Galileo's theorems, but they may also rearrange the entire intellectual terrain. New principles become fundamental and old ones become obvious. Entirely new terrain becomes accessible, and some old terrain becomes boring.

A Social View of Material Intelligence

When I first introduced social components of literacies, I made two basic points. First, as with any of the major intellectual accomplishments of society, there is always a gradual, cumulative development that involves many people. The second point is about the conversion of a material intelligence in a technical sense (which Newton and Leibniz had) into a true widespread literacy. The simplest version of the latter story is that a community decides a material intelligence is powerful and valuable enough that it is worth the considerable effort of teaching it to all newcomers. The community then puts in place an infrastructure for teaching it—freshman calculus or ninth-grade algebra, for example.

In this section, I want to expand these points, particularly the second, into a more faithfully complex view of the social processes surrounding literacy. We need, most of all, to begin to address the following central questions:

1. What determines whether a literacy can exist? and
2. What determines its nature?

In this way, we may be able to make a more intelligent assessment of whether computational literacies can come to exist, what they may be like, and, as important, how we may design and foster them.

Let me begin with a modest first try at a definition of literacy:

Literacy is a socially widespread patterned deployment of skills and capabilities in a context of material support (that is, an exercise of material intelligence) to achieve valued intellectual ends.

The “intellectual” part is merely to emphasize that we’re not talking about skillfully operating a piece of heavy equipment to dig a hole. The “patterned deployment” part is to avoid lumping all versions of widespread material intelligence under one umbrella. Unless we distinguish, for example, patterns in using algebra from patterns in using ordinary text, we can’t be specific enough to rule in or rule out particular future literacies. Essentially different patterns in the deployment of literacy skills need to be understood separately, possibly on different principles. Algebra doesn’t work cognitively or socially like reading and writing natural language. Computational literacy will exhibit still other patterns.

This initial definition turns out to have a great deal of ambiguity in it. Keep in mind that the important thing we want to do is think about possible future literacies, rather than present ones where we have a better sense of what is included in a literacy and what is not. Ambiguity makes it difficult to decide what literacies are sensible and possible. Looking for ambiguities, start in the middle, with “material support.” What materials? What support? With conventional literacy, presumably we mean text, but do we mean text in newspapers, in books, on notepads, on computer screens, on blackboards, or indiscriminately all of these at once? And what kind of support? In the previous section, I listed a fairly big collection of ways algebra supports intellectual accomplishments, yet that list is scarcely complete. Indeed, I do not believe that it can in principle be complete, for new physical inscription systems bring about new possibilities.

Scanning across the definition we also meet “skills and capabilities.” Which ones? No respectable account of all the skills and capabilities that humans possess has been produced, and just as with “support,” we cannot expect any closed list to suffice. Innovation in the material means of possible literacies may make any list of essential skills obsolete. The invention of graphing made curve recognition skills relevant to intellectual pursuits in a whole new way, and in fact it redefined those skills with a new vocabulary. Trivially, but not inconsequentially, a certain kind of manual dexterity and hand-eye coordination became relevant with the invention and adoption of the computer mouse. More profoundly, any given skills may change their effect on and relevance to valued accomplishments with the development of other skills. For example, arithmetic

may be a valued skill, but it changes its entire context—its community association, if not its essential meaning—when quantitative sciences give arithmetical computation new reach. Not just accountants but also engineers and scientists use arithmetic, and for each, it is relevant in a different way. For accountants, arithmetic may be “keeping track”; for engineers, it may be “deciding on an element of design”; and for scientists, it may be “tracing implications of a theory.”

Finally, values, as in “valued intellectual ends,” add another dimension of ambiguity in the proposed definition of literacy. Whose values, and of what sort? Scientists’ parsimony, citizens’ political empowerment, artists’ aesthetics, a child’s joyfulness in play?

Although I wring just a bit more specificity out of our preliminary definition in a moment, there is a fundamental lesson here. We must recognize an inescapable diversity in the phenomenon of literacy. There is no essential, common basis of literacy along any of the dimensions listed or along any other similar ones. There are no fixed basic human skills on which it builds. If oral language is a central competency, it is one among an open set of competencies we have or can build. Even oral language itself is open to innovation; we talk in different ways about different things depending on many other components of our material (and immaterial) intelligence. For example, we anticipate and build on non-speech intelligences in our talk—say, reciting $F = ma$ or announcing preliminary guesses of successor equations whose ultimate value will be tested substantially in different, more material modes. This is not to say that intrinsic human intelligence is infinitely malleable, but that existing and future intelligences draw on and engage it in such complex and intricate ways that guessing essential commonalties is not much more than an entertaining parlor game.

Similarly, saying what we get out of literacies is at best a tentative and culturally relative pursuit. We might identify intellectual powers (e.g., improved memory, more “logical” reasoning capability, precision in expression, metadiscursive competencies such as better understanding or manipulation of context dependencies in expression, etc.) or instrumental capabilities (say, “mastering nature”). However, these outcomes certainly vary across different material forms and practices; they are value related and hence depend on culture.

Construed scientifically, this claim of fundamental diversity is contentious and probably unpopular. In a different context, it would deserve a lot of exposition in defense. In this context, however, I believe the claim is properly conservative and at least heuristically correct. Whether or not the claim proves ultimately true, we simply cannot afford to limit our explorations of possible future literacies to extrapolations of what we think we understand about literacy now. Every claim for the essence of literacy can suggest how we may do better with computers, but computer-supported literacies may also work in completely different ways. At this stage, we need generative ideas as much as we need restrictive ones.

Still, can't we do better than "anything goes"? Yes, we can. What do the following have in common: newspapers, magazines (from *People* to *Soldier of Fortune* to *National Geographic*), scientific papers, pulp fiction, poetry, advertisements, tax forms, instruction manuals, and financial prospectuses? The seemingly innocuous but essential observation is that, although they use mostly the same basic material form, they each serve different groups of people in different ways. Variations in form and patterns of use from one to another are comprehensible as adaptations to serve particular purposes in particular contexts.

Let me introduce some terminology. I call each of the specialized forms in which we find literacy exercised in production and consumption a *genre*. This use of the term is a little different from the conventional use in literary criticism, especially when we extend *genre* to cover patterns in the production and consumption of algebra or of new computational inscription forms. However, the basic idea of a recognizably distinct use of a common material substrate is preserved as long as we also emphasize that genres serve particular groups of people in particular ways.

This latter idea—that any genre fits the needs and circumstances of a community—I describe with the phrase *the genre fits a social niche*.

Consider the following example, which I call the subway romance novel-reading niche. A few years ago when I rode the subway regularly in Boston, I undertook an informal study. I noted each day how many people were in my car, how many were reading, and what they read. I noticed that a large percentage of people read (a surprisingly small proportion of these read newspapers), and a reasonable proportion of these

riders read romance novels. Think about all of the factors that go into the creation and perpetuation of this genre in its niche.

1. It goes without saying that the romance novel niche rests on the well-established universal literacy basis developed in public education. I doubt this niche could self-generate without that prerequisite; the effort to learn to read is too great for the incremental value of being able to read a romance novel.

2. Almost all subway romance novel readers are women. This says a lot about the position of women in our society.

3. The Western concept of romantic love is an essential constituent. Whatever currents created and sustained the idea, romance is at the heart of romance novel reading. Other cultures would not recognize the sense or value of this genre.

4. Similarly, whatever personal value is perceived in the genre, it is important that there is no public social sanction against reading such novels. There's a delicate balance here. How many fewer public readers of *Playboy* are there because of the very modest and sporadic disapproval it brings? Religious fundamentalist cultures disdain and suppress both romance novels and pornographic magazines.

5. The price of production and cost of paper are relevant. A fifty-dollar romance novel wouldn't sell. Similarly, it is important that writers of these novels can make a living writing or else that it is possible to write while moonlighting. How important is the ubiquitous corner drugstore or newsstand to distribution?

6. The invention of the printing press and paper are relevant technical accomplishments. Cuneiform tablets just wouldn't work.

7. The requisite unoccupied commuting time relies on the existence of mass transit and whatever public values and political processes were necessary to create it. I haven't any idea what proportion of romance novel consumption comes from subway reading, but I'll bet it is significant enough that the demise of subways would be a blow, maybe even a fatal one, to publishers. It is also important that the trains are not outrageously crowded or noisy and that the readers' investment in and nature of their jobs doesn't force out pleasure reading.

These observations are almost the opposite of any claim that there is an essence to the operation and power of literacy. The conditions for

creating and sustaining a genre in its social niche reach deeply into and depend delicately on all sorts of physical, social, cultural, institutional, and historical conditions.

We can consolidate this view of literacy in a central hypothesis.

A literacy is the convergence of a large number of genres and social niches on a common, underlying representational form.

Genres are the variously refined and specialized styles of the underlying form, as a romance novel is a specialized sort of text. The social niche defines the complex web of motivating, enabling, and constraining factors that, first and foremost, allow a stability in the form of the genre and in its characteristic pattern of production and consumption. The social niche not only establishes the conditions for existence, but should also explain the defining characteristics of a genre. Existence and nature were the two basic questions that started this inquiry into the social basis of literacies.

The term *niche* is borrowed from ecology, where species—their characteristics and their survival—are studied according to the niche they occupy in the complex web of dependencies in which they participate. Does a particular species have enough land to forage; is it physically adapted to eat available food; are conditions right for the production of that food; are natural predators limited in some way? Genre is to social niche as species is to ecological niche. The challenging game in both inquiries is to discover and identify the necessary and possible types of interdependency. More than any other aspect of this metaphor, I believe that the complexity and range of types of interdependency for social niches of current and future literacies will match or exceed the complexity and range we are still discovering in ecology.

One aspect of a social niches inquiry is manifestly even more complex than for ecological niches. At least for biological niches, the basic chemistry of life is stable. We are all carbon-based life forms that use DNA to pass information from generation to generation. In contrast, our interest in genres and social niches is predicated on a substantial change of the basic material substrate—from static and mainly linear forms to essentially dynamic, multiply connected, and interactive computational media. This change is the main reason for the inquiry, but it also makes the inquiry more difficult and less definitive.

What we know and what we don't know is put in high relief by the concept of social niches. Multiple genres and niches explicitly represent inescapable diversity that strongly motivates broad exploration into new niches now that the "chemical basis of life" in this new "ecology" is moving to electronic forms. On the other hand, social niches also emphasize limits and our scientific accountability, stemming from the basic requirement to assess and explain viability of new social niches. We can't make just any new literacy, no matter how good it might be for us. Social viability is a harsh master. The "skills basis" and "support for intellectual ends" of our first proposed definition of literacy are put in a larger context, including dimensions such as economics and cultural history. Everything we know about each of these dimensions is relevant in principle.

To summarize, a social niches view of literacy comprehends the variability we know from conventional literacy as inescapable. We need to make room for both pulp novels and scientific papers. Each genre fits a different context in a different way. Recognizing that diversity, we are prepared for a future that could be very different. At the same time, we know that not everything can work. In understanding what works and what might work, we need to examine many perspectives on the viability of a niche and the fit of a genre to it.

I wish to cover three other general issues about literacy, genres, and social niches here. The first emphasizes the uncertainty of the central hypothesis concerning social niches and literacy. The question is whether a large number of genres and niches must be involved in a literacy, or would a few—or even one very important one—do? It seems clear that the current widespread textual literacy works because of the existence of a large number of niches that use basically one common representational form. What about all possible future literacies? My bet is that the most important literacies will always work in this way, and the work described in this book assumes that. This issue marks an important choice point determining the kind of software systems we design. Do we design a large number of pieces and kinds of software to fit into a diversity of niches? Or do we follow the pattern from the case of written text and try to create a rich medium capable of supporting a protean array of niches? The work described in this book follows the latter course—aiming to

change minds with a single, if extremely versatile, material form. The issue of forms and niches is further explored, for example, under the banner of multifunctionality in chapters 6 and 7, and then also in chapter 9.

Second, I want to make explicit yet another layer of complexity in the analysis of social niches. Start with an image. Think of a grand canyon of textual literacy carved up into quasi-hierarchical subliteracies and sub-subliteracies, which are, metaphorically, branches, gulches, rivulets, and microrivulets built into the texture of the canyon. People read; they read novels (or scientific works); they read pulp fiction (or historical novels); they read pulp romance novels (or science fiction). Down the scientific reading branch, there are subbranches for distinct genres such as treatises, papers, and so on.

The grand canyon view of literacy concentrates on form. But it hides both many uniformities and irregularities. Consider this irregularity: treatises in both philosophy and in mathematics are treatises and not papers or novels, yet they are noticeably different from each other, for understandable reasons: the forms of argument in mathematics and philosophy are different, and this difference propagates into the literary form. Consider also the following (hidden) regularity: you may be tempted to think the influence of the scientific community resides only in “its” genres, but sometimes scientists behave more or less as a bloc with respect to other genres. This might happen because of communitywide characteristics such as economic class, level of education, and so on, so the influence of the scientific community on nonscientific genres is scattered about the grand canyon and not made clear in its structure. We can’t see, for example, that the scientific community as a whole is irrelevant to the subway romance novel–reading niche, but that it may be very relevant to science fiction or film. We can’t tell that wiping out scientists wouldn’t affect romance novels, but that making secretaries’ jobs more interesting perhaps would.

I think of this blindness as a problem of perspective. If we choose to think of social niches as geometric forms, then they are forms in a high-dimensional space, not the two or three of the grand canyon. If we choose to look at them from one perspective or another—say, material form, community, or values—we will see the niches grouped in different ways, with different relations among them. Take the grand canyon view (form,

subform, etc.), tilt it on its side to get a community view, and you may notice that Jane Austen novels and books such as *Relativistic Quantum Fields* have much more in common than you might have thought.

Finally, let me list a few perspectives on social niches and comment briefly on them mainly because they are relevant to future prospects. We'll come back with more extensive discussion of these perspectives in chapter 9, after we've prepared a better understanding of the possibilities for computational literacies.

1. *Values, interests, motivations.* I know of no really good scientific theory of these, but without question we must take them into account in designing or studying social niches. Of the list of values I mentioned earlier—including scientific, political, artistic, and playful sensibilities—I take two to be most important. Naturally, my personal interests are building on and developing scientific aesthetics—for example, wanting to understand how things work and a great appreciation for the power and parsimony of theories. The other central kind of value may be surprising. It is, at least emblematically, whatever interests can lead a child into extended, self-motivated activity. I think the dawn of computational media is precisely the right time to remake the experience of science and mathematics learning in schools so that interests and values are not ignored. This revision will be a major topic in chapters 4 and 5.

2. *Skills and capabilities.* Textual literacy draws on certain human competencies and not others. For example, the immense competence of humans in dealing with both dynamic and spatial configurations is barely engaged by conventional literacies. We can do better electronically.

3. *Materials.* The material form of future computational literacies is a huge open question, and it may be the place where our directed skills as designers can have most leverage. We can wait for things to happen by accident, or, with due respect for what we do not know, we can move deliberately in the direction of the best we can imagine.

The form of future inscription systems is one thing, but delivery and use are another. With delivery at least, we seem in fine shape technologically for many possible literacies. Unlimited inexpensive or free distribution on CD-ROM (or DVD or other future versions) and via network is already a reality. A really portable personal computer for every teacher and school child could make an immense difference in richness of social niches, I am convinced. We are within a short technological and economical hop of ubiquitous availability. Politically and with respect to a sufficiently clear and convincing public image of what we might achieve, we probably have a longer distance to go.

4. *Community and communal practices.* Current community structures are important, but future possibilities are equally important, if not more so. It is probably too arrogant to think we can design new communities, but as network communications become universal on computers, we may be able to promote productive changes.

5. *Economics.* Hardware is much less the issue than software. It is still difficult to make money with educational software. The research and development of future literacies is an issue of public trust if ever there was one, but the issue doesn't even appear on the agenda of any government agency. If the conclusions of this book are correct—or even a responsible good guess—we are making a terrible mistake by this omission.

6. *History.* Cultural and technical history are powerful currents. The development of a computational basis for new literacies is orthogonal, if not antithetical, to most current trends. I discuss history (and some other of the issues listed above) more in chapter 9. In the best case, blindly following current directions means a delay, possibly a long one. In the worst cases, we'll do things such as standardize suboptimal technology, of which the awkward QWERTY keyboard, which we all are stuck using, is emblematic.