

June 1999

Changing Minds

Computers, Learning, and Literacy

To Appear
MIT Press
1999

Andrea A. diSessa
University of California at Berkeley
diSessa@soe.berkeley.edu

1. 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, and 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 professional. I keep up with some aspects of my work that don’t get covered in professional journals (like what features one gets in a cheap, 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. A lot of people buy newspapers, and I'm sure many of them 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 upon, 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 15 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. It makes 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 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. Not only students, but also teachers 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 20 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 like science and business. Just look at texts, tests and assignments from core subjects. They really have changed little so far. Numbers tell a more optimistic but still muted story of penetration. In 1995, K-12 schools in the U.S. had about three computers per “average” 30-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 1/3 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 1/2 computer per classroom per year. At that rate, average schools can easily meet my benchmark in a decade and a half. Over 10% of the high schools in the country are *already* above the threshold benchmark.

I fully expect the rate of computer acquisition to accelerate. That 1/2 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, 10 years, computers will be easily 10 times more powerful (30 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, unless our society is suicidally reluctant to share the future with its young, schools will soon enough be computer-rich communities.

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 deeply felt or well-thought-out objections. 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 immediately 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 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 “computer 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—except 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, like 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 talking about is not artificial intelligence in the sense of placing our own intelligence or knowledge, or some enhanced version of it onto 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 macro-components. The micro-focus will show a little about how traditional literacy actually works in episodes of thinking with a materially enhanced intelligence. The macro-focus will introduce some large-scale and irreducibly social considerations that determine whether a new literacy is achievable, and how. Much of the rest of the book will build 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 will 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 pillar of material. That is, literacy involves external, materially based signs, symbols, depictions, or representations. This last set of terms, and others besides, 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 hard 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 U.S. 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 exhibit limitations in what they may allow us to think about. They have associated with them particular modes of mediated thought and connections to other subsystems. Written language, the prototype of literacy, has an alphabet, a lexicon, a grammar, and 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, say, 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 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), 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 a lot to gain by thinking carefully about what the whole game of literacy is and about what we can do with computers that can hasten or undermine new possibilities.

The second pillar of literacy is mental, 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 must be 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 “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 in a way that conventional literacy barely touches. I will 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 of social and individual to get a feeling for the issues.

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 where a new material intelligence emerged out of the specific needs of an investigation, and that new intelligence clearly contributed to Newton's ability to state and to validate his new scientific accomplishments.

Fundamentally, the calculus is a way of writing down and drawing inferences about (e.g., 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 non-uniformly 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 on a blank slate. He borrowed and extended techniques, even graphical techniques, that had been around certainly since Galileo (1564-1642), 50 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 with Newton. Leibnitz (1646-1716), most believe, independently developed the calculus at about the same time. Indeed, the notational form mainly in use today is Leibnitz's, not Newton's. Although I can't prove it, I believe the reasons for this fact are in important measure pedagogic. Leibnitz'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 Leibnitz'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, $\frac{dx}{dt}$. (The d in dx and dt stands for a change, or "delta," in the quantity.)

Newton's notation is opaque, \dot{x} . In Leibnitz's notation, the "change of variable theorem,"

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

looks obvious, even if it is not. “Cancel the *dy*” 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 $\frac{dz}{dy}$ and $\frac{dx}{dy}$,

whereas Leibnitz’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 Leibnitz’ contributions are both obvious and subtle. Obviously, once again, science is a strongly cumulative social enterprise. Contributions must be both shared and extended by others to reach their full potential. 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, pales in comparison to the possibilities of popular new literacies.

The second moral from Leibnitz is more subtle, but it explains why I will spend two chapters on the material basis of computational media. The inscribed form of thought is critically important. I’ve suggested that Leibnitz 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” (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% harder to learn. Would those early innovators have had the courage to guess it might succeed? Similarly, at the second phase, if 10% fewer students “got it,” would they have declared success; would others have followed and had enough success for the whole project to succeed? Finally, might Leibnitz’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 on details of material form and on social forces in an essential way, one that 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 will illustrate 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 will be plenty rich and complex enough dealing with mathematics and science. As a bonus this little story will lead directly into my own experiences teaching children about motion using computers.

The last reason to look at Galileo, returning 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 Galileo's *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 below.)

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:

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

Theorem 1:

In the case $r_1 = r_2$, the r terms cancel, leaving $\frac{d_1}{d_2} = \frac{t_1}{t_2}$.

Theorem 2:

In the case $t_1 = t_2$, the t terms cancel, leaving $\frac{d_1}{d_2} = \frac{r_1}{r_2}$. Conversely, if $\frac{d_1}{d_2} = \frac{r_1}{r_2}$

then $\frac{t_1}{t_2} = 1$ or $t_1 = t_2$.

Theorem 3:

In the case of $d_1 = d_2$, the d terms cancel, leaving $\frac{r_1}{r_2} \frac{t_1}{t_2} = 1$, or $\frac{t_1}{t_2} = \frac{r_2}{r_1}$.

Theorem 4:

This is precisely our little ratio lemma, $\frac{d_1}{d_2} = \frac{r_1}{r_2} \frac{t_1}{t_2}$.

Theorem 5:

Solve the equation above for $\frac{t_1}{t_2}$; $\frac{t_1}{t_2} = \frac{d_1}{d_2} \frac{r_2}{r_1}$.

Theorem 6:


$$\text{Solve for } \frac{r_1}{r_2}; \frac{r_1}{r_2} = \frac{d_1}{d_2} \frac{t_2}{t_1}.$$

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

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 *Dialogues Concerning Two New Sciences*, Galileo. Translated by H. Crew and A. de Salvio, Northwestern University, 1939.)

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, nor the education provided by his parents, nor 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 later than Galileo’s main work, 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.

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 said as they are, are less compact and less memorable. Our memories are better with good external inscriptions, even if you are not using the material form as memory by rereading what you 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 sub-literacies 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' analytic geometry can learn all of Galileo's accomplishments concerning motion in very short order.

A piece 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. But these details are irrelevant; the algebraically expressed relations apply to any pair of motions. An even higher level of abstraction than that of equations turns out 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 ($\nabla \times \mathbf{E} = -\frac{1}{c} \frac{\partial \mathbf{B}}{\partial t}$, 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 that: 1) Algebra did not always exist; it was invented. 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

representation systems. Systematic representational systems aid discovery because they convert abstract “intellectual” patterns into spatial, visible ones. 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 sub-literacies can effectively improve our memories, even without our rereading what we wrote. They may be well-adapted to saying the particular things that need to be said in a particular field of study, to saying them clearly, precisely, and compactly. 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 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. So we struggle to derive a result once, then essentially memorize it so that we have it quickly available whenever we need it. But if reasoning suddenly becomes cheap, we can keep just the definitions and axioms in our own minds and derive particular theorems at need. In this case reasoning became cheap because of a new modicum of material intelligence. We can see many implications of an algebraic expression quickly and easily 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 Leibnitz 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 enrich 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, it 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. More deeply, 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 first-try definition turns out to have an awful lot 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 last 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 are these? 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 and relevance to “valued accomplishments” with the development of other skills. For example, arithmetic may be a valued skill, but it changes its whole 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 of these, 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, arithmetic 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, of what sort? Scientists' parsimony, citizens' political empowerment, artists' aesthetics, a child's joyfulness in play?

Although I will 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 any other similar ones. There are no fixed basic human skills that it builds on. If oral language is a central competency, it is one among an open set of competencies we have or can build. And 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 commonalities 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, meta-discursive competencies like better understanding and manipulating context dependencies in expression, etc.) or instrumental capabilities (say, “mastering nature”). However, these outcomes certainly vary across different material forms and practices, and they are value-related and hence will 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, financial prospectuses, and so on? 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 is a little different from the conventional use of the term in literary criticism, especially when we extend "genre" to cover patterns in the production and consumption of algebra, or of new computational inscription forms. But 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 will describe by saying *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 that brings? Religious fundamentalist cultures disdain and suppress both romance novels and girly magazines.

5) The price of production and cost of paper are relevant. A \$50 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, if not 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/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 these 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 like 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 is to underline the uncertainty of the central hypothesis concerning social niches and literacy. The question is whether a “large number” of

genres/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. We'll explore the issue of forms and niches further, for example, under the banner of multi-functionality 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 a quasi-hierarchical sub-literacies and sub-sub-literacies, which are, metaphorically, branches, gulches, rivulets and micro-rivulets 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). Following down the scientific reading branch, there are sub-branches for distinct genres like 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 philosophy and in mathematics are both 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 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 community-wide characteristics like economic class, level of education, and so on. So the influence of the scientific community on non-scientific 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 may be very relevant to science fiction or film. We couldn't tell that wiping out scientists wouldn't affect romance novels, but that if secretaries' jobs were more interesting, perhaps that would.

I think of this blindness as 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 2 or 3 of the grand canyon. And 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, sub-form, etc.), tilt it on its side to get a community view, and you may notice that novels of Jane Austen and books like “Relativistic Quantum Fields” have a lot more in common than you might have thought.

Finally, let me list a few perspectives on social niches and comment briefly on them mainly as they are relevant to future prospects. We'll come back with more extensive discussion of these 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 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—like 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 (DVD or other future versions) and via network is already a reality. Really portable and personal computers for every teacher and school child could make an immense difference in richness of social niches, I am convinced. We are within a short hop of ubiquitous availability technologically and economically. 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'll discuss history (and some other of the issues 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 like standardize on sub-optimal technology, of which the awkward QWERTY keyboard, which we all are stuck using, is emblematic.

2. HOW IT MIGHT BE

Chapter 1 set the scale of the enterprise—powerful new material intelligence for human civilization—and opened inquiry into the material, cognitive, and social foundations for new literacies. But what might a computational literacy be like? This chapter presents two concrete images of what the future may hold. One will be oriented mainly cognitively and the other mainly socially. Because it is too easy to wax poetic about unimaginable future revolutions, I want to strike a careful balance between present and future. Both of these examples will be rooted in experiences we have already had, but I also want to extrapolate toward what I believe to be genuine but profound future possibilities.

Beyond Algebra

Algebra brought certain kinds of reasoning from the province of geniuses like Galileo into the grasp of average high school students. But, suppose there were some new representational form that was (a) easier to learn, (b) learnable at earlier ages, (c) even better adapted to describing important aspects of motion, (d) synthetic—that is, it could produce motion for all to see, rather than just analyze or describe it, (e) useful for even a broader range of subjects compared to algebra, and (f) an incredible lot of fun to play with? Surely such an advance would be immediately adopted by schools.

Let me jump in the deep end with both feet. Suppose that by the end of elementary school, students were literate with computer programming. I am taking computer programming languages to be a material form for a hypothetical new literacy, and I'm assuming that programming is within the grasp of elementary school students. I want to look directly at how such students could approach something like Galileo's six motion theorems, and also the next step in his research program—defining and exploring the properties of uniformly accelerating motion. Evidently we're looking at going algebra one better by taking the introduction of some basic ideas of motion back about 5 grades (from high school physics) on the

basis of a new and different literacy. I will use Boxer, the prototype computational medium my colleagues and I have been working on as the example notational system.

The three elements in Figure 1 constitute a computer program that represents uniform motion. We call it the *tick model* as, at its core, it shows what happens at each “tick of the clock.”




Figure 1. The “tick model,” which defines uniform motion.

The tick model, like almost any program, has two parts. The *informational* part, or data, is simply a number that represents the speed of the moving object. In Boxer, this corresponds to a data box with an appropriate name, say **s**, for speed. That is the first box in Figure 1.

The second and third boxes define the *procedural* part of the program, what happens with the data. These are *doit* (do it) boxes. The second box in Figure 1 represents the overall “shape” of the process. In this case, it is very simple. In order to **run**, you **loop** (repeat over and over) the simple action, **tick**. **Tick**, in turn, is also very simple. **Tick** does a **move s**, which commands a graphical object to move the short distance specified by the speed, **s**. Every tick of the clock (perhaps **tick** should sound a little click), our object moves a little bit. For simplicity we can make our computer loop through one tick each second.

Each of these parts is independently important. **Run** tells us that exactly the same thing happens over and over—the process is uniform. **Tick**, of course, says what happens over and over. And **s** quantifies the physical process. A bigger **s** means a proportionally greater speed. Figure 2 shows the program “from top to bottom” with dotted lines to emphasize dependencies.




Figure 2. Showing the dependencies in the tick model.

What I'd like to do now is consider the properties—the advantages, and some disadvantages—of this representational form as an aid to thinking and learning about motion. I will use more or less the same set of properties we discussed earlier in considering calculus, and more particularly algebra. In fact, I will use algebra, along with text and to a minor degree calculus, in making comparisons to programming, which will highlight particular features of each representational scheme.

Let us start with expressiveness. Programming is, simply and directly, marvelously adapted to expressing many things about motion. The tick model illustrates the fit between motion and programming well. All the important features of uniform motion appear, and they appear pretty much in independent places. I would call the latter property “information factoring”; you need to look in only one place to see one of the essential features. Look at **run** and you see the uniformity of the process. Look at **tick** and you see the nature of the process at each instant. **s** shows a simple, numerical parameter, and the use of **s** in tick shows exactly the nature of the parameter's influence.

Compared to text, the tick model shows the characteristic strengths of a technical representational system. It is concise, with all the advantages accrued by compact form, including, usually, memorability. It is also precise. Not only do numbers have a natural fit in the representation, but the process depicted is quite unambiguous, if the “reader” is familiar with the computational meaning of the program. The latter conceptual clarity is hardly mysterious. Programming languages like Boxer are precisely languages that express certain kinds of processes without

ambiguity. Learning the language means learning how to interpret these inscriptions. Once you have learned that, every instance where the language aptly describes a situation is an instance where you reuse your basic understanding effectively.

Aptness, of course, is critical. Every inscription system is apt for some things and less apt for others. Programming, like arithmetic, is not evidently very apt for poetry nor for discussing the nature of computational media. But it happens to be quite superb for motion. Being apt for a wide range of things, of course, is pretty much the question of multiple niches for a computational medium. I won't try to discuss the extent and limits of programming's aptness in any detail here.

There is a slightly tricky point about precision and conciseness when we are talking about text and programming as media. Language always affords you the option of developing an intricate technical vocabulary and expressive figures of speech to tune its general capabilities to a particular use. Not only can you invent technical terms, but you can talk about other representational forms. For example, you can speak an equation by saying "distance equals rate times time." But there are limits to this enfolding of other inscription systems. In the first place, you have to do the work of building the technical vocabulary or learning the other representational system talked about. And then you still don't get all of its advantages. I have frequently written little programs in my head at night in bed before falling asleep. But then, unless the program is completely trivial, I realize I'll have to try it in the morning to see whether it does what I wanted correctly.

Programming, like text, happens to be excellent at enfolding more specific representation subsystems, and at accepting "tuning" toward specific purposes (as in adding a specialized vocabulary). These turn out to be among its most important strengths. See Chapter 7.

Comparing to algebra illuminates other important properties of the tick model. Algebra—for instance, $d = rt$ —is arguably a bit more concise and better adapted to a familiar class of word problems: "How long does it take Johnny to...?"

But its very high level of abstraction, important to some central aspects of its power, is also a serious problem in other ways. Algebra does not distinguish at all effectively among motion ($d = rt$), converting meters to inches ($i = 39.37 \times m$), defining coordinates of a straight line ($y = mx$) or a host of other conceptually varied situations. Distinguishing these contexts is critical in learning, although it is probably nearly irrelevant in fluid, routine work for experts. The tick model picks apart uniform motion and sets it in a clear, well-developed process framework. Given how poorly adapted both language and algebra are to the central conceptual details of physics concepts generally (and given other issues, described below), I am amazed at how well we have managed to teach these things with older media. Even so, it is no surprise that we have had to wait years more than I believe is necessary—until after algebra instruction—before teaching motion concepts. And even then, students almost always get stranded at the high level of abstraction of algebra and symbol pushing, at least for a time. “Plug and chug” is the well-known phenomenon of novices’ avoiding all conceptual analysis by grabbing a seemingly relevant equation and “grinding” the algebra.

The conceptual aptness of programming for motion has a third and equally important component, beyond a process orientation and an appropriate level of abstraction. I mentioned before that “one of the most fundamental and enduring moves of all time in physics” was the move to understanding Nature’s laws as local, as unfolding only at instants of time and independently at each point in space. The emblem for that grandly important principle here is the **tick** procedure itself, which represents the heart of motion, what happens at each instant. Programming turns out in general to be really well adapted to local principles. We previously looked at the insufficiency of algebra to describe local principles, which forced Newton to invent calculus. Our hypothetical sixth-grade children are not only getting started learning physics much earlier and in some ways better than they could with algebra, but they are also being introduced to some of the critical post-algebra ideas aptly captured in current instruction only by calculus.

I can at least sharpen focus on the tip of this iceberg with another observation about the tick model. Looking at the tick model and understanding what it means as a process describing motion allows the following observation. The total distance traveled by an object is just the sum of all the \mathbf{s} 's, which are the distances moved each "tick." In the simplest case this is just t times s , where t is the number of ticks. We have just derived the basic relationship that we used to redo Galileo's six theorems. So $d = rt$ (in this case $d = st$) is within easy inferential reach of the tick model. But so is a generalization if \mathbf{s} is changing each instant. Distance traveled is still the sum of all the speeds, \mathbf{s} . Calculus-literate readers will recognize this as an important special case of the Fundamental Theorem of Calculus. The conventional way of writing this theorem is $d = \int s dt$, where \int is Leibnitz's sign for the integral operation, corrupted in a wonderfully heuristic fashion from S, for sum.

Going back to Galileo, we shouldn't be surprised that he had difficulty maintaining a focus on the local nature of motion, although some of his best accomplishments came when he managed to do that. An excellent case in point is that he did not derive his versions of $d = rt$ as I did above, by focusing on the instants and small bits of motion that accumulate into a larger motion. Instead, he embedded his motions in still larger motions—going in the opposite direction! (Amazingly, this method works, although, as you are invited to check in the proof quoted in the previous chapter, it involves both complexity and cleverness.)

The tick model is not the same as knowing algebra and $d = rt$, nor is it the same as knowing calculus. In particular, both algebra and calculus allow some inferences that are difficult to express in computer programs. Yet programming prepares the ground for these later inscriptional accomplishments to an amazing degree.

The last feature of representational forms that I want to discuss with respect to the tick model is how particular forms help in detecting patterns and making discoveries. This modest-seeming consideration actually opens up into an advantage of the tick model that may dwarf any of its advantages in power and

generality. Programs are not just analytic and a basis for reasoning. They are in addition synthetic. They can be run.

Running the tick model causes an object to move across the screen at a certain speed. Adjusting **s** adjusts how fast it goes. Replacing **loop** with **repeat t** (where **t** represents equally the amount of time during which the process runs, or the number of repetitions) means we can play with how long an object moves and how far it goes. Programming turns analysis into experience and allows a connection between analytic forms and their experiential implications that algebra and even calculus can't touch.

“Make it experiential” is perhaps the single most powerful educational heuristic that I know. Experts aren't left out either. The synthetic power of computer programs, say, to simulate weather or global warming transcends all other inscription systems. The capacity to simulate is defining the modern information age for many scientists and engineers. The tick model may be a small step, but it crosses over into that new world.




Figure 3. An enhanced tick model, with acceleration.

Allow me two steps of elaboration of the tick model. Figure 3 shows the tick model enhanced by a new variable **a**. In an analytic mode, you can see that **a** gets added to **s** at each tick. If you think about it some, you'll be able to understand some properties of **a** and how it affects motion. **a** is, in fact, the tick-model version of the concept of acceleration. Beyond analysis, it is easy to imagine children playing with this enhanced tick model to get an excellent intuitive understanding of acceleration, as well as an analytic one. A salient example phenomenon is that **a** exerts “delayed” control on an object: If you set **a** to be a positive number, over time the object speeds

up more and more. If you then set \mathbf{a} to be negative, the object doesn't either stop or change its motion in any abrupt way. A negative \mathbf{a} just nibbles away at a big \mathbf{s} until it becomes small.

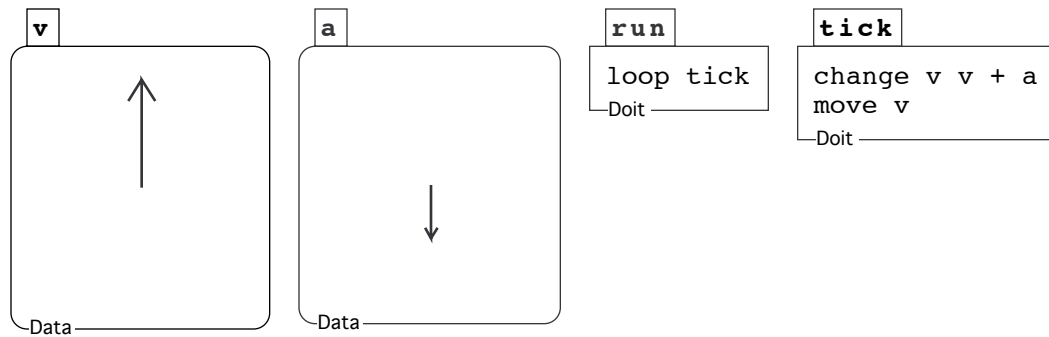


Figure 4. A tick model where numbers have been replaced with vectors.

One more step: Figure 4 shows a version of the tick model where numbers have become vectors. A vector is shown in Boxer as an arrow (in a box) that may be adjusted by grabbing and moving the tip. **Move** takes a vector and causes an object to move the length of the arrow in the direction pointed by the arrow; what could be easier? \mathbf{v} is a “fancy” version of speed, called velocity, that specifies direction as well as amount. You can see that at each instant, \mathbf{v} gets changed to the sum of the old \mathbf{v} and \mathbf{a} , just like the numbers in Figure 3. For those who don't know what it means to add vectors: If you had a computer, you could just click on the **change v v + a** line a few times to see what happens. You'd see that the tip of \mathbf{v} moves as specified by \mathbf{a} , in the direction \mathbf{a} points and a distance equal to its length, just as \mathbf{v} moves a graphical object. Figure 5 shows snapshots of \mathbf{v} being dragged downward, through $\mathbf{0}$ (no velocity, indicated by \bullet), by \mathbf{a} .




Figure 5. A sequence of velocities, each dragged downward from the previous by \mathbf{a} .

We have whizzed past a slew of Galileo’s accomplishments on the superhighway of a marvelous notation system. (Naturally, it would have been better to stop and look at some sights along the way.) But now that we’re here, consider two little examples. If we **run** the Figure 4 tick model as shown we’ll see an object rise (frame 1 and 2 in Figure 5 shows moving upward while slowing down) to a peak, seeming to stop for an instant (the middle frame in Figure 5 shows no velocity) and then fall down. This resembles what happens if we toss a ball. In fact, this is just how Galileo managed to describe a ball toss. Perhaps the most surprising thing about this description of a toss is that the form of the program is pretty much the same as where we started. It is describing a process that is *the same at every instant*, except that \mathbf{v} is changing. Most people, and almost all students, find this uniformity amazing. Few are prepared to believe “going up” works in exactly the same way as “falling.” Even more, essentially everyone believes the peak of the toss is very special, expressed by some sort of balancing (say, balancing of the upward toss force with downward gravity). In contrast, the top of a toss is really an innocuous happenstance, that adding \mathbf{a} to \mathbf{v} will drive \mathbf{v} through its zero-length position.

The tick model does not teach these things automatically and without a teacher. But it is an excellent way to experience, think and talk through how Galileo’s model works. What happens if you pull \mathbf{v} to the side so that it slants diagonally to the right? The moving object then makes a trajectory like a diagonally tossed ball. But since \mathbf{a} always moves the tip of \mathbf{v} directly downward, \mathbf{v} always slants

to the right, always has a component of rightward motion. Is this how the world works? It is, in an important idealization, but most find it initially unbelievable. Making unbelievable, but accurate, descriptions feel plausible is surely something to work through.

Dragging the tip of \mathbf{a} around while the tick model is running makes for an engaging and sometimes quite challenging game. Getting the object to stop where you want it to, for example, involves setting \mathbf{a} to bring \mathbf{v} to its no-length state exactly at the place you choose to stop. No mean feat. This is just the kind of job the pilot of a space ship has in order to dock with an orbiting satellite or space lab. In fact, according to Newton, \mathbf{a} exactly represents the effect of direction and magnitude of space ship thrusters.

We have come through expressive aptness, conceptual precision, and so on, to a really new place. From here, it is easy to imagine sixth grade students getting personally and creatively involved in designing space ships and all sorts of games. Galileo's and Newton's sometimes forbidding abstractions have been resituated in a fabric of doing—play—that can be owned by children. We can be instrumental and say that mathematics and science can be motivating and engaging in a way that far transcends words, algebra, and calculus. We can talk about “time on task” and notions of learning through rich feedback. Or we can say merely that we have managed to bring mathematics and science into a child's world in a way that shames “you'll need this for the next course” or “just do the exercises.” This last phrasing may be the most important, and I will spend a lot of time with it in Chapters 4 and 5.

Here is the point of this section in a nutshell. A new representational form, programming, as part of a new literacy can lead to deeper learning, much earlier with fewer unpleasant glitches, and in a way that transforms the experience of students substantially from doing what adults say in semi-comprehension into a really rich and appropriately kid-like experience, more like what they want to and can do without adults intruding awkwardly.

The hanging hypothetical component of the story is the key. Can sixth grade

children really become literate in programming well enough, with little enough trouble, that a lot of learning can be transformed as I suggested above? I have not been scrupulous in hiding for dramatic effect that this whole hypothetical scenario is a lot less hypothetical than it might seem. We have been through the path traced above with “genuine” sixth grade children, right down to some boggling surprises for them (e.g., that an upward toss and a downward fall work the same way) and the really pleasant times when students demanded more “work time” to finish their programmed motion games. It is particularly easy for me “to imagine sixth grade students getting personally and creatively involved,” because I have seen it again and again. Plus, we provided these students the necessary literacy in a few weeks. I’ll add relevant detail in Chapters 3 and 8.

Tool-rich Cultures

Having lived in two scientific communities—physics and cognitive/education research (and also computer science for a time)—I have been deeply struck by the relative impoverishment of schools with respect to tool creation and use. Scientific communities, in general, are without question tool-rich cultures. Scientists use telescopes, microscopes, particle accelerators, chemistry laboratory equipment like Bunsen burners, test tubes and analytical balances, vacuum pumps, multi-meters, oscilloscopes, stroboscopes, oceangoing submersibles, culture dishes, video machines for human and material performance analysis. And now thousands of pieces of software help mathematicians, medical researchers, physicists, psychologists, and so on, collect, sort, analyze, and present data. There are even meta-tools—tools to build tools—for example, in the university science department shop and in programming languages and software toolkits. Collecting and displaying tools of the trade is an excellent way of chronicling the history of the scientific enterprise.

Other important classes of tools are easy to miss. Beyond machines and equipment, it is particularly important in this book to recognize representational

tools such as algebra, calculus, tables, and all the rest. Without stretching too far we can include reference books and even libraries. Just slightly more exotic are all the many techniques and skills for operating physical and representational devices (graphing is a tool!). Arguably, especially in view of considerations I will take up directly, even “purely” intellectual capability, habits of mind tuned to particular services in scientific inquiry, should count as tools.

From our social niches point of view, the position of these devices within the scientific communities and cultures is more important than their nominal classification as tools. First and foremost, tools are instrumental to higher purposes. They are seldom, if ever, ends in themselves. For example, the moment a new, better tool or technique emerges to serve a needed function, the old tool is abandoned. As much as I valued my old, artfully made and trusted Pickett slide rule (colored yellow, just slightly tinted green, at the peak of human visual acuity), it fell instantly into disuse in some forgotten drawer as soon as cheap scientific calculators became available.

But as long as these tools serve their purposes well, they carry traces of the fundamental values and goals of the community. They accomplish the jobs that define and justify the very existence of the community. Tools are badges of membership, symbols of commitment and accomplishment, frequently tinged with affect like pride and sometimes (for beginners) embarrassment.

These tools are in a deep sense owned by their communities. The principles by which the tools work are frequently “community property” and taught to newcomers. Many are fabricated by and for the community.

Schools are quite another matter. In the category of physical tools, one begins to run out after books, black- or whiteboards, pencils, pens and papers. Yes, a few science classrooms have equipment for laboratory activities. But these are marginal to the main practices of lecture, reading and paper-and-pencil problem solving. Until computers, devices like VCRs, overhead projectors, video disk machines, and so on, barely show above the noise in widespread educational practice.

Representational tools and techniques at first present a more positive face. Indeed, the expressed purpose of schooling brings a reasonable dose of scientific and mathematical “tools” to classrooms. Yet even here, a closer look at social positioning reveals problematic differences compared to the scientific community. The gaps are easiest to see from the viewpoint of students. Scientific tools in school are almost always ends in themselves, or they are related vaguely and artificially to “doing well in school” or “preparing for the future.” Problems are assigned and understood by everyone as thinly veiled occasions to exercise tool knowledge or skills, rather than reasons for the tools’ existence. And many schools are really designed around incompetence in the sense that any real understanding is a sign to move on to the next topic. Pride in accomplishment is seldom reached.

The net result of these facts is that school—at least the academic component of school—feels artificial to most students. Students enjoy little sense of ownership, personal commitment, or exercise of competence. Because the purposes of the tools they learn are not their purposes, because “how they work” is seldom a topic for students, important learning mechanisms are undermined. (Knowing how tools work and understanding the jobs they do provide help in bootstrapping adequate skills in using them.) These points will be elaborated and extended in Chapters 4 through 7.

Teachers may be a little better off, as their agendas dominate school activities. But what intellectual and physical tools serve *their* aims and needs? Again, curricular scientific and mathematical tools are mainly targets rather than means to an end, although some few of them may count as instruments in the service of teachers’ higher goals. But most teachers make few if any of their own tools and materials; textbooks come from outside the community, filtered through political mandates and textbook manufacturers’ needs and perceptions of educational goals and means. The larger pattern of disenfranchisement and under-professionalization of teachers is a widely recognized structural problem in American education.

I believe computational media and associated new literacies may be exactly

the infrastructural change that can support converting schools, particularly mathematics and science classes, into vital communities of tool building and sharing. The new genre of software I have in mind is a set of open, reconfigurable, re-purposeable tools for student tasks, tasks that are more like design and student research than conventional activities such as exercises and short problem solving. Such a shift in student activity is consistent with much that is recommended in many current educational reform documents. Nevertheless, I believe those reform efforts could be immensely aided with suitable material support, a point that is almost never taken to heart. Reform needs implements, not just implementation plans.

For now, the Boxer vectors that we discussed earlier can serve as an example of such tools. Vectors, in fact, were produced as a simple toolkit by a graduate student member of our research group. The transition that I pointed out—from vectors as an expressive language for laws and properties of motion to a synthetic resource of students' own, self-motivated projects—is a critical one that needs to be generalized and repeated hundreds of times with different tools. Student programming literacy may be a common basis for these transitions. Starter kits for student design and research, kits built with open programming interfaces, will be a common form of product in this genre. In the next chapter, I discuss a teacher-created set of materials—the Infinity Box—which served as springing-off point for core student creativity via programming. (Anticipating the discussion directly below, I note that the Infinity Box and several other productions by its teacher-creator have served other teachers, as well as students, as resources.)

Making educational tools directly serve valued student goals, however, is only a piece of the image. I believe we can revitalize teachers' professional experience by fostering their central participation in the production of such tools, as well as in their roles as coach and mentor to students. We would be making a fairly radical departure from current assumptions and practice, where teachers are barely trusted to copy worksheets, let alone create (or even modify!) substantial new

materials. In what world might this really be possible?

In the first instance, my experience with teachers has convinced me that we have thousands of creative, energetic and competent teachers to draw on. We don't need to imagine that every teacher is exceptional. Five to ten percent of the teaching force as active participants in a culture of creative innovation through tools would be quite sufficient. Although we can always use more innovative teachers, I don't believe their numbers are the most critical barrier. Here is an alternative list of barriers:

1) Lowering the threshold of technical expertise necessary to produce cogent educational tools. This, of course, is nearly a defining goal for computational media. I won't say more here than if sixth-grade students can make substantial creative use of a medium, I believe teachers will be able to do so as well.

2) Increasing the skill level of teachers. Beyond thresholds, as in the item above, the more expert that teachers become, the more they will be able to accomplish. The presumption of new literacies is that we have many years ahead of bootstrapping, rising expectations, and corresponding rising accomplishments.

3) Changing the image of products. Current assumptions about software are that it comes in large units that are so slick and complex as to require many person-years of effort to create them, usually by highly technically competent software engineers. With sufficient resources in the underlying medium, general or more specific toolkits, and open products available to modify, such assumptions will change dramatically. Several of the Boxer examples that appear later in this book were based on teacher- and student-modifiable toolkits.

4) Changing assumptions about curriculum. I know that sixth-grade students can learn a lot about motion with a computational medium. The conventional wisdom, based on conventional media, is that this is impossible. Changing assumptions about what is possible to learn may be especially difficult for teachers. We would be asking them to teach more than they now know. Learning the subject might be more difficult for teachers because they may resist playing with the ideas

and programs the way children will.

5) Revising community affiliations and work groups. One main problem teachers have currently is not so much lack of creativity or insufficient numbers in a population of teacher-developers. Instead, it is a matter of sparseness in geographical distribution. It's unlikely that a threshold number of compatible people with similar goals work close together. As telecommunications becomes an integral part of computational media, as they are for Boxer, teacher colleagues will be able to work together no matter where they live. The Internet and World Wide Web demonstrate that both collaboration and product distribution can, in principle, be free and easy.

As I have suggested and will continue to illustrate, most Boxer teachers with whom we have worked have at least modified existing materials, and several have been autonomously creative in developing new materials. All it takes is one creative teacher and a small audience of collegial experimenters to offer help and criticism to begin a collaborative project. Little projects can grow gradually as their products are refined and judged worthwhile by a growing network community. Don't forget, also, that we are still in "prehistoric" times with respect to computational literacies. Dozens of excellent teacher creators could easily become thousands.

But suppose even that it takes a university research project to initiate a tool set. (Indeed, one would hope collaborative development communities would cross traditional boundaries of expertise and affiliation.) Still, programming a toolset is just the beginning. An incredibly important niche in which teachers can be especially effective is the creation and documentation of productive activities for students. It is fortunately easy in Boxer to create tutorials or interactive work sheets, say, to teach the basics of vector algebra or scaffold a student project. Yet another important service is to feed student work back to the community to refine materials, improve other teachers' preparation for difficulties and successes of students, and generally advance community know-how. (Again, see the Infinity Box in the next

chapter.) Collecting good examples of student work can be directly helpful to future students not only for ideas, but as models and even as starter kits and tools for their own projects.

Let me close this section by clarifying what I am and am not advocating here and what aspects of it are more or less certain. First, I want to share my personal enthusiasm for a certain new genre of educational software. Relatively simple educational tools can be created in fairly large numbers. These tools are open, flexible, and combinable and will serve students *both* by introducing and exercising powerful mathematical and scientific ideas *and* by opening avenues of personal and collective involvement in what are felt by students to be directly meaningful self- and community-expressive activities. I am convinced such tools are possible and that they are based on sound learning principles. We have had many, many experiences in working with students and Boxer to suggest that these possibilities are technically and educationally realistic, and that they can be powerful. In the remainder of this book, I will highlight some.

The teacher side of the critical social niches issues is more hypothetical. I have less doubt than many that there is a sufficient population of creative, energetic and potentially technically competent teachers who may engage each other in collaboratively fashioning such tools. The tele-community aspect of this scenario is already feasible. For reasons treated mostly later, I believe requisite technical skills in using Boxer or some future computational medium are attainable. The more difficult challenge is managing all the practical and cultural issues of the implied new social webs. Communities like this will likely have to arise out of a concerted effort to break current teacher isolation, and out of confidence and commitment that teachers want to and can take charge of parts of education that have not been their purview. These may well be best viewed as political issues of teacher professionalization. Can teachers, with help from some other groups, take hold of educational reform from within their group? Can they turn some very nontraditional opportunities offered by technology to their communal ends?

Finally, I do not offer the possibility of an educational culture of tool-builders and sharers—teachers and students alike—as a prescription or panacea. Instead, it is one among many possibilities that should be considered. Like the best possibilities, it has deep and complex technological and social roots that I hope it will be our privilege to cultivate and come to understand.