MINDSTORMS

Children, Computers, and Powerful Ideas

SEYMOUR PAPERT

BASIC BOOKS
New York

Cover design by Ann Kirchner

Basic Books
Hachette Book Group
1290 Avenue of the Americas, New York, NY 10104
www.basicbooks.com

Originally published in hardcover by Basic Books in 1980

Published by Basic Books, an imprint of Perseus Books, LLC, a subsidiary of Hachette Book Group, Inc. The Basic Books name and logo is a trademark of the Hachette Book Group.

Library of Congress Cataloging-in-Publication Data has been applied for.

ISBNs: 978-1-5416-7512-4 (paperback); 978-1-5416-7510-0 (ebook)

E3-20200903-JV-NF-ORI

Contents

BASIC
BOOKS

Frontispiece: LOGO Turtle.

FOREWORD

The Seeds That Seymour Sowed

IN REREADING *MINDSTORMS* TODAY, FORTY YEARS AFTER ITS PUBLICATION, I had two conflicting reactions:

On the one hand, so many of Seymour's ideas that were seen as radical in 1980 are now part of the education mainstream.

On the other hand, so many of Seymour's dreams remain unrealized and unfulfilled.

How is it that Seymour's ideas can be so aligned with today's realities but still so disconnected from them?

To make sense of this seeming contradiction, it's helpful to transport yourself back to 1980, when *Mindstorms* was published. The first personal computers had been developed just a few years earlier. No one had mobile phones or tablets or even laptops. The Web didn't exist, and few people had heard of the Internet. So it was truly radical to predict, as Seymour did then, that millions and millions of children around the world would soon be interacting with digital technologies every day, as they do now.

Even more radical were the ways in which Seymour imagined children using computers. In the small community of researchers who in 1980 were beginning to think about the use of computers in K–12 education, most focused on "computer-aided instruction," in which computers played the role of a traditional teacher: delivering information and instruction to students, conducting quizzes to measure what the students had learned, then adapting subsequent instruction based on student responses.

In *Mindstorms*, Seymour offered a radically different vision. For Seymour, computers were not a replacement for the teacher but a new medium that children could use for making things and expressing themselves. Using a particularly memorable turn of phrase in *Mindstorms*, Seymour rejected the computer-aided instruction approach in which “the computer is being used to program the child” and argued for an alternative approach in which “the child programs the computer.”

Seymour’s ideas about educational technologies have had a growing influence in the decades since the publication of *Mindstorms*. Schools everywhere are now adding makerspaces and coding classes, offering students opportunities that few could have imagined in 1980. Seymour’s work should be seen as the intellectual inspiration for the Maker Movement and the Coding Movement.

Yet if Seymour were alive today, I have no doubt that he would be very frustrated with the ways that making and coding are being introduced in schools. Seymour would view most of the current initiatives as “technocentric” (a term that Seymour popularized). That is, the initiatives focus too much on helping children develop technical skills: how to use a 3d printer, how to define an algorithm, how to write efficient computer code.

For Seymour, technical skills were never the goal. In the Introduction to *Mindstorms*, he wrote: “My central focus is not on the machine but on the mind.” Seymour was certainly interested in machines and new technologies, but only insofar as they could support learning or lead to new insights about learning.

A significant portion of *Mindstorms* focuses on logo, the first programming language designed specifically for children. But at the core of *Mindstorms* are Seymour’s ideas about education and learning, not technical issues. In the book, he lays the intellectual foundation for an educational theory that he later named “constructionism.” The theory builds on the work of Jean Piaget, the great child-development pioneer, who Seymour had collaborated with in the early 1960s. Piaget’s great insight was that knowledge is not delivered from teacher to learner; rather, children are constantly constructing knowledge through their everyday interactions with people and objects around them. Seymour’s constructionism theory adds a second type of construction, arguing that children construct knowledge most effectively when they are actively engaged in constructing things in the world. As children construct things in the world, they construct new ideas and theories in their minds, which motivates them to construct new things in the world, and on and on.

Seymour saw rich learning opportunities in all different types of "construction" activities: building sand castles on the beach, writing stories in a diary, drawing pictures in a sketchbook. Why was Seymour so interested in computational technologies? Because he recognized that computational technologies can greatly expand the range of what and how children create. With computers, children can create things that move, interact, and change over time, such as animations, simulations, and interactive games. In the process, children can gain new insights into the workings of dynamic systems in the world around them—including the workings of their own minds. In addition, computers enable children to modify, duplicate, document, and share their creations in ways they never could before, providing new ways for them to explore and understand the creative process.

It continually frustrated Seymour that people seemed to hear only part of his message, often focusing on the technology at the expense of the ideas. In an article Seymour wrote twenty years after the publication of *Mindstorms*, he lamented that the three parts of the book's subtitle—*Children, Computers, and Powerful Ideas*—had not been equally appreciated: "Most of the many educators who found inspiration and affirmation in the book (as well as those who hated it) discussed it as if it were about children and computers, as if the third term was there as a sound bite, the kind of shibboleth that pervades the discourse of technology in education. I did not mean it to be that: I actually thought I was writing a book about ideas!"

Of course, proponents of today's learn-to-code initiatives would argue that they, too, are interested in ideas, not just technical skills. Many of them frame their work around the idea of "computational thinking"—aiming to introduce children to problem-solving strategies that come from the field of computer science but are applicable across many other domains. Learning problem-solving strategies is certainly valuable. But Seymour had a bigger, broader vision. He wanted to support children not only in developing their thinking but also in developing their voice.

Seymour saw the computer not just as a problem-solving tool but as an expressive medium. He believed that learning to program was analogous to learning to write, providing children with new ways of organizing and expressing their ideas. Seymour wanted to help all children, from all backgrounds, have opportunities to express and share their ideas so that they could be full and active participants in society.

And what would Seymour think about today's widespread excitement about introducing artificial intelligence (ai) in k–12 education? In chapter 7 of *Mindstorms*, Seymour described how ai research was an important source of inspiration for his work in education and learning. Again, Seymour's writings in 1980 seem prescient. But, still, I believe that Seymour would be very frustrated with the approach of today's ai-in-education efforts. Seymour was interested in applying ideas from ai to engage children in thinking about their own thinking—and learning about their own learning. Most of today's ai-in-education initiatives have a very different set of goals, focusing on the use of machine intelligence, rather than the understanding of human intelligence.

So what has limited the spread of Seymour's ideas? Why haven't his ideas had an even bigger influence? One challenge is the general resistance to change in educational systems. But there are also challenges in the ways that Seymour's ideas have been communicated, supported, and interpreted. For instance, Seymour often used examples from mathematics and computer science, leading some people to interpret his ideas too narrowly, even though Seymour intended his ideas to apply across all disciplines. And Seymour often highlighted examples of what children can create and learn on their own, leading some people to try to implement his ideas without paying enough attention to the role of teachers, parents, and peers in the learning process.

When I think about the spread of Seymour's ideas, I like to think of Seymour planting ideas much as a farmer plants seeds. Some were mathematical ideas, some were pedagogical ideas, some were technological ideas, some were epistemological ideas. Some of Seymour's ideas spread like wildflowers around the world. Some took root in a few places, but not in others. Some of his seeds still lie dormant in the ground.

I'm part of a community of researchers and educators who continue to believe deeply in Seymour's ideas and vision. We're dedicated to nurturing the seeds that Seymour sowed, and we're trying to provide the right conditions for them to grow. I worked closely with Seymour for many years, first as an mit graduate student and then as a faculty colleague. My research today continues to be deeply influenced by Seymour's ideas, and I continue to explore ways to support his ideas in different learning contexts. As a framework for my work, I developed a set of four guiding principles that are all inspired by Seymour's ideas:

Projects. Seymour provocatively argued for "projects over problems." Of course, Seymour understood the importance of problem solving. But he believed that people learn to solve problems (and learn new concepts and strategies) most effectively while they are actively engaged in meaningful projects. Too often, schools start by teaching concepts to students, and only then give students a chance to work on projects. Seymour argued that it is best for children to learn new ideas *through* working on projects, not *before* working on projects.

Passion. In the Preface to *Mindstorms*, Seymour described how his childhood fascination with gears provided him a way to explore important mathematical concepts. For me, the most important and memorable line in the Preface is Seymour's statement: "I fell in love with the gears." Seymour understood the importance of learners building on their interests and passions. He knew that people will work longer and harder, and make deeper connections to ideas, when they're working on projects that they're passionate about. Seymour once said: "Education has very little to do with explanation, it has to do with engagement, with falling in love with the material."

Peers. In the final chapter of *Mindstorms*, titled "Images of the Learning Society," Seymour wrote about the Brazilian samba schools, where people come together to create music and dance routines for the annual carnival festival. What intrigued Seymour most was the way that samba schools bring together people of all different ages and all different levels of experience: children and adults, novices and experts, all working together, learning with and from one another. For Seymour, this type of peer-based learning was at the core of a Learning Society. Seymour, like Piaget before him, is sometimes criticized for focusing too much on the individual learner. But his writings on samba schools show another side of Seymour. The technologies of the *Mindstorms* era weren't quite ready for the type of peer-based online collaboration that we see today, but Seymour recognized the importance of the social dimension of learning.

Play. Often, people associate play with laughter and fun. But for Seymour, play meant more than that. It involved experimenting, taking risks, testing the boundaries, and iteratively adapting when things go wrong. Seymour sometimes referred to this process as "hard fun." He recognized that children don't want things to be easy: They're willing to work very hard on things that they find meaningful. And Seymour didn't just encourage play and hard fun for others; he lived it himself. He was always playing with ideas, wrestling with ideas, experimenting with ideas. I never met anyone who was, at once, so playful and so serious about ideas.

Seymour was not just a mathematician, educator, philosopher, and computer scientist. He was also an activist. From his teenage years in the 1940s, when he battled apartheid in his native South Africa, Seymour was always looking to bring about change. When he initiated new educational projects, he often exclaimed: "Now is the time!" He was a person of big ideas who wanted big changes. In *Mindstorms* and other writings, he advocated for revolutionary changes in the ways we think about children, learning, and education. In the forty years since *Mindstorms*, the changes have been more evolutionary than revolutionary. Many of Seymour's dreams have been unfulfilled.

But I believe that the environment is becoming more fertile for many of the seeds that Seymour sowed. There is a growing recognition that educational systems are not meeting the needs of today's fast-changing society. More educational reformers are advocating for changes that are aligned with Seymour's ideas, such as providing children with more opportunities to explore, experiment, and express themselves, so that they can develop as creative thinkers. The changes are evolutionary, not revolutionary, but the long-term trends are heading in the direction of Seymour's vision.

As you read *Mindstorms*, don't get distracted by the details of the 1980-era technologies that are described in the book. Rather, think about the ways that Seymour's ideas can be integrated into today's discussions about educational strategies and policies. And think about what you might do to nurture the seeds that Seymour sowed.

MITCHEL RESNICK
LEGO Papert Professor of Learning Research
MIT Media Lab

PREFACE

The Gears of My Childhood

BEFORE I WAS TWO YEARS OLD I HAD DEVELOPED AN INTENSE involvement with automobiles. The names of car parts made up a very substantial portion of my vocabulary: I was particularly proud of knowing about the parts of the transmission system, the gearbox, and most especially the differential. It was, of course, many years later before I understood how gears work; but once I did, playing with gears became a favorite pastime. I loved rotating circular objects against one another in gearlike motions and, naturally, my first "erector set" project was a crude gear system.

I became adept at turning wheels in my head and at making chains of cause and effect: "This one turns this way so that must turn that way so..." I found particular pleasure in such systems as the differential gear, which does not follow a simple linear chain of causality since the motion in the transmission shaft can be distributed in many different ways to the two wheels depending on what resistance they encounter. I remember quite vividly my excitement at discovering that a system could be lawful and completely comprehensible without being rigidly deterministic.

I believe that working with differentials did more for my mathematical development than anything I was taught in elementary school. Gears, serving as models, carried many otherwise abstract ideas into my head. I clearly remember two examples from school math. I saw multiplication tables as gears, and my first brush with equations in two variables (e.g., $3x + 4y = 10$) immediately evoked the differential. By the time I had made a mental gear model of the relation between x and y, figuring how many teeth each gear needed, the equation had become a comfortable friend.

Many years later when I read Piaget, this incident served me as a model for his notion of assimilation, except I was immediately struck by the fact that his discussion does not do full justice to his own idea. He talks almost entirely about cognitive aspects of assimilation. But there is also an affective component. Assimilating equations to gears certainly is a powerful way to bring old knowledge to bear on a new object. But it does more as well. I am sure that such assimilations helped to endow mathematics, for me, with a positive affective tone that can be traced back to my infantile experiences with cars. I believe Piaget really agrees. As I came to know him personally I understood that his neglect of the affective comes more from a modest sense that little is known about it than from an arrogant sense of its irrelevance. But let me return to my childhood.

One day I was surprised to discover that some adults—even *most* adults—did not understand or even care about the magic of the gears. I no longer think much about gears, but I have never turned away from the questions that started with that discovery: How could what was so simple for me be incomprehensible to other people? My proud father suggested "being clever" as an explanation. But I was painfully aware that some people who could not understand the differential could easily do things I found much more difficult. Slowly I began to formulate what I still consider the fundamental fact about learning: Anything is easy if you can assimilate it to your collection of models. If you can't, anything can be painfully difficult. Here too I was developing a way of thinking that would be resonant with Piaget's. *The understanding of learning must be genetic.* It must refer to the genesis of knowledge. What an individual can learn, and how he learns it, depends on what models he has available. This raises, recursively, the question of how he learned these models. Thus the "laws of learning" must be about how intellectual structures grow out of one another and about how, in the process, they acquire both logical and emotional form.

This book is an exercise in an applied genetic epistemology expanded beyond Piaget's cognitive emphasis to include a concern with the affective. It develops a new perspective for education research focused on creating the conditions under which intellectual models will take root. For the last two decades this is what I have been trying to do. And in doing so I find myself frequently reminded of several aspects of my encounter with the differential gear. First, I remember that no one told me to learn about differential gears. Second, I remember that there was *feeling, love,* as well as understanding in my relationship with gears. Third, I remember that my first encounter with them was in my second year. If any "scientific" educational psychologist had tried to "measure" the effects of this encounter, he would probably have failed. It had profound consequences but, I conjecture, only very many years later. A "pre- and post-" test at age two would have missed them.

Piaget's work gave me a new framework for looking at the gears of my childhood. The gear can be used to illustrate many powerful "advanced" mathematical ideas, such as groups or relative motion. But it does more than this. As well as connecting with the formal knowledge of mathematics, it also connects with the "body knowledge," the sensorimotor schemata of a child. You can *be* the gear, you can understand how it turns by projecting yourself into its place and turning with it. It is this double relationship—both abstract and sensory—that gives the gear the power to carry powerful mathematics into the mind. In a terminology I shall develop in later chapters, the gear acts here as a *transitional object.*

A modern-day Montessori might propose, if convinced by my story, to create a gear set for children. Thus every child might have the experience I had. But to hope for this would be to miss the essence of the story. *I fell in love with the gears.* This is something that cannot be reduced to purely "cognitive" terms. Something very personal happened, and one cannot assume that it would be repeated for other children in exactly the same form.

My thesis could be summarized as: What the gears cannot do the computer might. The computer is the Proteus of machines. Its essence is its universality, its power to simulate. Because it can take on a thousand forms and can serve a thousand functions, it can appeal to a thousand tastes. This book is the result of my own attempts over the past decade to turn computers into instruments flexible enough so that many children can each create for themselves something like what the gears were for me.

INTRODUCTION

Computers for Children

JUST A FEW YEARS AGO PEOPLE THOUGHT OF COMPUTERS AS EXPENSIVE and exotic devices. Their commercial and industrial uses affected ordinary people, but hardly anyone expected computers to become part of day-to-day life. This view has changed dramatically and rapidly as the public has come to accept the reality of the personal computer, small and inexpensive enough to take its place in every living room or even in every breast pocket. The appearance of the first rather primitive machines in this class was enough to catch the imagination of journalists and produce a rash of speculative articles about life in the computer-rich world to come. The main subject of these articles was what people will be able to do with their computers. Most writers emphasized using computers for games, entertainment, income tax, electronic mail, shopping, and banking. A few talked about the computer as a teaching machine.

This book too poses the question of what will be done with personal computers, but in a very different way. I shall be talking about how computers may affect the way people think and learn. I begin to characterize my perspective by noting a distinction between two ways computers might enhance thinking and change patterns of access to knowledge.

Instrumental uses of the computer to help people think have been dramatized in science fiction. For example, as millions of "Star Trek" fans know, the starship *Enterprise* has a computer that gives rapid and accurate answers to complex questions posed to it. But no attempt is made in "Star Trek" to suggest that the human characters aboard think in ways very different from the manner in which people in the twentieth century think. Contact with the computer has not, as far as we are allowed to see in these episodes, changed how these people think about themselves or how they approach problems. In this book I discuss ways in which the computer presence could contribute to mental processes not only instrumentally but in more essential, conceptual ways, influencing how people think even when they are far removed from physical contact with a computer (just as the gears shaped my understanding of algebra although they were not physically present in the math class). It is about an end to the culture that makes science and technology alien to the vast majority of people. Many cultural barriers impede children from making scientific knowledge their own. Among these barriers the most visible are the physically brutal effects of deprivation and isolation. Other barriers are more political. Many children who grow up in our cities are surrounded by the artifacts of science but have good reason to see them as belonging to "the others"; in many cases they are perceived as belonging to the social enemy. Still other obstacles are more abstract, though ultimately of the same nature. Most branches of the most sophisticated modern culture of Europe and the United States are so deeply "mathophobic" that many privileged children are as effectively (if more gently) kept from appropriating science as their own. In my vision, space-age objects, in the form of small computers, will cross these cultural barriers to enter the private worlds of children everywhere. They will do so not as mere physical objects. This book is about how computers can be carriers of powerful ideas and of the seeds of cultural change, how they can help people form new relationships with knowledge that cuts across the traditional lines separating humanities from sciences and knowledge of the self from both of these. It is about using computers to challenge current beliefs about who can understand what and at what age. It is about using computers to question standard assumptions in developmental psychology and in the psychology of aptitudes and attitudes. It is about whether personal computers and the cultures in which they are used will continue to be the creatures of "engineers" alone or whether we can construct intellectual environments in which people who today think of themselves as "humanists" will feel part of, not alienated from, the process of constructing computational cultures.

But there is a world of difference between what computers can do and what society will choose to do with them. Society has many ways to resist fundamental and threatening change. Thus, this book is about facing choices that are ultimately political. It looks at some of the forces of change and of reaction to those forces that are called into play as the computer presence begins to enter the politically charged world of education.

Much of the book is devoted to building up images of the role of the computer that are very different from current stereotypes. All of us, professionals as well as laymen, must consciously break the habits we bring to thinking about the computer. Computation is in its infancy. It is hard to think about computers of the future without projecting onto them the properties and the limitations of those we think we know today. And nowhere is this more true than in imagining how computers can enter the world of education. It is not true to say that the image of a child's relationship with a computer I shall develop here goes far beyond what is common in today's schools. My image does not go beyond: It goes in the opposite direction.

In many schools today, the phrase "computer-aided instruction" means making the computer teach the child. One might say the *computer is being used to program* the child. In my vision, *the child programs the computer* and, in doing so, both acquires a sense of mastery over a piece of the most modern and powerful technology and establishes an intimate contact with some of the deepest ideas from science, from mathematics, and from the art of intellectual model building.

I shall describe learning paths that have led hundreds of children to becoming quite sophisticated programmers. Once programming is seen in the proper perspective, there is nothing very surprising about the fact that this should happen. Programming a computer means nothing more or less than communicating to it in a language that it and the human user can both "understand." And learning languages is one of the things children do best. Every normal child learns to talk. Why then should a child not learn to "talk" to a computer?

There are many reasons why someone might expect it to be difficult. For example, although babies learn to speak their native language with spectacular ease, most children have great difficulty learning foreign languages in schools and, indeed, often learn the written version of their own language none too successfully. Isn't learning a computer language more like the difficult process of learning a foreign written language than the easy one of learning to speak one's own language? And isn't the problem further compounded by all the difficulties most people encounter learning mathematics?

Two fundamental ideas run through this book. The first is that it is possible to design computers so that learning to communicate with them can be a natural process, more like learning French by living in France than like trying to learn it through the unnatural process of American foreign-language instruction in classrooms. Second, learning to communicate with a computer may change the way other learning takes place. The computer can be a mathematics-speaking and an alphabetic-speaking entity. We are learning how to make computers with which children love to communicate. When this communication occurs, children learn mathematics as a living language. Moreover, mathematical communication and alphabetic communication are thereby both transformed from the alien and therefore difficult things they are for most children into natural and therefore easy ones. The idea of "talking mathematics" to a computer can be generalized to a view of learning mathematics in "Mathland"—that is to say, in a context which is to learning mathematics what living in France is to learning French.

In this book the Mathland metaphor will be used to question deeply engrained assumptions about human abilities. It is generally assumed that children cannot learn formal geometry until well into their school years and that most cannot learn it too well even then. But we can quickly see that these assumptions are based on extremely weak evidence by asking analogous questions about the ability of children to learn French. If we had to base our opinions on observation of how poorly children learned French in American schools, we would have to conclude that most people were incapable of mastering it. But we know that all normal children would learn it very easily if they lived in France. My conjecture is that much of what we now see as too "formal" or "too mathematical" will be learned just as easily when children grow up in the computer-rich world of the very near future.

I use the examination of our relationship with mathematics as a thematic example of how technological and social processes interact in the construction of ideas about human capacities. And mathematical examples will also help to describe a theory of how learning works and of how it goes wrong.

I take from Jean Piaget[1] a model of children as builders of their own intellectual structures. Children seem to be innately gifted learners, acquiring long before they go to school a vast quantity of knowledge by a process I call "Piagetian learning," or "learning without being taught." For example, children learn to speak, learn the intuitive geometry needed to get around in space, and learn enough of logic and rhetorics to get around parents—all this without being "taught." We must ask why some learning takes place so early and spontaneously while some is delayed many years or does not happen at all without deliberately imposed formal instruction.

If we really look at the "child as builder," we are on our way to an answer. All builders need materials to build with. Where I am at variance with Piaget is in the role I attribute to the surrounding cultures as a source of these materials. In some cases the culture supplies them in abundance, thus facilitating constructive Piagetian learning. For example, the fact that so many important things (knives and forks, mothers and fathers, shoes and socks) come in pairs is a "material" for the construction of an intuitive sense of number. But in many cases where Piaget would explain the slower development of a particular concept by its greater complexity or formality, I see the critical factor as the relative poverty of the culture in those materials that would make the concept simple and concrete. In yet other cases, the culture may provide materials but block their use. In the case of formal mathematics, there is both a shortage of formal materials and a cultural block as well. The mathophobia endemic in contemporary culture blocks many people from learning anything they recognize as "math," although they may have no trouble with mathematical knowledge they do not perceive as such.

cycle without creating a dependence on machines. My discussion differs from most arguments about "nature versus nurture" in two ways. I shall be much more specific both about what kinds of nurturance are needed for intellectual growth and about what can be done to create such nurturance in the home as well as in the wider social context.

Thus this book is really about how a culture, a way of thinking, an idea comes to inhabit a young mind. I am suspicious of thinking about such problems too abstractly, and I shall write here with particular restricted focus. I shall in fact concentrate on those ways of thinking that I know best. I begin by looking at what I know about my own development. I do this in all humility, without any implication that what I have become is what everyone should become. But I think that the best way to understand learning is first to understand specific, well-chosen cases and then to worry afterward about how to generalize from this understanding. You can't think seriously about thinking without thinking about thinking about something. And the something I know best how to think about is mathematics. When in this book I write of mathematics, I do not think of myself as writing for an audience of mathematicians interested in mathematical thinking for its own sake. My interest is in universal issues of how people think and how they learn to think.

When I trace how I came to be a mathematician, I see much that was idiosyncratic, much that could not be duplicated as part of a generalized vision of education reform. And I certainly don't think that we would want everyone to become a mathematician. But I think that the kind of pleasure I take in mathematics should be part of a general vision of what education should be about. If we can grasp the essence of one person's experiences, we may be able to replicate its consequences in other ways, and in particular this consequence of finding beauty in abstract things. And so I shall be writing quite a bit about mathematics. I give my apologies to readers who hate mathematics, but I couple that apology with an offer to help them learn to like it a little better—or at least to change their image of what "speaking mathematics" can be all about.

In the Foreword of this book I described how gears helped mathematical ideas to enter my life. Several qualities contributed to their effectiveness. First, they were part of my natural "landscape," embedded in the culture around me. This made it possible for me to find them myself and relate to them in my own fashion. Second, gears were part of the world of adults around me, and through them I could relate to these people. Third, I could use my body to think about the gears. I could feel how gears turn by imagining by body turning. This made it possible for me to draw on my "body knowledge" to think about gear systems. And finally, because, in a very real sense, the relationship between gears contains a great deal of mathematical information, I could use the gears to think about formal systems. I have described the way in which the gears served as an "object-to-think-with." I made them that for myself in my own development as a mathematician. The gears have also served me as an object-to-think-with in my work as an educational researcher. My goal has been the design of other objects that children can make theirs for themselves and in their own ways. Much of this book will describe my path through this kind of research. I begin by describing one example of a constructed computational "object-to-think-with." This is the "Turtle."[3]

The central role of the Turtle in this book should not be taken to mean that I propose it as a panacea for all educational problems. I see it as a valuable educational object, but its principal role here is to serve as a model for other objects, yet to be invented. My interest is in the process of invention of "objects-to-think-with," objects in which there is an intersection of cultural presence, embedded knowledge, and the possibility for personal identification.

The Turtle is a computer-controlled cybernetic animal. It exists within the cognitive minicultures of the "LOGO environment," LOGO being the computer language in which communication with the Turtle takes place. The Turtle serves no other purpose than of being good to program and good to think with. Some Turtles are abstract objects that live on computer screens. Others, like the floor Turtles shown in the frontispiece are physical objects that can be picked up like any mechanical toy. A first encounter often begins by showing the child how a Turtle can be made to move by typing commands at a keyboard. FORWARD 100 makes the Turtle move in a straight line a distance of 100 Turtle steps of about a millimeter each. Typing RIGHT 90 causes the Turtle to pivot in place through 90 degrees. Typing PENDOWN causes the Turtle to lower a pen so as to leave a visible trace of its path while PENUP instructs it to raise the pen. Of course the child needs to explore a great deal before gaining mastery of what the numbers mean. But the task is engaging enough to carry most children through this learning process.

The idea of programming is introduced through the metaphor of teaching the Turtle a new word. This is simply done, and children often begin their programming experience by programming the Turtle to respond to new commands invented by the child such as SQUARE or TRIANGLE or SQ or TRI or whatever the child wishes, by drawing the appropriate shapes. New commands once defined can be used to define others. For example just as the house in Figure 1 is built out of a triangle and a square, the program for drawing it is built out of the commands for drawing a square and a triangle. Figure 1 shows four steps in the evolution of this program. From these simple drawings the young programmer can go on in many different directions. Some work on more complex drawings, either figural or abstract. Some abandon the use of the Turtle as a drawing instrument and learn to use its touch sensors to program it to seek out or avoid objects.[4] Later children learn that the computer can be programmed to make music as well as move Turtles and combine the two activities by programming Turtles to dance. Or they can move on from floor Turtles to "screen Turtles," which they program to draw moving pictures in bright colors. The examples are infinitely varied, but in each the child is learning how to exercise control over an exceptionally rich and sophisticated "micro-world."

Readers who have never seen an interactive computer display might find it hard to imagine where this can lead. As a mental exercise they might like to imagine an electronic sketchpad, a computer graphics display of the not-too-distant future. This is a television screen that can display moving pictures in color. You can also "draw" on it, giving it instructions, perhaps by typing, perhaps by speaking, or perhaps by pointing with a wand. On request, a palette of colors could appear on the screen. You can choose a color by pointing at it with the wand. Until you change your choice, the wand draws in that color. Up to this point the distinction from traditional art materials may seem slight, but the distinction becomes very real when you begin to think about editing the drawing. You can "talk to your drawing" in computer language. You can "tell" it to replace this color with that. Or set a drawing in motion. Or make two copies and set them in counterrotating motion. Or replace the color palette with a sound palette and "draw" a piece of music. You can file your work in computer memory and retrieve it at your pleasure or have it delivered into the memory of any of the many millions of other computers linked to the central communication network for the pleasure of your friends.

That all this would be fun needs no argument. But it is more than fun. Very powerful kinds of learning are taking place. Children working with an electronic sketchpad are learning a language for talking about shapes and fluxes of shapes, about velocities and rates of change, about processes and procedures. They are learning to speak mathematics and acquiring a new image of themselves as mathematicians.

In my description of children working with Turtles, I implied that children can learn to program. For some readers this might be tantamount to the suspension of disbelief required when we enter a theater to watch a play. For them programming is a complex and marketable skill acquired by some mathematically gifted adults. But my experience is very different. I have seen hundreds of elementary schoolchildren learn very easily to program, and evidence is accumulating to indicate that much younger children could do so as well. The children in these studies are not exceptional, or rather, they are exceptional in every conceivable way. Some of the children were highly successful in school, some were diagnosed as emotionally or cognitively disabled. Some of the children were so severely afflicted by cerebral palsy that they had never purposefully manipulated physical objects. Some of them had expressed their talents in "mathematical" forms, some in "verbal" forms, and some in artistically "visual" or in "musical" forms.

week or two on vacation in France than like living there. But like children who have spent a vacation with foreign-speaking cousins, they were clearly on their way to "speaking computer."

When I have thought about what these studies mean, I am left with two clear impressions. First, that all children will, under the right conditions, acquire a proficiency with programming that will make it one of their more advanced intellectual accomplishments. Second, that the "right conditions" are very different from the kind of access to computers that is now becoming established as the norm in schools. The conditions necessary for the kind of relationships with a computer that I will be writing about in this book require more and freer access to the computer than educational planners currently anticipate. And they require a kind of computer language and a learning environment around that language very different from those the schools are now providing. They even require a kind of computer rather different from those that the schools are currently buying.

It will take most of this book for me to convey some sense of the choices among computers, computer languages, and more generally, among computer cultures, that influence how well children will learn from working with computation and what benefits they will get from doing so. But the question of the *economic* feasibility of free access to computers for every child can be dealt with immediately. In doing so I hope to remove any doubts readers may have about the "economic realism" of the "vision of education" I have been talking about.

My vision of a new kind of learning environment demands free contact between children and computers. This could happen because the child's family buys one or a child's friends have one. For purposes of discussion here (and to extend our discussion to all social groups) let us assume that it happens because schools give every one of their students his or her own powerful personal computer. Most "practical" people (including parents, teachers, school principals, and foundation administrators) react to this idea in much the same way: "Even if computers could have all the effects you talk about, it would still be impossible to put your ideas into action. Where would the money come from?"

What these people are saying needs to be faced squarely. They are wrong. Let's consider the cohort of children who will enter kindergarten in the year 1987, the "Class of 2000," and let's do some arithmetic. The direct public cost of schooling a child for thirteen years, from kindergarten through twelfth grade is over $20,000 today (and for the class of 2000, it may be closer to $30,000). A conservatively high estimate of the cost of supplying each of these children with a personal computer with enough power for it to serve the kinds of educational ends described in this book, and of upgrading, repairing, and replacing it when necessary would be about $1,000 per student, distributed over thirteen years in school. Thus, "computer costs" for the class of 2000 would represent only about 5 percent of the total public expenditure on education, and this would be the case even if nothing else in the structure of educational costs changed because of the computer presence. But in fact computers in education stand a good chance of making other aspects of education cheaper. Schools might be able to reduce their cycle from thirteen years to twelve years; they might be able to take advantage of the greater autonomy the computer gives students and increase the size of classes by one or two students without decreasing the personal attention each student is given. Either of these two moves would "recuperate" the computer cost.

My goal is not educational economies: It is not to use computation to shave a year off the time a child spends in an otherwise unchanged school or to push an extra child into an elementary school classroom. The point of this little exercise in educational "budget balancing" is to do something to the state of mind of my readers as they turn to the first chapter of this book. I have described myself as an educational utopian—not because I have projected a future of education in which children are surrounded by high technology, but because I believe that certain uses of very powerful computational technology and computational ideas can provide children with new possibilities for learning, thinking, and growing emotionally as well as cognitively. In the chapters that follow, I shall try to give you some idea of these possibilities, many of which are dependent on a computer-rich future, a future where a computer will be a significant part of every child's life. But I want my readers to be very clear that what is “utopian” in my vision and in this book is a particular way of using computers, of forging new relationships between computers and people—that the computer will be there to be used is simply a conservative premise.