

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

MINDSTORMS

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 cut 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 developmen-

tal 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 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

MINDSTORMS

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 ex-

tremely 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 Piaget¹ 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

MINDSTORMS

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.

We shall see again and again that the consequences of mathophobia go far beyond obstructing the learning of mathematics and science. They interact with other endemic “cultural toxins,” for example, with popular theories of aptitudes, to contaminate peoples’ images of themselves as learners. Difficulty with school math is often the first step of an invasive intellectual process that leads us all to define ourselves as bundles of aptitudes and ineptitudes, as being “mathematical” or “not mathematical,” “artistic” or “not artistic,” “musical” or “not musical,” “profound” or “superficial,” “intelligent” or “dumb.” Thus deficiency becomes identity and learning is transformed from the early child’s free exploration of the world to a chore beset by insecurities and self-imposed restrictions.

Two major themes—that children can learn to use computers in a masterful way, and that learning to use computers can change the way they learn everything else—have shaped my research agenda on computers and education. Over the past ten years I have had the good fortune to work with a group of colleagues and students at MIT (the LOGO² group in the Artificial Intelligence Laboratory) to create environments in which children can learn to communicate with computers. The metaphor of imitating the way the child learns to talk has been constantly with us in this work and has led to a vision of education and of education research very different from the traditional ones. For people in the teaching professions, the word “education” tends to evoke “teaching,” particularly classroom teaching. The goal of education research tends therefore to be focused on how to improve classroom teaching. But if, as I have stressed here, the model of successful learning is the way a child learns to talk, a process that takes place without deliberate and organized teaching, the goal set is very different. I see the classroom as an artificial and inefficient learning environment that society has been forced to invent because its informal environments fail in certain essential learning domains, such as writing or gram-

mar or school math. I believe that the computer presence will enable us to so modify the learning environment outside the classrooms that much if not all the knowledge schools presently try to teach with such pain and expense and such limited success will be learned, as the child learns to talk, painlessly, successfully, and without organized instruction. This obviously implies that schools as we know them today will have no place in the future. But it is an open question whether they will adapt by transforming themselves into something new or wither away and be replaced.

Although technology will play an essential role in the realization of my vision of the future of education, my central focus is not on the machine but on the mind, and particularly on the way in which intellectual movements and cultures define themselves and grow. Indeed, the role I give to the computer is that of a *carrier* of cultural “germs” or “seeds” whose intellectual products will not need technological support once they take root in an actively growing mind. Many if not all the children who grow up with a love and aptitude for mathematics owe this feeling, at least in part, to the fact that they happened to acquire “germs” of the “math culture” from adults, who, one might say, knew how to speak mathematics, even if only in the way that Moliere had M. Jourdain speak prose without knowing it. These “math-speaking” adults do not necessarily know how to solve equations; rather, they are marked by a turn of mind that shows up in the logic of their arguments and in the fact that for them to play is often to play with such things as puzzles, puns, and paradoxes. Those children who prove recalcitrant to math and science education include many whose environments happened to be relatively poor in math-speaking adults. Such children come to school lacking elements necessary for the easy learning of school math. School has been unable to supply these missing elements, and, by forcing the children into learning situations doomed in advance, it generates powerful negative feelings about mathematics and perhaps about learning in general. Thus is set up a vicious self-perpetuating cycle. For these same children will one day be parents and will not only fail to pass on mathematical germs but will almost certainly infect their children with the opposing and intellectually destructive germs of mathophobia.

Fortunately it is sufficient to break the self-perpetuating cycle at

MINDSTORMS

one point for it to remain broken forever. I shall show how computers might enable us to do this, thereby breaking the vicious 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.”⁸

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

MINDSTORMS

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

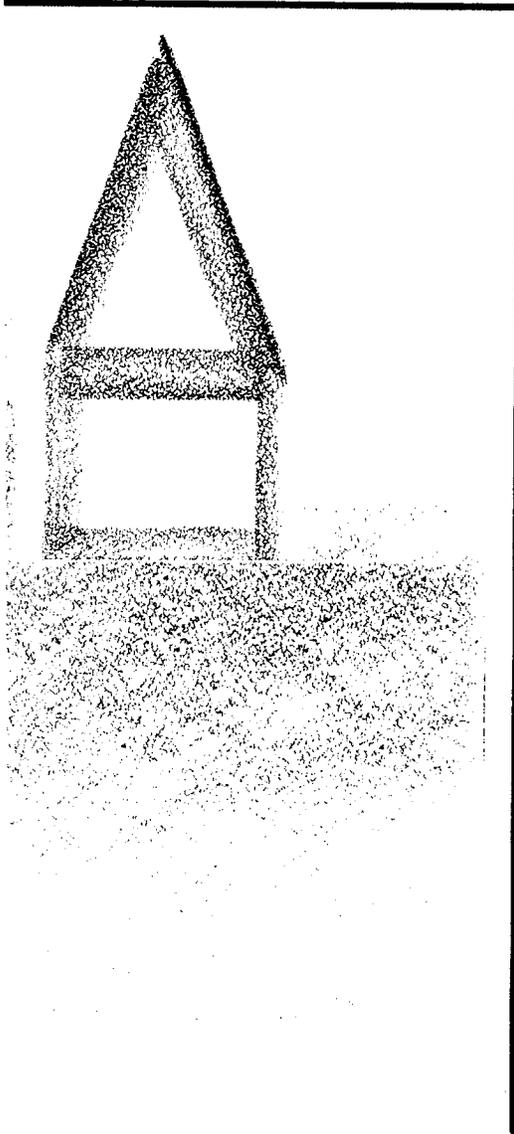
speaking, or perhaps by pointing with a wand. On request, a palette of colors could appear on the screen. You can choose a color by pointing at it with the wand. Until you change your choice, the wand draws in that color. Up to this point the distinction from traditional art materials may seem slight, 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.

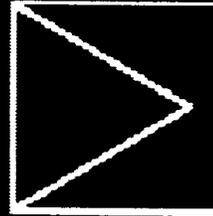
In my description of children working with Turtles, I implied that children can learn to program. For some readers this might be tantamount to the suspension of disbelief required when we enter a theater to watch a play. For them programming is a complex and marketable skill acquired by some mathematically gifted adults. But my experience is very different. I have seen hundreds of elementary school children 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.

Figure 1

A Plan

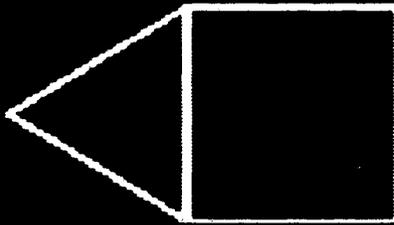


A Bug

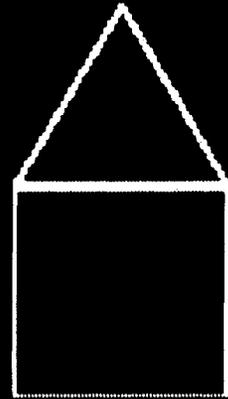


TO HOUSE
SQ
TRI

And a Compromise



TO HOUSE
SQ
RIGHT 30
TR1
END



TO HOUSE
RIGHT 90
SQ
RIGHT 30
TR1
END

MINDSTORMS

Of course these children did not achieve a fluency in programming that came close to matching their use of spoken language. If we take the Mathland metaphor seriously, their computer experience was more like learning French by spending a 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 2,000 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 emotion-

MINDSTORMS

ally 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.