Seymour Papert, Mindstorms - Children, Computers + Powerful Ideas, 1980

Chapter 7

Logo's Roots: Piaget and AI

THE READER has already met a variety of learning situations drawn together by a common set of ideas about what makes for effective learning. In this chapter we turn directly to these ideas and to the theoretical sources by which they are informed. Of these we focus on two: first, the Piagetian influence, and second, the influence of computational theory and artificial intelligence.

I have previously spoken of "Piagetian learning," the natural, spontaneous learning of people in interaction with their environment, and contrasted it with the curriculum-driven learning characteristic of traditional schools. But Piaget's contribution to my work has been much deeper, more theoretical and philosophical. In this chapter I will present a Piaget very different from the one most people have come to expect. There will be no talk of stages, no emphasis on what children at certain ages can or cannot learn to do. Rather I shall be concerned with Piaget the epistemologist, as his ideas have contributed toward the knowledge-based theory of learning that I have been describing, a theory that does not divorce the study of how mathematics is learned from the study of mathematics itself.

I think these epistemological aspects of Piaget's thought have been underplayed because up until now they offered no possibilities for action in the world of traditional education. But in a computer-rich educational environment, the educational environment of the

next decade, this will not be the case. In chapter 5 and in the development of the Turtle idea itself we saw examples of how an epistemological inquiry into what is fundamental in a sector of mathematics, the mathematics of differential systems, has already paid off in concrete, effective educational designs. The Piaget of the stage theory is essentially conservative, almost reactionary, in emphasizing what children cannot do. I strive to uncover a more revolutionary Piaget, one whose epistemological ideas might expand known bounds of the human mind. For all these years they could not do so for lack of a means of implementation, a technology which the mathetic computer now begins to make available.

The Piaget as presented in this chapter is new in another sense as well. He is placed in a theoretical framework drawn from a side of the computer world of which we have not spoken directly, but whose perspectives have been implicit throughout this book, that of artificial intelligence, or AI. The definition of artificial intelligence can be narrow or broad. In the narrow sense, AI is concerned with extending the capacity of machines to perform functions that would be considered intelligent if performed by people. Its goal is to construct machines and, in doing so, it can be thought of as a branch of advanced engineering. But in order to construct such machines, it is usually necesary to reflect not only on the nature of machines but on the nature of the intelligent functions to be performed.

For example, to make a machine that can be instructed in natural language, it is necessary to probe deeply into the nature of language. In order to make a machine capable of learning, we have to probe deeply into the nature of learning. And from this kind of research comes the broader definition of artificial intelligence: that of a cognitive science. In this sense, AI shares its domain with the older disciplines such as linguistics and psychology. But what is distinctive in AI is that its methodology and style of theorizing draw heavily on theories of computation. In this chapter we shall use this style of theorizing in several ways: first, to reinterpret Piaget; second, to develop the theories of learning and understanding that inform our design of educational situations; and third, in a somewhat more unusual way. The aim of AI is to give concrete form to ideas about thinking that previously might have seemed abstract, even metaphysical. It is this concretizing quality that has made ideas from AI so attractive to many contemporary psychologists. We propose to teach AI to children so that they, too, can think more concretely about mental processes. While psychologists use ideas from AI to build formal, scientific theories about mental processes, children use the same ideas in a more informal and personal way to think about themselves. And obviously I believe this to be a good thing in that the ability to articulate the processes of thinking enables us to improve them.

Piaget has described himself as an epistemologist. What does he mean by that? When he talks about the developing child, he is really talking as much about the development of knowledge. This statement leads us to a contrast between epistemological and psychological ways of understanding learning. In the psychological perspective, the focus is on the laws that govern the learner rather than on what is being learned. Behaviorists study reinforcement schedules, motivation theorists study drive, gestalt theorists study good form. For Piaget, the separation between the learning process and what is being learned is a mistake. To understand how a child learns number, we have to study number. And we have to study number in a particular way: We have to study the structure of number, a mathematically serious undertaking. This is why it is not at all unusual to find Piaget referring in one and the same paragraph to the behavior of small children and to the concerns of theoretical mathematicians. To make more concrete the idea of studying learning by focusing on the structure of what is learned, we look at a very concrete piece of learning from everyday life and see how different it appears from a psychological and from an epistemological perspective.

We will consider learning to ride a bicycle. If we did not know better riding a bicycle would seem to be a really remarkable thing. What makes it possible? One could pursue this question by studying the rider to find out what special attributes (speed of reaction, complexity of brain functioning, intensity of motivation) contribute to his performance. This inquiry, interesting though it might be, is irrelevant to the real solution to the problem. People can ride bicycles because the bicycle, once in motion, is inherently stable. A bicycle without a rider pushed off on a steep downgrade will not fall over; it will run indefinitely down the hill. The geometrical construction of the front fork ensures that if the bicycle leans to the left the wheel will rotate to the left, thus causing that bicycle to turn and produce a centrifugal force that throws the bicycle to the right, counteracting the tendency to fall. The bicycle without a rider balances perfectly well. With a novice rider it will fall. This is because the novice has the wrong intuitions about balancing and freezes the position of the bicycle so that its own corrective mechanism cannot work freely. Thus learning to ride does not mean learning to balance, it means learning not to unbalance, learning not to interfere.

What we have done here is understand a process of learning by acquiring deeper insight into what was being learned. Psychological principles had nothing to do with it. And just as we have understood how people ride bicycles by studying bicycles, Piaget has taught us that we should understand how children learn number through a deeper understanding of what number is.

Mathematicians interested in the nature of number have looked at the problem from different standpoints. One approach, associated with the formalists, seeks to understand number by setting up axioms to capture it. A second approach, associated with Bertrand Russell, seeks to define number by reducing it to something more fundamental, for example, logic and set theory. Although both of these approaches are valid, important chapters in the history of mathematics, neither casts light on the question of why number is learnable. But there is a school of mathematics that does do so, although this was not its intention. This is the structuralism of the Bourbaki school.[1] Bourbaki is a pseudonym taken by a group of French mathematicians who set out to articulate a uniform theory for mathematics. Mathematics was to be one, not a collection of subdisciplines each with its own language and line of development. The school moved in this direction by recognizing a number of building blocks that it called the "mother structures." These structures have something in common with our idea of microworlds. Imagine a microworld in which things can be ordered but have no other properties. The knowledge of how to work the world is, in terms of the Bourbaki school, the mother structure of order. A second microworld allows relations of proximity, and this is the mother structure of topology. A third has to do with combining entities to produce new entities; this is the algebraic microstructure. The Bourbaki school's unification of mathematics is achieved by seeing more complex structures, such as arithmetic, as combinations of simpler structures of which the most important are the three mother structures. This school had no intention of making a theory of learning. They intended their structural analysis to be a technical tool for mathematicians to use in their day-to-day work. But the theory of mother structures *is* a theory of learning. It is a theory of how number is learnable. By showing how the structure of arithmetic can be decomposed into simpler, but still meaningful and coherent, structures, the mathematicians are showing a mathetic pathway into numerical knowledge. It is not surprising that Piaget, who was explicitly searching for a theory of number that would explain its development in children, developed a similar, parallel set of constructs, and then, upon "discovering" the Bourbaki school was able to use its constructs to elaborate his own.

Piaget observed that children develop coherent intellectual structures that seemed to correspond very closely to the Bourbaki mother structures. For example, recall the Bourbaki structure of order; indeed, from the earliest ages, children begin to develop expertise in ordering things. The topological and algebraic mother structures have similar developmental precursors. What makes them learnable? First of all, each represents a coherent activity in the child's life that could in principle be learned and made sense of independent of the others.

Second, the knowledge structure of each has a kind of internal simplicity that Piaget has elaborated in his theory of *groupements*, and which will be discussed in slightly different terms later. Third, although these mother structures are independent, the fact that they are learned in parallel and that they share a common formalism are clues that they are mutually supportive; the learning of each facilitates the learning of the others.

Piaget has used these ideas to give an account of the development of a variety of domains of knowledge in terms of a coherent, lawful set of structures as processes within the child's mind. He describes these internal structures as always in interaction with the external world, but his theoretical emphasis has been the internal events. My perspective is more interventionist. My goals are education, not just understanding. So, in my own thinking I have placed a greater emphasis on two dimensions implicit but not elaborated in Piaget's own work: an interest in intellectual structures that could develop as opposed to those that actually at present do develop in the child, and the design of learning environments that are resonant with them. The Turtle can be used to illustrate both of these interests: first, the identification of a powerful set of mathematical ideas that we do not presume to be represented, at least not in a developed form, in children; second, the creation of a *transitional object,* the Turtle, that can exist in the child's environment and make contact with the ideas. As a mathematician I know that one of the most powerful ideas in the history of science was that of differential analysis. From Newton onward, the relationship between the local and the global pretty well set the agenda for mathematics. Yet this idea has had no place in the world of children, largely because traditional access to it depends on an infrastructure of formal, mathematical training. For most people, nothing is more natural than that the most advanced ideas in mathematics should be inaccessible to children. From the perspective I took from Piaget, we would expect to find connections. So we set out to find some. But finding the connections did not simply mean inventing a new kind of clever, "motivating" pedagogy. It meant a research agenda that included separating what was most powerful in the idea of *differential* from the accidents of inaccessible formalisms. The goal was then to connect these scientifically fundamental structures with psychologically powerful ones. And of course these were the ideas that underlay the Turtle circle, the physics microworlds, and the touch-sensor Turtle.

In what sense is the natural environment a source of microworlds, indeed a source for a network of microworlds? Let's narrow the whole natural environment to those things in it that may serve as a source for one specific microworld, a microworld of pairing, of one-to-one correspondence. Much of what children see comes in pairs: mothers and fathers, knives and forks, eggs and egg cups. And they, too, are asked to be active constructors of pairs. They are asked to sort socks, lay the table with one place setting for each person, and distribute candies. When children focus attention on pairs they are in a self-constructed microworld, a microworld of pairs, in the same sense as we placed our students in the microworlds of geometry and physics Turtles. In both cases the relevant microworld is stripped of complexity, is simple, graspable. In both cases the child is allowed to play freely with its elements. Although there are constraints on the materials, there are no constraints on the exploration of combinations. And in both cases the power of the environment is that it is "discovery rich."

Working with computers can make it more apparent that children construct their own personal microworlds. The story of Deborah at the end of chapter 4 is a good example. LOGO gave her the opportunity to construct a particularly tidy microworld, her "RIGHT 30 world." But she might have done something like this in her head without a computer. For example, she might have decided to understand directions in the real world in terms of a simple set of operations. Such intellectual events are not usually visible to observers, any more than my algebra teachers knew that I used gears to think about equations. But they can be seen if one looks closely enough. Robert Lawler, a member of the Massachusetts Institute of Technology LOGO group, demonstrated this most clearly in his doctoral research. Lawler set out to observe everything a six-year-old child, his daughter Miriam, did during a six-month period. The wealth of information he obtained allowed him to piece together a picture of the microstructure of Miriam's growing abilities. For example, during this period Miriam learned to add, and Lawler was able to show that this did not consist of acquiring one logically uniform procedure. A better model of her learning to add is that she brought into a working relationship a number of idiosyncratic microworlds, each of which could be traced to identifiable, previous experiences.

I have said that Piaget is an epistemologist, but have not elaborated on what kind. Epistemology is the theory of knowledge. The

term epistemology could, according to its etymology, be used to cover all knowledge about knowledge, but traditionally it has been used in a rather special way; that is, to describe the study of the conditions of validity of knowledge. Piaget's epistemology is concerned not with the validity of knowledge but with its origin and growth. He is concerned with the genesis and evolution of knowledge, and marks this fact by describing his field of study as "genetic epistemology." Traditional epistemology has often been taken as a branch of philosophy. Genetic epistemology works to assert itself as a science. Its students gather data and develop theories about how knowledge developed, sometimes focusing on the evolution of knowledge in history, sometimes on the evolution of knowledge in the individual. But it does not see the two realms as distinct: It seeks to understand relations between them. These relations can take different forms.

In the simplest case the individual development is parallel to the historical development, recalling the biologists' dictum, ontogeny recapitulates phylogeny. For example, children uniformly represent the physical world in an Aristotelian manner, thinking, for example, that forces act on position rather than on velocities. In other cases, the relation is more complex, indeed to the point of reversal. Intellectual structures that appear first in a child's development are sometimes characteristic not of early science but of the most modern science. So, for example, the mother structure topology appears very early in the child's development, but topology itself appeared as a mathematical subdiscipline only in modern times. Only when mathematics becomes sufficiently advanced is it able to discover its own origins.

In the early part of the twentieth century, formal logic was seen as synonymous with the foundation of mathematics. Not until Bourbaki's structuralist theory appeared do we see an internal development in mathematics that opened the field up to "remembering" its genetic roots. And through the work of genetic epistemology, this "remembering" puts mathematics in the closest possible relationship to the development of research about how children construct their reality.

Genetic epistemology has come to assert a set of homologies between the structures of knowledge and the structures of the mind that come into being to grasp this knowledge. Bourbaki's mother structures are not simply the elements that underly the concept of number; rather, homologies are found in the mind as it constructs number for itself. Thus, the importance of studying the structure of knowledge is not just to better understand the knowledge itself, but to understand the person.

Research on the structure of this dialectical process translates into the belief that neither people nor knowledge—including mathematics—can be fully grasped separately from the other, a belief that was eloquently expressed by Warren McCulloch, who, together with Norbert Wiener, should have credit for founding cybernetics. When asked, as a youth, what question would guide his scientific life, McCulloch answered: "What is a man so made that he can understand number and what is number so made that a man can understand it?"

For McCulloch as for Piaget, the study of people and the study of what they learn and think are inseparable. Perhaps paradoxically for some, research on the nature of that inseparable relationship has been advanced by the study of machines and the knowledge they can embody. And it is to this research methodology, that of artificial intelligence, that we now turn.

In artificial intelligence, researchers use computational models to gain insight into human psychology as well as reflect on human psychology as a source of ideas about how to make mechanisms emulate human intelligence. This enterprise strikes many as illogical: Even when the performance looks identical, is there any reason to think that underlying processes are the same? Others find it illicit: The line between man and machine is seen as immutable by both theology and mythology. There is a fear that we will dehumanize what is essentially human by inappropriate analogies between our "judgments" and those computer "calculations." I take these objections very seriously, but feel that they are based on a view of artificial intelligence that is more reductionist that anything I myself am interested in. A brief parable and some only half-humored reasoning by analogy express my own views on the matter.

Men have always been interested in flying. Once upon a time, scientists determined to understand how birds fly. First they watched them, hoping to correlate the motion of a bird's wings with its upward movement. Then they proceeded to experiment and found that when its feathers were plucked, a bird could no longer fly. Having thus determined that feathers were the organ of flight, the scientists then focused their efforts on microscopic and ultramicroscopic investigation of feathers in order to discover the nature of their flight-giving power.

In reality our current understanding of how birds fly did not come through a study narrowly focused on birds and gained nothing at all from the study of feathers. Rather, it came from studying phenomena of different kinds and requiring different methodologies. Some research involved highly mathematical studies in the laws of motion of idealized fluids. Other research, closest to our central point here, consisted of building machines for "artificial flight." And, of course, we must add to the list the actual observation of bird flight. All these research activities synergistically gave rise to aeronautical science through what we understand of the "natural flight" of birds and the "artificial flight" of airplanes. And it is in much the same spirit that I imagine diverse investigations in mathematics and in machine intelligence to act synergistically with psychology in giving rise to a discipline of cognitive science whose principles would apply to natural and to artificial intelligence.

It is instructive to transpose to the context of flying the common objections raised against AI. This leads us to imagine skeptics who would say, "You mathematicians deal with idealized fluids—the real atmosphere is vastly more complicated," or "You have no reason to suppose that airplanes and birds work the same way—birds have no propellors, airplanes have no feathers." But the premises of these criticisms are true only in the most superficial sense: the same *principle* (e.g., Bernoulli's law) applies to real as well as ideal fluids, and they apply whether the fluid flows over a feather or an aluminum wing.

Workers in the "cognitive studies" branch of AI do not share any one way of thinking about thinking, any more than traditional psychologists do. For some, the computer model is used to reduce all thinking to the formal operations of powerful deductive systems.

Aristotle succeeded in formulating the deductive rules for a small corner of human thinking in such simple syllogisms as "If all men are mortal and Socrates is a man, then Socrates is mortal." In the nineteenth century, mathematicians were able to extend this kind of reasoning to a somewhat larger but still restricted area. But only in the context of computational methods has there been a serious attempt to extend deductive logic to cover all forms of reasoning, including common-sense reasoning and reasoning by analogy. Working with this kind of deductive model was very popular in the early days of AI. In recent years, however, many workers in the field have adopted an almost diametrically opposed strategy. Instead of seeking powerful deductive methods that would enable surprising conclusions to be drawn from general principles, the new approach assumes that people are able to think only because they can draw on larger pools of specific, particular knowledge. More often than we realize, we solve problems by "almost knowing the answer" already. Some researchers try to make programs be intelligent by giving them such quantities of knowledge that the greater part of solving a problem becomes its retrieval from somewhere in the memory.

Given my background as a mathematician and Piagetian psychologist, I naturally became most interested in the kinds of computational models that might lead me to better thinking about powerful developmental processes: the acquisition of spatial thinking and the ability to deal with size and quantity. The rival approaches—deductive and knowledge based—tended to address performance of a given intellectual system whose structure, if not whose content, remained static. The kind of developmental questions I was interested in needed a dynamic model for how intellectual structures themselves could come into being and change. I believe that these are the kind of models that are most relevant to education.

The best way I know to characterize this approach is to give a sample of a theory heavily influenced by ideas from computation that can help us understand a specific psychological phenomenon: Piagetian conservation. We recall that children up until the age of six or seven believe that a quantity of liquid can increase or decrease when it is poured from one container to another. Specifically, when the second container is taller and narrower than the first, the children unanimously assert that the quantity of liquid has increased. And then, as if by magic, at about the same age, all children change their mind: They now just as unequivocally insist that the amount of liquid remains the same.

Many theories have been advanced for how this could come to pass. One of them, which may sound most familiar because it draws on traditional psychological categories, attributes the preconservationist position to the child's being dominated by "appearances." The child's "reason" cannot override how things "seem to be." Perception rules.

Let us now turn to another theory, this time one inspired by computational methods. Again we ask the question: Why does height in a narrow vessel seem like more to the child, and how does this change?

Let us posit the existence of three agents in the child's mind, each of which judges quantities in a different "simple-minded" way.* The first, A_{height} judges the quantity of liquids and of anything else by its vertical extent. A_{height} is a practical agent in the life of the child. It is accustomed to comparing children by standing them back to back and of equalizing the quantities of Coca-Cola and chocolate milk in children's glasses. We emphasize that A_{height} does not do anything as complicated as "perceive" the quantity of liquid. Rather, it is fanatically dedicated to an abstract principle: Anything that goes higher is more.

There is a second agent, called A_{width}, that judges by the horizontal extent. It is not as "practiced" as A_{height}. It gets its chance to judge that there is a lot of water in the sea, but in the mind of the child this principle is less "influential" than A_{height}.

Finally, there is an agent called $A_{history}$ that says that the quantities are the same because once they were the same. $A_{history}$ seems to speak like a conservationist child, but this is an illusion. $A_{history}$ has

* The computational perspective on conservation that follows is a highly schematized and simplified overview of how this phenomenon would be explained by a theory, "The Society of Mind," being developed by Marvin Minsky and the author and to be discussed in our forthcoming book.

no understanding and would say the quantity is the same even if some had indeed been added.

In the experiment with the preconservationist child, each of the three agents makes its own "decision" and clamors for it to be adopted. As we know, A_{height}'s voice speaks the loudest. But this changes as the child moves on to the next stage.

There are three ways, given our assumption of the presence of agents, for this change to take place. A_{height} and A_{width} could become more "sophisticated," so that, for example, A_{height} would disqualify itself except when all other things are equal. This would mean that A_{height} would only step forward to judge by height those things that have equal cross sections. Second, there could be a change in "seniority," in prerogative: $A_{history}$ could become the dominant voice. Neither of these two modes of change is impossible. But there is a third mode that produces the same effect in a simpler way. Its key idea is that A_{height} and A_{width} neutralize one another by giving contradictory opinions. The idea is attractive (and close to Piaget's own concept of grouplike compositions of operations) but raises some problems. Why do all three agents not neutralize one another so that the child would have no opinion at all? The question is answered by a further postulate (which has much in common with Piaget's idea that intellectual operators be organized into *groupements*). The principle of neutralization becomes workable if enough structure is imposed on the agents for A_{height} and A_{width} to be in a special relationship with one another but not with $A_{history}$. We have seen that the technique of creating a new entity works powerfully in programming systems. And this is the process we postulate here. A new entity, a new agent comes into being. This is A_{geom}, which acts as the supervisor for A_{height} and A_{width}. In cases where A_{height} and A_{width} agree, A_{geom} passes on their message with great "authority." But if they disagree, A_{geom} is undermined and the voices of the underlings are neutralized. It must be emphasized that A_{geom} is not meant to "understand" the reasons for decision making by A_{height} and A_{width}. A_{geom} knows nothing except whether they agree and, if so, in which direction.

This model is absurdly oversimplified in suggesting that even so simple a piece of a child's thinking (such as this conservation) can

be understood in terms of interactions of four agents. Dozens or hundreds are needed to account for the complexity of the real process. But, despite its simplicity, the model accurately conveys some of the principles of the theory: in particular, that the components of the system are more like people than they are like propositions and their interactions are more like social interactions than like the operations of mathematical logic. This shift in perspective allows us to solve many technical problems in developmental psychology. In particular, we can understand logical learning as continuous with social and bodily learning.

I have said that this theory is inspired by a computational metaphor. One might ask how. The "theory" might appear to be nothing but anthropomorphic talk. But we have already seen that anthropomorphic descriptions are often a step toward computational theories. And the thrust of the society-of-mind theory is that agents can be translated into precise computational models. As long as we only think about these agents as "people," the theory is circular. It explains the behavior of people in terms of the behavior of people. But, if we can think of the agents as well-defined computational entities similar to the subprocedures VEE, LINE, and HEAD in the procedure MAN, everything becomes clearer. We saw even in small programs how very simple modules can be put together to produce complex results.

This computational argument saves the society-of-mind theory from the charge of relying on a vicious circle. But it does not save it from being circular: On the contrary, like recursive programs in the style of the procedure SPI of chapter 3, the theory derives much of its power from a constructive use of "circular logic." A traditional logician looking at how SPI was defined by reference to SPI might have objected, but the computer programmers and genetic epistemologists share a vision in which this kind of self-reference is not only legitimate but necessary. And both see it as having an element of paradox that is only very partially captured by noting how children use their "inferior" logic to construct the "superior" logic of their next phase of development. To an increasing extent throughout his long career Piaget has emphasized the importance for intellectual growth of children's ability to reflect on their own thinking.

The "mathetic paradox" lies in the fact that this reflection must be from within the child's current intellectual system.

Despite its oversimplified, almost metaphorical status, the four-agent account of conservation captures an element of the paradox. A mathematical logician might like to impose on A_{height} and A_{width} a superior agent capable of calculating, or at least estimating, volume from height and cross-section. Many educators might like to impose such a formula on the child. But this would be introducing an element alien to the pre-conservationist child's intellectual system. Our A_{geom} belongs firmly in the child's system. It might even be derived from the model of a father not quite succeeding in imposing order on the family. It is possible to speculate, though I have no evidence, that the emergence of conservation is related to the child's oedipal crisis through the salience it gives to this model. I feel on firmer ground in guessing that something like A_{geom} can become important because it so strongly has the two-sided relationship that was used to conceive the Turtle: It is related both to structures that are firmly in place, such as the child's representation of authority figures, and to germs of important mathematical ideas, such as the idea of "cancellation."

Readers who are familiar with Piaget's technical writings will recognize this concept germ as one of the principles in his "groupments." They may therefore see our model as not very different from Piaget's. In a fundamental sense they would be right. But a new element is introduced in giving a special role to computational structures: The theme of this book has been the idea of exploiting this special role by giving children access to computational cultures. If, and only if, these have the right structure they may greatly enhance children's ability to represent the structures-in-place in ways that will mobilize their conceptual potential.

To recapitulate our reinterpretation of Piaget's theory makes three points. First, it provides a specific psychological theory, highly competitive in its parsimony and explanatory power with others in the field. Second, it shows us the power of a specific computational principle, in this case the theory of pure procedures, that is, procedures that can be closed off and used in a modular way. Third, it concretizes my argument about how different languages

can influence the cultures that can grow up around them. Not all programming languages embody this theory of pure procedures. When they do not, their role as metaphors for psychological issues is severely biased. The analogy between artificial intelligence and artificial flying made the point that the same principles could underlie the artificial and natural phenomena, however different these phenomena might appear. The dynamics of lift are fundamental to flight as such, whether the flyers are of flesh and blood or of metal. We have just seen a principle that may be fundamental both to human and artificial intelligence: the principle of epistemological modularity. There have been many arguments about whether the ideal machine for the achievement of intelligence would be analog or digital, and about whether the brain is analog or digital. From the point of view of the theory I am advancing here, these arguments are beside the point. The important question is not whether the brain or the computer is discrete but whether knowledge is modularizable.

For me, our ability to use computational metaphors in this way, as carriers for new psychological theories, has implications concerning where theories of knowledge are going and where we are going as producers and carriers of knowledge. These areas are not independent. In earlier chapters it was suggested that how we think about knowledge affects how we think about ourselves. In particular, our image of knowledge as divided up into different kinds leads us to a view of people as divided up according to what their aptitudes are. This in turn leads to a balkanization of our culture.

Perhaps the fact that I have spoken so negatively about the balkanization of our culture and so positively about the modularization of knowledge requires some clarification. When knowledge can be broken up into "mind-size bites," it is more communicable, more assimilable, more simply constructable. The fact that we divide knowledge up into scientific and humanistic worlds defines some knowledge as being a priori uncommunicable to certain kinds of people. Our commitment to communication is not only expressed through our commitment to modularization, which facilitates it, but through our attempt to find a language for such domains as physics and mathematics, which have as their essence communication between constructed entities. By restating Newton's laws as assertions about how particles (or "Newtonian Turtles") communicate with one another, we give it a handle that can be more easily grabbed by a child or by a poet.

Consider another example of how our images of knowledge can subvert our sense of ourselves as intellectual agents. Educators sometimes hold up an ideal of knowledge as having the kind of coherence defined by formal logic. But these ideals bear little resemblance to the way in which most people experience themselves. The subjective experience of knowledge is more similar to the chaos and controversy of competing agents than to the certitude and orderliness of *p's* implying *q's*. The discrepancy between our experience of ourselves and our idealizations of knowledge has an effect: It intimidates us, it lessens the sense of our own competence, and it leads us into counterproductive strategies for learning and thinking.

Many older students have been intimidated to the point of dropping out, and what is true for adults is doubly true for children. We have already seen that despite their experience of themselves as theory builders, children are not respected as such. And these contradictions are compounded by holding out an ideal of knowledge to which no one's thinking conforms. Many children and college students who decide "I can never be a mathematician or a scientist" are reflecting a discrepancy between the way they are led to believe the mathematician must think and the way they know they do. In fact the truth is otherwise: Their own thinking is much more like the mathematician's than either is like the logical ideal.

I have spoken of the importance of powerful ideas in grasping the world. But we could hardly ever learn a new idea if every time we did we had to totally reorganize our cognitive structures in order to use it or if we even had to insure that no inconsistencies had been introduced. Although powerful ideas have the capacity to help us organize our way of thinking about a particular class of problems (such as physics problems), we *don't* have to reorganize ourselves in order to use them. We put our skills and heuristic strategies into a kind of tool box—and while their interaction can, in the course of time, give rise to global changes, the act of learning is itself a local event.

The local nature of learning is seen in my description of the acquisition of conservation. The necessary agents entered the system locally; their top goals were in contradiction with each other; the agent that finally reconciles them leaves them in place. There is no reason why this "patchwork theory" of theory building should be considered appropriate only for describing the learning of children. Research in artificial intelligence is gradually giving us a surer sense of the range of problems that can be meaningfully solved on the pattern we have sketched for the conservation problem: with modular agents, each of them simple-minded in its own way, many of them in conflict with one other. The conflicts are regulated and kept in check rather than "resolved" through the intervention of special agents no less simple-minded than the original ones. Their way of reconciling differences does not involve forcing the system into a logically consistent mold.

The process reminds one of tinkering; learning consists of building up a set of materials and tools that one can handle and manipulate. Perhaps most central of all, it is a process of working with what you've got. We're all familiar with this process on the conscious level, for example, when we attack a problem empirically, trying out all the things that we have ever known to have worked on similar problems before. But here I suggest that working with what you've got is a shorthand for deeper, even unconscious learning processes. Anthropologist Claude Lévi-Strauss[2] has spoken in similar terms of the kind of theory building that is characteristic of primitive science. This is a science of the concrete, where the relationships between natural objects in all their combinations and recombinations provide a conceptual vocabulary for building scientific theories. Here I am suggesting that in the most fundamental sense, we, as learners, are all *bricoleurs*.[3] This leads us into the second kind of implication of our computational theory of agents. If the first implications had to do with impacts on our ideas about knowledge and learning, the second have to do with possible impacts on our images of ourselves as learners. If *bricolage* is a model for how scientifically legitimate theories are built, then we can begin to develop a greater respect for ourselves as *bricoleurs*. And of course this joins our central theme of the importance and power of Piagetian learning. In order to create the conditions for bringing what is now non-Piagetian learning to the Piagetian side, we have to be able to act in good faith. We have to feel that we are not denaturing knowledge in the process.

I end this chapter on cognitive theory and people with a conjecture. Earlier I said that I would not present Piaget as a theorist of stages. But thinking about Piagetian stages does provide a context in which to make an important point about a possible impact of a computational culture on people. Piaget sees his stages of cognitive development as invariable, and numerous cross-cultural investigations have seemed to confirm the validity of his belief. In society after society, children seem to develop cognitive capacities in the same order. In particular, his stage of concrete operations, to which the conservations typically belong, begins four or more years earlier than the next and final stage, the stage of formal operations. The construct of a stage of concrete operations is supported by the observation that, typically, children in our society at six or seven make a breakthrough in many realms, and seemingly all at once. They are able to use units of numbers, space, and time; to reason by transitivity; to build up classificatory systems. But there are things they cannot do. In particular, they flounder in situations that call for thinking not about how things are but about all the ways they could be. Let us consider the following example, which I anticipated in the introduction.

A child is given a collection of beads of different colors, say green, red, blue, and black, and is asked to construct all the possible pairs of colors: green-blue, green-red, green-black, and then the triplets and so on. Just as children do not acquire conservation until their seventh year, children around the world are unable to carry out such combinatorial tasks before their eleventh or twelfth year. Indeed, many adults who are "intelligent" enough to live normal lives never acquire this ability.

What is the nature of the difference between the so-called "concrete" operations involved in conservation and the so-called "formal" operations involved in the combinatorial task? The names given them by Piaget and the empirical data suggest a deep and essential difference. But looking at the problem through the prism of the ideas developed here gives a much different impression.

From a computational point of view, the most salient ingredients of the combinatorial task are related to the idea of procedure—systematicity and debugging. A successful solution consists of following some such procedure as:

1. Separate the beads into colors.
2. Choose a color A as color 1.
3. Form all the pairs that can be formed with color 1.
4. Choose color 2.
5. Form all the pairs that can be formed with color 2.
6. Do this for each color.
7. Go back and remove the duplicates.[3]

So what is really involved is writing and executing a program including the all-important debugging step. This observation suggests a reason for the fact that children acquire this ability so late: Contemporary culture provides relatively little opportunity for *bricolage* with the elements of systematic procedures of this type. I do not mean to say that there are no such opportunities. Some are encountered; for example, in games where a child can create his own "combinatorial microworlds." But the opportunities, the incentives, and the help offered the child in this area are very significantly less than in such areas as number. In our culture number is richly represented, systematic procedure is poorly represented.[4]

I see no reason to doubt that this difference could account for a gap of five years or more between the ages at which conservation of number and combinatorial abilities are acquired.

The standard methodology for investigating such hypotheses as this is to compare children in different cultures. This has, of course, been done for the Piagetian stages. Children at all the levels of development anthropologists have been able to distinguish, and in over a hundred different societies from all the continents, have been asked to pour liquids and sort beads. In all cases, if conservation and combinatorial skills came at all, conservation of numbers was evidenced by children five or more years younger than those evidencing combinatorial skills. Yet this observation casts no doubt on my hypothesis. It may well be universally true of precomputer societies that *numerical* knowledge would be more richly represented than *programming* knowledge. It is not hard to invent plausible explanations of such a cognitive-social universal. But things may be different in the computer-rich cultures of the future. If computers and programming become a part of the daily life of children, the conservation-combinatorial gap will surely close and could conceivably be reversed: Children may learn to be systematic before they learn to be quantitative!

Chapter 8

Images of the Learning Society

THE VISION I HAVE PRESENTED is of a particular computer culture, a mathetic one, that is, one that helps us not only to learn but to learn about learning. I have shown how this culture can humanize learning by permitting more personal, less alienating relationships with knowledge and have given some examples of how it can improve relationships with other people encountered in the learning process: fellow students and teachers. But I have made only passing remarks about the social context in which this learning might take place. It is time to face (though I cannot answer) a question that must be in many readers' minds: Will this context be school?

The suggestion that there might come a day when schools no longer exist elicits strong response from many people. There are many obstacles to thinking clearly about a world without schools. Some are highly personal. Most of us spent a larger fraction of our lives going to school than we care to think about. For example, I am over fifty and yet the number of my postschool years has barely caught up with my preschool and school years. The concept of a world without school is highly dissonant with out experiences of our own lives. Other obstacles are more conceptual. One cannot define such a world negatively, that is by simply removing school and putting nothing in its place. Doing so leaves a thought vacuum that the mind has to fill one way or another, often with vague but