

New Theories for New Learnings

by

Seymour Papert *

Copyright 1984 Seymour Papert

* This paper is the edited transcript of a speech given 18 April 1984 at the National Association of School Psychologists' Conference. It will also be published in the October 1984 issue of the *School Psychology Review*.

The computer is going to be a catalyst of very deep and radical change in the educational system. Guiding that change will require far more than new facts, new statistics, and new policies. We need new conceptual frameworks for thinking about the learning process itself. These are big theses and deserve a more integrated and scholarly treatment than I could dream of attempting today. My goal is to expose you informally to some images that may stimulate your own thinking.

I will give examples of some such new concepts, but will first talk about the need to cast off much that is ingrained in current approaches to educational psychology. A simple example is a research methodology that looks for 'the effect' of such-and-such intervention, e.g.: "What is the effect of learning set theory on number skills?" A literature is beginning on 'the effect' of 'the computer' on 'cognitive development,' or on 'learning mathematics,' or on 'schools,' or on 'the classroom.' Such questions are totally meaningless.

A typical example comes from a much quoted paper from the Bank Street College of Education's Children and Technology Center. This work is certainly more elaborate than much other research across the country, but it suffers from the same fault. In the central paper by Pea and Kurland¹ on the cognitive effects of programming, it is assumed that programming as such could have a definite cognitive effect, and that this effect—if it exists—would show itself in improved abilities to plan. I am quoted as having advanced such a view, though the position I have really advanced is quite different.

If programming does influence thinking, this influence must not be conceptualized as a direct consequence of learning to program. The kind of effect I have in mind is something more indirect. Fluency in programming provides an opportunity for teachers to teach in new ways and for students to learn in new ways. But seizing this opportunity requires *action* by teachers and by learners.

To warm up to the point, consider an imaginary argument about the effect of "happiness" on

1. Roy D. Pea and D. Midian Kurland: "On the Cognitive Effects of Learning Computer Programming," *New Ideas in Psychology*, Vol. 2, No. 2, 1984.

"learning." Other things being equal, it's obvious that everybody learns better when relaxed and happy. A happy state is undoubtedly conducive to learning when learning is taking place. But if you blindly looked for correlations between degree of happiness and amount of learning, you would probably find a negative result. Many people are happiest when they are not learning. Similarly, I suppose that one can be deeply immersed in programming without much learning going on. But this proves nothing about whether programming can be used to enhance learning.

Much current research is marred by another flaw as well. Imagine (if you can) that we lived in a world without writing—and, of course, without pencils, pens and books. Then one day, somebody invents writing and the pencil, and people say, "Wow, this would be great for education. Let's give these things to all the children and teach them to write." So then somebody else says, "Hey, wait a minute. You can't just do that. You can't just give every child a pencil. You'd better start by doing some rigorous experiments on a small scale. So we'll put one pencil in a classroom and we'll see what happens. If great things happen, we'll put two pencils in a classroom, and if greater things happen, then we'll put in more."

But this would be missing the whole point of the pencil—indeed, many points. Pencils are an extension of you and whatever you do. They are everywhere, and are deeply embedded in our culture. Even a baby picks up a pencil and scribbles on the wall, trying to use this adult instrument. Babies don't know what a pencil is good for in the adult world, but find something that looks and feels like what it seems to be used for by adults. The mother is not too happy when her baby scribbles on every wall in sight, but this messing is important—it's part of the glue that joins the child to the culture and the culture to the child. Later on, the pencil is used for increasingly diverse purposes: writing and drawing and arithmetic and essays and illicit notes in the classroom. It's there all the time. It doesn't have 'a purpose;' it doesn't have 'an effect.'

As I see the computer, it will come to be used in that way. For example, the computer is already the only writing instrument seriously considered by people (like many reporters and authors) for whom writing is an important activity. Sooner or later, we are going to see the same for children—and only then will the computer begin to achieve its true level of

significance. But one computer in the classroom is not even a vague first approximation to this.

I am not arguing against research. How else can we pursue the idea of computer as pencil? I am myself a researcher. What I sometimes ridicule is the use of experimental methods that have the trappings of science, but no relation to the important issues. For example, experiments on one pencil per class would have *no* bearing on the effect of pencils in our culture. Nor do experiments on the effect of a small degree of artificial access to computers have any bearing on the effect computers could have when they are culturally integrated.

To illustrate what I mean by culturally integrated, I turn to another flaw in much contemporary research. Researchers want to know if a word processor is good for creative writing. So in a typical experiment, for 6 weeks, 3 times a week, 40 minutes a day, the children come in, sit down in front of a word processor, and somebody says "Create," or "Write," while a person hovers about with a note pad writing down what the children do.

It's a wonder those experiments don't show total paralysis. They certainly don't show nearly as much effect on creativity as we see through informal, ethnographic observation of what children and adults do in computer-rich environments. Those children *are* creative when they're alone in a corner with nobody looking, when they feel comfortable with the environment and get an idea and start getting absorbed in what they're doing. Again, this example shows that the psychologist trying to observe has got to get away from the idea 'the computer produces an effect.' Hypothesis: 'Having a word processor improves creativity'—so get the word processor and measure creativity and that's all you need to think about. But for the computer to serve that purpose, it's got to be suitably integrated. At the very least, it must be freely available and its user must be relaxed. The experiment can't be in a fixed time, delimited, with people watching you and telling you, "Now go create."

So much for criticism of old methodologies. What should be done? I think that the important point to realize is that the penetration of computers into learning is still very new and very partially accomplished. We are only now getting to the point where significant, large scale experiments become possible. One condition for that is an abundance of computers.

I think we are moving toward a time when a computer for every child, rather than one in every classroom, will be widely accepted. When a sufficient number of schools have freely available computers, we will start to see dramatic consequences—and then it will begin to spread. By 1990, one computer per child should be an extremely common state of affairs.

In the meantime, there's a lot of balihoo in the press about this computer revolution—that computers are everywhere in the schools. But in fact, there is scarcely one for every 100 children—which is like no computer at all if you average it out. A very small number of schools are thinking in terms of one for every 30 children because that means each child can get an hour a week at the computer—which is a little better. But think of one hour a week for the pencil, and it's obvious that this is still absurd. There's a tiny number of schools who are thinking that one per 5 or 6 children is a minimum, because that way each child can have access to it every day.

This scarcity of computers is a very serious state of affairs. On the whole, computers are being used in ways that couldn't possibly do any serious good to anybody—and so we are beginning to see a backlash. An article called "Computer Worship" in *Science* '84² expresses this backlash: Computers are everywhere and everybody is talking about computer literacy, but what good does it do? My answer is that as *educational computing is usually practiced* it does no good, and can even do harm because using computers in puny and trivial ways undermines their image, and reinforces false perceptions. Most children are simply not getting an opportunity to become familiar with the computer as a powerful entity they can control.

I don't think that backlash will wipe out computers altogether; very strong social forces are moving toward imposing them on the world. But certainly a lot of confusion is generated about the whole set of psychological questions concerning the relationship between computers and children. We need to keep very clear—both for our own thinking and for the

2. Menosky, "Computer Worship," *Science* '84, May 1984, Vol. 5, No. 4, p. 40. See also the Summer 1984 issue of *Teachers College Record* (Vol. 85, No. 4, published by Teachers College of Columbia University), which also contains a variety of articles of this nature.

public—that 'computers' doesn't mean 'a computer sitting there' under arbitrary conditions. Just like in the pencil analogy, it's crucial that we also consider children's access to them.

But throwing out entrenched, misleading ideas and providing children with open access to computers is not enough in itself. We also need a whole new set of concepts and paradigms for studying the learning process in this burgeoning computer age.

What I am saying about computers is no doubt true about other materials and other ways of teaching. But the computer's impact is far more powerful. It gives us the opportunity of making much more radical changes in the conditions of learning than any other means we have had in the past—so the need for radical thinking becomes more urgent and the problems become more fundamental. The computer's presence and capabilities demand a deeper reexamination of accepted theory—as well as experimental methodologies—in educational psychology. In the rest of this talk, I shall illustrate how theoretical concepts and methodologies have to be drawn from various human sciences, including ethnography, psychoanalysis, political science and, of course, epistemology. And not just borrowed from these disciplines, either. We need to find ways of integrating them through a new conceptual framework for education.

Possibly the most important way that computer technology can change learning is by suggesting new and different kinds of relationships to knowledge and to learning. An example of this is seen in research by sociologist Sherry Turkle.³ It concerns a 7 year old, extremely hyperactive child who wouldn't sit still, had been doing nothing at school, wouldn't sit and write and probably couldn't. He was playing with a version of Logo that has objects that move on the screen—in this case, little bright, colored balls. They would fly out from the center and then return, over and over. This child found a very exciting pulsating pattern, and he sat there dancing to it.

3. Her book, *The Second Self: Computers and the Human Spirit* (New York: Simon & Schuster, 1984), describes her studies of people's emotional relations to computers, and discusses the ways in which computers are changing not only our society and culture, but even our ideas about who we are, what our world is like, and how we think and learn.

However, he didn't quite like that pattern. The balls traveled out at a certain speed for a certain amount of time, and he knew that you could control the speed and the time by giving instructions to the computer. This child said he wanted it to be 'perfect'—he meant he wanted the balls to fly out exactly to the edge of the screen and then come back in at a rhythm he could dance to. Now imagine a scene of several hours—quite long for this child—about three or four sessions of 45 minutes each. For ninety percent or more of this time, the child was dancing, watching the balls move on the screen, and calling friends over to show them what the balls were doing. Ten percent of the time, he was talking about what was happening on the screen, and about one percent of the time he was actually intervening—sticking in numbers that would make the pattern work better.

Now what's really amazing is that he was working with numbers. It is not amazing that his concept of what to do with the numbers was *extremely* rudimentary. That is, he knew that at a certain speed like 40, they would go out to a certain place—and he wanted them to go further, to the edge of the screen. He knew he could achieve this by changing the speed, but he wasn't too sure whether he should increase it or decrease it. Once he decided to increase it, he didn't know if he should change 40 to 41—or 40 to a million. So he tried all sorts of numbers, an effort that was worthwhile for him because the changes often had dramatic effects, even if not the one he was looking for. Sometimes the balls didn't do anything at all, and he found that funny—or they flew out so fast that he could hardly see them, and he found that funny too. In any case, he was involved with using numbers to move these things around—manipulating numbers like another child manipulates mud pies.

This child became absorbed in a typical piece of mathematics usually done (if at all) by much older students. He was dealing with two variables, or perhaps three—speed, time and distance—trying to feel out and explore the relationships between them. And it's highly unlikely that this child would ever have involved himself in that particular intellectual situation in the normal school setting.

There's a traditional term for thinking about such situations: motivation. The term 'motivation' leads you to think that he wouldn't normally work with numbers because he wasn't motivated, because it didn't have a payoff for him. This concept suggests giving him a payoff: chocolate

bars or gold stars or punishments or praise. But something much more fundamental is involved. In *Mindstorms*,⁴ I borrowed a term from psychoanalysis which I tried to propagate—not too successfully, but maybe it will still catch on. We need a term. The one I used is *syntonic*, ego-syntonic. What does this mean?

Some things feel like a part of you. I borrow from Robert Lawler the analogy of interior psychic space being like a big country. Certain parts of it are the center, the capital. Those are the important parts, the ones that feel most like yourself. Then there are outlying regions that happen to be there but they are more like an embarrassment—you hate to go there, you feel uncomfortable when you do, and it is all messy and slummy and you feel just awful. 'Alienated' refers to what is "out there" and 'syntonic' refers to what is "at the center." As a rough approximation, 'syntonic' means the opposite of alienated. And for this child, mathematics as we do it in the classroom is thoroughly alienated.

Now *why* is it alienated? I think there are two reasons. Both are relevant to this particular case and they interlock. One is a personality reason. The other is best described in cultural terms.

School as we set it up is tailor-made for certain personality types. I like to think of certain scales of oppositions: like obsessional-compulsives and hysterics.⁵ This is a little over-simplified, but will help us talk quickly. Our hyperactive child is toward the extreme end of the hysteric scale. The hysteric likes generalities, likes dramatic effects, doesn't like precise detail. The obsessional-compulsive likes little detail, likes the static, likes things you can examine closely.

When we look at school math, it's tailor-made for the obsessional-compulsive. You sit there and fill numbers into little squares on paper. Some people *do* like this kind of activity, but others don't—and some not only don't like it, they can't stand it. Our only technology for

4. Seymour Papert, *Mindstorms: Children, Computers and Powerful Ideas*, New York: Basic Books, 1980.

5. See, for example, David Shapiro, *Neurotic Styles*, New York: Basic Books, 1965.

teaching math was, until recently, this pencil-and-paper activity—which made the entryways for learning mathematics (and many other formal school subjects) extremely difficult for anyone who *didn't* like this obsessional style. Even if they do get to enjoy math, the learning style still doesn't match their personalities—and if it doesn't match their personalities, they don't do as well. So traditional school math isn't simply syntonic or alienated. It is more syntonic for obsessional-compulsive types, and more alienated for hysteric ones. Using dynamic, moving objects to present mathematics is ideal for people with hysteric styles—like the hyperactive child—but until we had computers, we didn't have any way of making a mathematics of moving things. Throwing tennis balls doesn't qualify in itself as mathematics because it isn't formalized—even informally; for example, it isn't tied to numbers and ways of manipulating them.

The computer allows us for the first time to match the subject matter and learning style to the personality type. We see the same syntonic and alienated responses with cultural types as well. People from different sub-cultures and different cultural backgrounds also have, in the same kind of spirit, different relationships to knowledge.

I'd like to quote a typical example of a cultural difference that's very marked in African cultures—which are essentially oral and very transactional. This example concerns adults, but we see exactly the same kinds of differing styles in children. Discussions between people from African cultures and Europeans or Americans often break down in a predictable way. One African described it like this to me once: "I can't talk to Americans because they always say what they mean."

Now you would think that's absurd because we were all trained to think the opposite, that communication breaks down because people *don't* say what they mean. But what he meant was very simple, very important. For him, the model of Americans having a discussion is like this: A says what he means, B says what she means, and then there's a confrontation. A and B have a long argument and in the end, one is right and the other is wrong. Whereas in his model, you don't say anything so confrontational at all. Instead, C says something—and he made a little spiral with his fingers—and then D says something, and they talk for a long time. His model is one of convergence—that there *is* a long time to talk and eventually you all agree.

You arrive at a point of view that isn't yours and isn't mine—it's consensual. This kind of transactional, consensual approach to knowledge is very typical of a lot of cultures. Even inside America, some sub-cultures are more consensual than others.

However, our presentation of knowledge in school is very unconsensual. There's the right and there's the wrong. The teacher says something and the pupil has to pick it up. You write things down on a piece of paper—and once they are written, they are frozen.

The computer is able to make that relationship to knowledge much more transactional because you are interacting with this machine: you are doing things and changing them and seeing stuff happen. The little hyperactive child is a clear example of just that. But the computer can be used in either style—and most of the time, it is used in a very confrontational way: "Here's the question. Here's the answer. You are wrong, Johnny. Try again." Traditional classroom procedures are simply transferred to the new technology. Very few people are looking for ways to use the computer that allow knowledge to be acquired and presented in a form that matches different personality and cultural types.

The story of the hyperactive child also raises another very important point about the nature of our relationship to knowledge and to learning. People have made an artificial split between the cognitive and other aspects of psychology, especially the affective. The word 'cognitive' is increasingly used to refer to knowledge—facts and skills. It is associated with very neutral terms, very general and cold-blooded. Its categories are true/false, relevant/irrelevant, similar/dissimilar. But in the examples we've talked about, the child is involved in a very hot-blooded way with a particular, personal form of mathematical thinking.

In order to understand such situations, we really have to see the cultural, the affective, and the cognitive welded together in the most intimate way. There's no sense in which you can describe the important part of what happened to that little hyperactive child as simply the acquisition of knowledge or as "the processing of information." It's just not of that nature. What he acquired is much more than facts or concepts or understandings. It's a *relationship*—with all the emotional qualities that entails. It's feeling good about numbers, for example. Of course, cognitive terms are relevant—in part, it's his knowledge of numbers that he feels

good about—but you've also got to think of the relationship in affective terms.

Another issue of this sort is raised very acutely by the presence of computers. In *Mindstorms*, I presented a concept that I called 'Piagetian learning' or 'learning without instruction' or 'learning without teaching.' I'd like to clarify what I meant by that because I've been misinterpreted as saying, "Give the child a computer and the child will learn all alone."

The analogy that I used doesn't really lead to that. I said to think of how the baby learns to speak. There's no curriculum, no professional teacher, and it's very effective. Almost all babies learn to speak. Far more of them learn to speak than children in school learn mathematics, or writing, or many other things that we teach by so-called professional, curriculum-driven, and expensive processes in the school. What makes this inexpensive, unprofessional process of the baby learning to speak so effective? I think the answer is very clear: what makes it effective is that it's melded so intimately into the social and emotional life of the child. The child is surrounded by other people who speak, and learns to speak because of this, through being immersed in a language culture. The analogy for computers is that the child learns computing by being immersed in a computer culture, mathematics by being immersed in a math culture.

We now have a few observations from school settings with a high density of computers. When children have fairly free access to them over a long enough period, a computer culture develops where many children in the school know about computing and a lot of knowledge is resident in the community of children. Much crucial learning happens without teaching. But as this cultural knowledge grows and develops, the amount of teaching done by the adult teacher doesn't diminish but changes. The teacher is not displaced but enhanced.

The adult teacher no longer has to focus on elementary, simple things and can spend whatever time is available on more subtle things: on spotting that this child is in trouble, that this child is blocking, that this child could be doing something more exciting, and so on. There are many situations where the children are not effective teachers. But in the ideal computer culture, maximum use is made of the children's abilities to teach their peers. In the experimental school at the Learning Research Development Center at Pittsburgh, Lesley

Thyberg introduced a rule about the computer: "Ask 3 before you ask me." That is, ask three other children before you ask the teacher. I have seldom seen second graders learn to work with the computer as effectively as they did there!

The computer has very special qualities because it's so interactive and experimental. A child can be stuck and can ask another child what to do, and the other child doesn't have to be very articulate to say, "Do this." Now "do this" doesn't convey everything that a highly articulate teacher might, but it does transmit a little bit of knowledge which is enough for that first child to get moving. And when you get moving and start using it, your knowledge grows and these bits accumulate. This is how culturally diffused knowledge spreads.

This was one of our central goals in designing Logo: to maximize the extent to which learning can be taken up in pieces, can be spread in this way. One concept which I'm beginning to think is *the* key concept for thinking about the epistemology of the learning process is what I call 'fragmented' or 'fractured' knowledge. It is also a central epistemological issue for education.

By contrast, in the standard teaching situation, the teacher is going to show you a whole package of knowledge. The teacher has this algorithm—the square root, say—and stands there at a blackboard to explain and transmit it to you, the student, and you'll pick it up and put it in your head. Well, that might or might not work sometimes, but it goes wrong in ways that you could very easily and obviously describe as: "It gets broken in transmission." The teacher might be trying and the child might be trying too, but maybe the child didn't hear a little, or maybe couldn't understand it, or didn't want to, or didn't have some prerequisite. What comes across is fractured knowledge, not the whole piece but a broken piece.

So how do these pieces get put together? We are beginning to understand that some knowledge is more easily put together than other knowledge—also that we can design learning situations that *make it easier* to put these broken pieces together. I'd like to tell you a final story illustrating this.⁶

In one school, children in the first through fourth grades were using a Logo system like the one that I described before. You could write programs that made these objects—called

sprites—move around, you could give them different shapes, and you could also make them draw aesthetically interesting designs on the screen. Some children—the more obsessional ones, to take that theme again—prefer making the static patterns. So there were several different kinds of things the children could do.

However, the teachers of the first and second graders decided that they wouldn't show these children how to make the sprites move. Their reason was quite straightforward. To make the sprites move, you had to give them both a speed and a direction. The way you gave them direction was to say `SETHEADING 270` to go west, and `SETHEADING 90` to go east (or left and right on the screen). The teachers felt that numbers like 270 and the idea of angles were simply too complex for the younger children to understand.

The third and fourth graders were shown how to work with speed and direction, however. Pretty soon, the younger children saw the sprites moving on the older children's screens and began to ask them: "How do you do that?" The first few times they asked, they didn't get enough of an answer to be able to do anything with it. But they kept at it. Gradually, the older children became more confident and articulate, and the younger ones also become more understanding about the whole system. So one day, a threshold was crossed where it was now possible for one second grader to bring back just enough knowledge about dynamics so that his classmates could work with it. The second grader didn't really understand what the older child said—and it's very important to realize this. He came over to the investigator in the middle of the next day and said, very proudly too, "I got it. Numbers are secret codes for directions." He explained that they didn't yet know the code, "but we're working on it."

In fact, there were fourteen of them working on this code, and what was exciting about it was that yes, they didn't know what 270 meant, and yes, they didn't know what degrees meant, but they did know that they could bring this into a conceptual frame that they understood very well: codes. By making this connection between the known and the new, they could work on it. A few weeks later, they were doing pretty well—and a few months later, they had objects

6. This incident also comes from Sherry Turkle's research, and is reported in her book, *The Second Self*.

flying all over the screen under perfect control, they knew about 270 degrees, they knew about degrees and angles and these big numbers, and they knew all sorts of things because they had broken this code.

These students were all working with the same system and they also had open access to the computers, so they could share little bits of knowledge and share their discoveries. At the same time that they were building their own understanding of how to make the sprites move, they were also building a computer culture from which they all learned.⁷

In conclusion, I want to apply my own concepts to my own talk. I have presented here some fragments of ideas. I expect some of them to take root and grow in some of your minds—not so much because what I have said is correct or compelling in itself, rather because I think that ideas like these have cultural resonance. They resonate with the problems that are in the air today, with the issues many of you must be encountering in your professional lives. This vision of the way ideas spread is as central to my model of how to talk to peers as it is to the creation of learning environments for children. The vision of knowledge fragments is equally central. In a different setting, I might be apologetic about bombarding you with so many unintegrated pieces of ideas. But I am not. If you will pick up some of these pieces and put them together in your own ways, I shall have been much more effective than if I had come here with a monolithically coherent exposition. No one learns by taking over another's point of view. One learns only by building one's own.

7. I discuss several other epistemological issues this incident raises in "Microworlds: Transforming Education," to appear in *The Computer Culture* (in press): the compendium of papers presented at the ITT Key Issues Conference, March 1984, at The Annenberg School of Communications, University of Southern California.