

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
A.I. LABORATORY

October 1971

Artificial Intelligence
Memo No. 246

LOGO
Memo No. 1

A COMPUTER LABORATORY FOR ELEMENTARY SCHOOLS

Seymour Papert

This paper is the substance of a proposal to the N.S.F. for support of research on children's thinking and elementary education.

This work was supported by the National Science Foundation under grant number GJ-1049 and conducted at the Artificial Intelligence Laboratory, a Massachusetts Institute of Technology research program supported in part by the Advanced Research Projects Agency of the Department of Defense and monitored by the Office of Naval Research under Contract Number N00014-70-A-0362-0002.

1. Abstract

This is a research project on elementary education whose immediate objective is the development of new methods and materials for teaching in an environment of computers and computer-controlled devices. Longer term objectives are related to theories of cognitive processes and to conjectures about the possibility of producing much larger changes than are usually thought possible in the expected intellectual achievement of children. This proposal is formulated in terms of the self-sufficient immediate objectives.

2. General Principles

The methods already developed in our project add new dimensions to the possibilities of curriculum reform in mathematics, physics, biology, and other conventional school subjects. They allow us to remove artificial barriers between "subjects", and so to integrate mathematics with the other sciences, to integrate science with linguistics and other academic areas and even to establish significant links between "academic" work and freer activities such as music and gymnastics. Partly through removing these barriers, partly independently, we are able to achieve a more involved and personal participation of children in their work.

The key to achieving all this is recognizing three new roles for the computer in education. These are:

(a) Mathematical Technologies for Children

We find that the intention behind this is most effectively conveyed by a fantasy. One might dream of having children learn mathematics by giving them a ship to sail the ocean, a sextant to fix their position and a cargo to trade with distant peoples. A large part of our work is directed at trying to make this dream come true (at least in principle) by creating mathematical instruments more manageable than ships and sextants, but which still allow the child to develop and exercise mathematical arts in the course of meaningful, challenging and personally motivated projects.

In our context the computer is not merely a device for manipulating symbols. It actually controls real, physical processes: motors that turn, trucks that move, boxes that emit sounds. By programming it, the child is able to produce an endless variety of actions in a completely intelligible, controlled way. New mathematical concepts translate directly into new power for action. Self-generated projects induce an immediate and practical need to understand the mathematics of movements, the physics of moving bodies and the formal structure of sound patterns.¹


(b) A Model for Learning About Learning and Other Cognitive Skills

Because "children learn by doing," research in education has the task of inventing better things for children to do, and we

¹A less lyrical, more technical account of some of the actions we have already made accessible, can be found in the appended paper, "Twenty Things To Do With A Computer."

see our technologies as a particularly good example that is completely justified on this score alone: one can do infinitely more with our machines than with the rods and blocks and shiny but passive laboratory apparatus of the curriculum reform of the previous two decades. But we are also influenced by another dimension of educational philosophy indicated by an addition to the slogan: "Children learn by doing, and by thinking about what they do." So, another dimension of the task facing research in education is to give children better ways and means to think about what they do; and this includes looking for activities whose structure might allow the child a particularly clear view of his own intellectual activity and so help him achieve a more articulate understanding of "doing."

To see how computer-controlled devices excel in this function, consider a simple task of the kind we assign during the first week of an elementary school course using a device known as a turtle:



The turtle "understands" simple commands (typed at the computer terminal) such as

FORWARD 73	Which makes it advance 73 units
RIGHT 90	Which makes it rotate around its center through a 90° angle
PENDOWN	Which makes the turtle lower its pen and leave a trace of its path
PENUP	

To make the turtle do anything more complex, the child must write a procedure. For example, the child might want to "teach the turtle" the command "BOX", whose effect is intended to make the turtle move around a square. To define it, the child must "tell the computer how to BOX." So he types:

TO BOX	The word "TO" indicates that we are defining. So the turtle does not go forward and turn right when the next lines are typed. Instead, it "remembers" these directions until it is given the command BOX. Notice the use of recursion in line 3 to set up a self-perpetuating process.
1 FORWARD 100	
2 RIGHT 90	
3 BOX	
END	


The last line says the definition is finished. The child now has only to type:

BOX	This command causes the turtle to go round and round and round . . . a square whose side is 100 units long. For this example we do not need to know how to make the turtle stop -- though the children do.
-----	--

Now suppose, to develop our example, that the child wants to make the turtle draw a triangle. He writes:


```
TO TRI
1 FORWARD 100
2 RIGHT 60
3 TRI
END
```

But when the command TRI is given, instead of drawing a triangle, the turtle draws this hexagon:



When a child makes mistakes in arithmetic classes, he (and, alas, often his teachers) might conclude, "I'm dumb" or "I'm not mathematically minded" or "I never could understand that" or (at best) "to hell with it." The child rarely engages in constructive thinking about how and why the mistake happened and what can be done about it. But when the turtle draws a hexagon instead of a triangle, the reaction is much more constructive (or at least easily becomes so with a little encouragement). Children almost unanimously see the turtle (rather than themselves) as doing the "wrong thing." And, of course, we strongly encourage this, for the child is then much more ready to be objective about what happened. Moreover, we are able to urge them to understand exactly why the turtle did what it did . . . rather than

merely make it do the "right thing." The diagram below will help the reader understand what did happen -- for an elementary school child with a past history of unbroken failure in math, it needs a longer time, and often comes as a revelation that it is ever possible to understand anything "exactly."



We have observed many children gradually transfer the objectivity and skill acquired in debugging the turtle to "debugging" their own mistakes and thus acquire a more constructive approach to their own learning.² So, in a sense, they are learning something about the practical psychology of learning as well as about programming and geometry. They are also learning something about the nature of thinking by seeing an objective, externalized example of that curious process called "formal" or "rigorous" thinking which remains so deeply mysterious to most people who pass through school without ever understanding what mathematical thinking is, or why it should be.

Finally, under this heading, we mention for completeness an area of our work with children that lies beyond the scope of this proposal. This is research on how people learn -- and can become better at learning -- physical skills such as juggling, balancing feats and other "circus arts." Our conjecture, (for which we are gradually collecting confirmation) is that these skills can be

²Other examples can be found in the appended papers "Teaching Children Thinking" and "Teaching Children To Be Mathematicians."

learned with special ease by subjects who have acquired sophistication in the "debugging model of learning" and in a number of similar concepts related to computation.

(c) Theoretical Computer Science as a Source for Research on Elementary Education

There has been very little interaction between the elementary curriculum reform movement and the conceptual, theoretical wings of computer science. We believe that by opening this dialogue we may be unleashing an intellectual force of great power; for education might be the area of research and application needed for certain germinating ideas in the theory of computation to acquire purpose and maturity. Artificial Intelligence springs to mind immediately as an area of computer science that should interact with education. But, there are other, less obvious, examples with a greater potential for immediate impact. One of these, which we shall use to illustrate our point, is computational geometry.

This is an essentially new branch of mathematics that derives from such sources as automata theory, pattern recognition and computational complexity. From these diverse origins, it has gradually moved towards a coherent conceptual personality. It so happens that the Artificial Intelligence Laboratory at M.I.T. is a focal point of development of computational geometry and of the thinking about education embodied in this proposal. So, it is not surprising that there should be an interaction, and that the best developed piece of mathematics in our new curriculum is "turtle geometry,"

which is firmly rooted in both areas.³

Enough other examples are germinating in our laboratory to convince us (as we suspected from general ideas) that many branches of computational mathematics will come into being and many of these will prove to have pedagogic virtues. This is not an appropriate place to develop this point; but we think we are entitled to emphasize how rare turtle geometry is amongst innovations in mathematical curricula in actually being a new piece of mathematics, created for children. This is starkly different from the general character of what is called the "New Math". The creators of that movement did not create very much (if any!) new mathematics, but rather seem to have scratched about in the mathematician's cupboard for fragments of already existing (and generally rather old!) mathematics that seemed suitable for children.

It is also appropriate to remark that computer scientists have not been any more imaginative in developing mathematical ideas for teaching their subject at any level. We see the traditional Boolean-Algebra-Finite-State-Machine-lambda-calculus course content as uninspired and uninspiring fare for young computer scientists. We would rather teach turtle geometry -- and other topics like it -- at all levels, whether in fifth grade or graduate school.*

³Remarks on Turtle Geometry are contained in all the appended papers, which use it as the best developed example to illustrate technical aspects of our curricula. The best exposition is in "Teaching Children To Be Mathematicians."

*A monograph by M. Minsky and S. Papert developing ideas of this sort will soon be published by the American Mathematical Society.

3. Immediate Research Plans

(a) General Problems and Proposals

We have scarcely begun to tap the three sources of educational innovation mentioned above. To a first approximation, our plan for the next year can be described as "more of the same kind." We find that some of our illustrations of new technologies are sufficiently compelling to attract considerable interest from research groups and educators including some who had become disenchanted with curriculum reform and computers. We believe this interest may start a wave of more research strong enough to deflect the national view of the role of computers in education. However, we have not yet carried the job far enough to be confident in this belief.

For our own future work, we have set ourselves the goal of developing enough material in the form of technology (like turtles) and ideas (like turtle geometry) to include the major part of a full elementary school curriculum in mathematics, science, music, circus arts, formal linguistics, and some new subjects one could call "cybernetic models" for use on the elementary school level.

On the technological side, this task is not as formidable as it might seem. For our engineering work does not any longer have the form of developing particular devices such as turtles. We have gained enough experience to shift over to developing modular kits out of which many different kinds of devices can be constructed. To this end we have already built a first model of a universal controller to allow extremely easy interfacing (of motors, relays, sensing devices, and so on) to the computer. We shall soon be able to install one of these in our elementary school computer laboratory, and the children will then be able to motorize under computer

control any construction they care to make, say with an erector set.

Developing theoretical ideas and teaching methods has become a strain on our computational resources. We have exciting ideas for a full physics course and a biology course, both deriving from the kind of project illustrated in section 17 of "Twenty Things To Do With A Computer." We have an extremely strong group working on teaching and studying music in the LOGO environment; and, we have an increasing pool of mathematical talent eager to develop new ideas. For these ideas to flourish, we need freer access than we can now provide for our research staff and students to computers and to children. We need a computer facility of our own, with real-time interactive facilities designed for our needs. At the moment, we rely on using a research computer setup at the Artificial Intelligence Laboratory for different purposes. This computer is now heavily overloaded -- so that, besides restrictions on our access to it, we suffer from poor response times that interfere with our real-time projects.

So, we have designed, and request funds for a new time-shared system based on a PDP-11 computer, which we shall equip with ten terminals including five displays suitable for turtle geometry and interfacing for real-time work of the kind we have already mentioned.

Our main reason for designing this system is that we need it for our own work. We have, however, kept in mind the possibility that it might serve as a model for general use in schools and other educational environments (colleges, museums, vocational

training, etc.) and have designed it (we believe without ever compromising our own needs) to facilitate this. The system will be very much ahead of the "state of the art" for mini-computers in power of language, in real-time facilities and in displays. On the other hand, its cost (excluding the once-only costs of engineering and programming) will be comparable to currently available systems. So, it will almost inevitably be used much more widely than in our laboratory. The important technical features of the system are:

(i) Language

The system will be dedicated to a version of LOGO extended to provide full list structure, algebraic infix notation and floating-point arithmetic.

(ii) Users

The system is designed for 10 users, but will probably be extendable to 16. The design takes advantage of the fact that many users will be children working with very small programs.

(iii) Real Time

The system will allow users to declare priorities for response time. The intention of the design is that the top priority program should be able to provide a response in less than 30 milliseconds for experiments such as the inverted pendulum discussed below.

(iv) Displays

The system will have five display tubes of adequate

quality for turtle geometry and similar graphics, but of considerably lower quality (and cost) than the standards usually considered necessary by the computer industry. (If suitable Plasma storage tubes become available in time we might switch to them.)

(v) Autonomy

The PDP-11 system will be able to operate autonomously, but will be connected by a single high speed line to a central computer facility to allow more convenient back-up storage, monitoring of programs for research purposes, etc.

(b) Examples of What a Future School Computer-Laboratory Might Contain

In this section we shall talk in a discursive and impressionistic way about kinds of new devices we see in a future school laboratory. We shall try to convey some images that guide our present thinking. But we warn the reader emphatically that we are firmly committed to the need for flexibility in this kind of research, and will certainly not stay by a pre-conceived idea when better prospects suggest themselves. So the following pages must not be read as a promise to do any particular thing. They are really trying to say that this is a rich area of research.

Build-An-Animal Kit

We imagine an environment in which children's study of biology is augmented and deepened by embodying various biological functions and mechanisms in cybernetic model animals. The idea of studying biology by building models is not, of course, new. Our contribution is to show how it can be done easily and cleanly by a child.

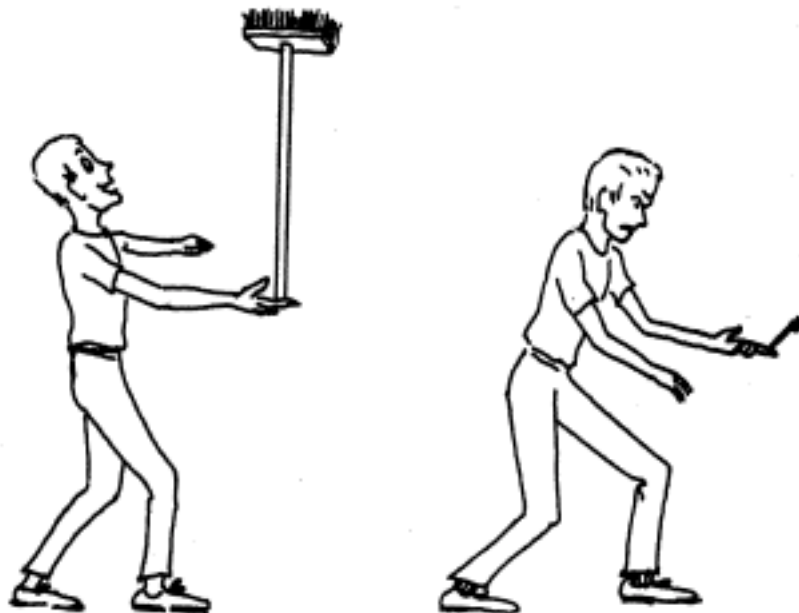
An aspect of what is meant by "clean" is illustrated by the relative importance in the construction of mechanical precision and of functional features. The child should not be frustrated by having to work to great mechanical precision. More important: the process of debugging sub-systems of the model should not be confused by bugs due to mechanical accidents.

A good example of a biological function suitable for study in this way (and perhaps in no other way at an elementary school level) is balance. To understand its mechanisms in any real sense one needs to have mastered a small set of very powerful concepts such as feedback, stability, momentum and a few others. In an abstract setting these concepts are difficult. Concretized in suitable projects in our kind of laboratory they are perfectly accessible. To illustrate the way in which one might start children off on understanding them we recall one of the projects mentioned in "Twenty Things To Do With A Computer" -- to program a turtle to balance an inverted pendulum.

The following illustrations from "Twenty Things..." show how we insert this in a limited project (the picture of the familiar turtle is symbolic -- this particular device happens to be too slow for this experiment).



Sketch of a Lesson Unit on "Physics in the Finger-Tips"

The problem is to find out what objects are easy to balance . . . and why. We begin with some crude experiments on balancing objects ourselves:



(Try to generalize to related situations: can a child or a giant walk more easily on a tight-rope? What length of stilts is best? Etc.)

The next two pictures show stages in the construction of "models" of this balancing problem . . . (making models is often a good way to understand things!).



What kind of program can make the turtle balance the rod? It is remarkable (but very typical of this kind of work) that an incredibly simple procedure will work under restricted but realistic conditions. This procedure (believe it or not!) is simply:

```
TO BALANCE
  1 FORWARD ANGLE
  2 BALANCE
END
```

The procedure uses a sub-procedure:

```
TO ANGLE
  1 OUTPUT PORT 5
END
```


This sub-procedure causes the information from "PORT 5" to become the "value" or "output" of the procedure ANGLE. We guess that the child must have selected the fifth port as the inlet for his wire from the potentiometer (see last picture) and that this provides a number between -180 and 180. We also see that he decided that the state of the system would be represented by this angle.

State is one of the prime concepts in our fifth grade courses -- whether or not we mean to use it in physical contexts. We see this as a good example of a fundamental concept that would appear absurdly abstract in the context of traditional schools and curricula, but which becomes simple, concrete and powerful in ours.

The next sub-project for the child is to understand exactly when this will "work." The reader will easily see that it will work only for certain lengths of rod, and for sufficiently gentle perturbing forces. He will less easily see that the procedure allows a drift in the turtle that will eventually bring it to the end of its permitted travel. Each of these observations leads to a modification or extension of the procedure. When it is eventually debugged it may work too well i.e. much better than the boy in the original problem! The final step will be to degrade its performance -- perhaps by limiting its reaction time to the human range.

Sketch of an Extension to Postural Balance

People can stand up without moving their feet . . . so they have a different balancing method. But its principles are the same. A project directed at finding out how it works should be accessible to a student who has understood the previous experiment provided that he has suitable components. A set of components we propose for our kit includes a joint that looks like two bars hinged together, but which has, hidden in the bars, a "muscle" in the form of an electrically driven actuator, and a sense organ. Information and power signals would be passed internally in the bars. These joints connect together easily, in the manner of standard "construction kits." The kit also contains passive bars, and bars with other sense organs, such as accelerometers, and possibly other kinds of actuators.



Project: Design a program to balance this device. The letter A marks two places at which accelerometers are attached.

Extension to Visually Controlled Balance

The kit would obviously contain simple light sensitive devices and could balance by using these as a source of information. An excellent experiment for children who are aware of these issues is to study whether people balance by vision.

Clue: People can balance with their eyes closed, even on the ball of one foot. But often fall over when placed in this situation while looking through right-left inverting spectacles.

Simple Vision, Tropisms, Etc., Etc.

The topic is obviously open-ended . . .

Make a Radio-Controlled Motorcycle

We leave the reader to work this one all on his own. The only clue we give him is that the control theory of balancing bicycles is much easier than popular beliefs (even those of physicists and psychologists) would lead one to suppose.