



The Real Computer Revolution Hasn't Happened Yet

Alan Kay

[Remarks on being awarded an honorary degree from the University of Pisa in Italy (June 15, 2007)]

VPRI Memo M-2007-007-a

Viewpoints Research Institute, 1209 Grand Central Avenue, Glendale, CA 91201 t: (818) 332-3001 f: (818) 244-9761

The Real Computer Revolution Hasn't Happened Yet

by Alan Kay

32 years ago in 1975 I was one of several lucky Americans who were invited to Pisa to help celebrate 20 years of computer science in Italy. I presented a paper on the first fruits of our attempts to invent personal computing at Xerox PARC. Over the years I somehow lost that paper but, Professor Attardi, who was more organized than I, was able to locate his copy and it has been republished as part of our ceremonies today. It is tempting in this talk to go through that paper and see how this past work influenced today. But I would much rather talk about future possibilities, and so I wrote a few historical notes to provide some context for the 1975 paper, and now can try to discuss some of the more important, and mostly hidden, gifts that personal computing networked together around the world can bring to humanity.

One connection to the past is that the researchers who invented today's fundamental technologies of personal computers, bit-mapped screens, overlapping window, icon and pointing interfaces, object-oriented programming, laser printing, the Ethernet, and the Internet, were motivated by the highest transformational achievements of the printing press. Simply put, the press in the 15th century was first thought to be a less expensive automation of hand written documents, but by the 17th century its several special properties had gradually changed the way important ideas were thought about to the extent that most of the important ideas that followed and the way they were thought about had not even existed when the press was invented. The two most important ideas were the inventions of science and of new ways to organize politics in society (which, in several important cases, were themselves extensions of the scientific outlook).

These changes in thought also changed what "literacy" meant, because literacy is not just being able to read and write, but to fluently deal with the *kinds of ideas* important enough to write about and discuss. One of the special properties of the press was its ability to absolutely replicate the corrected galley of an author, and this allowed a very different form of argumentation to arise. One way to look at the real printing revolution in the 17th and 18th centuries is in the co-evolution in *what* was argued about and

how the argumentation was done. Increasingly, it was about how the real world was set up, both physically and psychologically, and the argumentation was done more and more by using and extending mathematics, and by trying to shape natural language into more logically connected and less story-like forms.

One of the realizations we had about computers in the 60s was that they give rise to new and more powerful forms of arguments about many important issues via *dynamic simulations*. That is, instead of making the fairly dry claims that can be stated in prose and mathematical equations, the computer *could carry out the implications of the claims* to provide a better sense of whether the claims constituted a worthwhile model of reality. And, if the general literacy of the future could include the writing of these new kinds of claims and not just the consumption (reading) of them, then we would have something like the next 500 year invention after the printing press that could very likely change human thought for the better.

These were very ambitious aspirations indeed! The time from the invention of the press to the big changes in the 17th century was about 150 years, and this meant that the revolution in larger society happened because children gradually grew up with the different outlook of being able to think, argue, learn, and communicate in terms of the crisply written word in gradually more connected forms. Our idea about this came from a visit in 1968 with Seymour Papert, a mathematician who had invented the LOGO programming language for children and was starting to show that there were certain forms of advanced mathematics that when presented in a dynamic computer form were perfectly matched with the way children could think.

As McLuhan had pointed out in the 50s, when a new medium comes along it is first rejected on the grounds of “too strange and different”, but then is often gradually accepted if it can take on old familiar content. Years (even centuries) later, the big surprise comes if the medium’s hidden properties cause changes in the way people think and it is revealed as a wolf in sheep’s clothing. These changes are sometimes beneficial (I think the printing press was, though the Catholic Church would probably disagree), and sometimes not (I think television is a disaster, though most marketing people would disagree).

So we had a sense that the personal computer’s ability to imitate other media (and inexpensively via Moore’s Law) would both help it to become established in society, and

that this would also make it very difficult for most people to understand what it actually was. Our thought was: but if we can get the children to learn the real thing then in a few generations the big change will happen. 32 years later the technologies that our research community invented are in general use by more than a billion people, and we have gradually learned how to teach children the real thing. But it looks as though the actual revolution will take longer than our optimism suggested, largely because the commercial and educational interests in the old media and modes of thought have frozen personal computing pretty much at the “imitation of paper, recordings, film and TV” level.

Meanwhile, what the computer can really do – both in terms of simulation and argumentation – has been taken up by the scientific, mathematical, engineering and design disciplines. And those interested in Papert’s visions for changing the nature of children’s thought for the better, by helping them learn “powerful ideas” via actually constructing them, have made quite a bit of progress over the last 3 decades. There is now much to say, show and teach about “what children can do”.

The shame of the computer vendors – both hardware and software – is that no commercial intellectual amplifier for children has been made. All of the machines and software tools are primarily aimed at business, and to some extent for the home. This is a gross ignorance of the needs of the world of the most disastrous kind. It’s the new ideas and ways of thinking that the children of the world need – a children’s computer is imperative because this is now the best way to help children learn these new ideas – and is also much less expensive than paper for books and other conventional media.

Two years ago several members of the 60s research community that invented personal computing decided to create an extremely inexpensive personal laptop computer – a Dynabook – for all the children of the world. This initiative – called One Laptop Per Child – was started by Nicholas Negroponte and involves researchers old and new, including Seymour Papert, our research institute, and many other interested and committed designers, all aimed at bypassing the huge gaps created by commercial interests.

This community has always been willing to design and build whatever is needed, regardless of whether vendors have the tools and materials or not. The Alto at Xerox PARC is a good case in point. All the hardware and all the software were done at PARC and the little assembly line that was set up eventually built about 2000 of these first modern

personal computers. Today, most laptops are built in Taiwan or mainland China, and different brand names (such as HP, Dell, Sony, and Macintosh) may all be built by the very same offshore manufacturer. So, if you want to make your own laptop, just put a design together, get some orders for a million or so, and get on a plane to Taiwan! The OLPC goal is to make a very inexpensive machine that can furnish full function, so it is interesting to see where the money you pay for your laptop is allocated.

For example, about 50% of the price of a standard laptop is in sales, marketing, distribution, and profit. OLPC is a non-profit and sells directly to countries. Another 25% of the price is from commercial software, much of it from Microsoft. But there is a world-side open source and free software community that is equivalent in many respects, especially for web and educational environments. The expensive items that remain include a disk drive and the display. But the Flash memory used in cameras and memory sticks can be less expensive than the least expensive disk drive (and is much more robust because it is solid state). The display is a special problem, because cost is not the only issue. A display for the 3rd world must require much less power and has to be visible in direct sunlight with the backlight turned off. OLPC researcher Mary Lou Jepson solved the display problem brilliantly by inventing a new kind of flat-screen display which has higher resolution than regular displays (200 pixels/inch), 1/7th the power, and 1/3rd the cost.



The result is a computer that currently costs \$170, can hold hundreds of books (many of them dynamic), at a cost of about 20 cents per book, and maintains an automatic mesh internetwork with other laptops. This computer was pooh-pooed by the hardware and software vendors, but they have now started to make similar inexpensive offerings themselves (e.g. Intel now has a “\$400 laptop”, and Microsoft recently announced that they would sell their software to the 3rd world for a few dollars). It would be nice to say that they have now seen the light, but it is more likely that they now simply feel threatened, and are responding.

One of the nice properties of doing this with a non-profit organization is that as costs of materials go down and manufacturing is made less expensive, all of the cost savings will simply be passed along to the children. Also, the first phase of this project is being done

in just a little over two years, so many of the customizations of materials that could be done are set for the next phase. It is quite possible to make a \$50 laptop or even less if all the available technologies and manufacturing techniques were brought to bear.

Of course, the hardware part of this project is only a small part, even though it is challenging to make a full featured “\$100 laptop”. There is also system software, end-user authoring environments, educational content, various kinds of packaging and documentation, and, most importantly, the mentors that are needed to help the children learn the powerful ideas.

I’ll come back to this critical part of the educational ecology. For now, let’s note that – for mathematics and science in the 1st and 2nd worlds, the percentage of elementary school teachers and parents who really know math and science is much too small for most children to be guided over the threshold. In the 3rd world the percentage of knowledgeable mentors is vanishingly small.

This leads to a frustrating impasse. As I will demonstrate in a minute, it is now known how to help 10 and 11 year old children become very fluent in powerful forms of calculus and other advanced mathematical thinking. But no child has ever invented calculus! The wonderful nature of modern knowledge, aided by writing and teaching, is that many ideas which require a genius to invent (in the case of calculus: two geniuses) can be learned by a much wider and less especially talented population. But it is very difficult to invent in a vacuum, even for a genius. (Imagine being born with a 500 IQ in 10,000 BC. Not a lot is going to happen! Even Leonardo couldn’t invent an engine for any of his vehicles. He was plenty smart enough but he lived in the wrong time and thus *didn’t know enough*.)

If children have learned how to read, then they can bypass adults to some extent – both in home and in school – by going to a library and learning from reading. There are many such cases, and it is likely that a fair amount of the printing revolution gradually happened that way. But it is much more difficult for a child to learn how to read without the assistance (or at least the cooperation of adults), and again we see the critical importance of mentoring. When Andrew Carnegie set up thousands of free public libraries in the US, each one of them had a special room where *the librarians taught reading* to whomever wished to learn!

Of course, children can learn many things without special mentoring just by experimentation, and by sharing knowledge amongst themselves. But we don't know of any examples where this includes the great inventions of humanity such as deductive mathematics and mathematically based empirical sciences. To use an analogy: what if we were to make an inexpensive piano and put it in every classroom? The children would certainly learn to do something with it by themselves – it could be fun, it could have really expressive elements, it would certainly be a kind of music. But it would quite miss what has been invented in music over centuries by great musicians. This would be a *shame* with regard to music – *but for science and mathematics it would be a disaster*. The special processes and outlook in the latter (particularly in science) are so critical and so hidden that it is crippling not to be taught them as “skills which allow the art”. As Ed Wilson has pointed out, our genetic makeup for social interests, motivations, communication, and invention, is essentially what humans were in the Pleistocene. Much of what we call modern civilization is made from *inventions* such as agriculture, writing and reading, math and science, governance based on equal rights, etc. These were hard to invent, and are best learned via guides.

So, we simply have to find ways to solve the mentoring problem, not just for the 3rd world, but for the 1st and 2nd worlds also. We can easily make 5 million of the OLPC laptops this Fall, but we couldn't produce 1000 new teachers with the required knowledge and skills for any amount of money (in part because it takes years for human beings to learn and practice what they need to know). This is one of the reasons that education lags science, technology and other advances in ideas so badly.

It is sometimes shocking to see what even very young children are able to do when they are in a good environment for learning. The most important principles in early children's learning are to try to find out what they can do, what representations of ideas work best for them, and what kind of social environment triggers their built-in urges to get competent in the world they live in.

First grade teacher Julia Nishijima, whom we met at one of the schools we worked with 15 years ago, was a little unusual in that she was a natural mathematician.

We don't think that she had studied math formally or had actually had ever taken a calculus course. But she was like a talented jazz musician who has never taken formal lessons. She really understood the music of mathematics.

She had a natural mathematical outlook on the world and here's one of the most interesting projects we saw in her classroom. She had the children pick a shape that they liked, and the idea was to, using just those shapes, make the next larger forms that had the same shape.

Here are diamonds, squares, triangles and trapezoids.



Julia next got her students to reflect on their creations. She treated



math as a kind of a science to get the children to create structures with interesting mathematical properties and then analyze them. Six year old Lauren noticed that it took one tile to make the first one, the total number of tiles was one. It took three more tiles to make the next shape, and the total number of files was four.

	1	2	3
1	1	1	
2	3	4	
3	5	9	
4	7	16	
5	9	25	
6	11	36	
7	13	49	
8	15	64	
9	17	81	
10	19	100	

And it took five more tiles to make the next one. So pretty soon she saw "Oh yeah, these are just the odd numbers, differing by two each time, I add two each time and I'll get the next one of these".

And the sum of these is the square numbers, up to at least 6 times 6 here (she wasn't quite sure about seven times seven).

Lauren had discovered two very interesting progressions that all the mathematicians and scientists in the audience will recognize.

Then the teacher had the kids all bring their projects up to the front of the room and put them on the floor so they could all look at them, and the kids were absolutely amazed because all of the progressions were exactly the same!

Every child had filled out a table that looked like Lauren's and it meant that the growth laws for all of these were exactly the same and the children had discovered a great generalization about growth.



	1	2
1	1	1
2	3	4
3	5	9
4	7	16
5	9	25
6	11	36
7	13	49
8	15	64
9	17	81
10	19	100

The mathematicians and scientists who are reading this will recognize that the odd numbers are produced as a first order differential relationship that gives a smooth, uniform progression – in the computer mathematics that we use, this is expressed as *doing over and over* with *increase-by*:

Doing over-and-over: Odds increase-by 2

And the Total number of tiles is produced by a second order (because it uses the results of a first order relationship) differential relationship:

Doing over-and-over: Totals increase-by Odds

The idea of *increase-by* is easy for everyone, because it is just putting things (that is, adding things) into a pile of things.


Mathematics is really “careful thinking about how representations of ideas could imply other representations of ideas”, and the most important process in helping anyone learn how to do mathematical thinking is to put them in many situations in which they can use how they think right now in a more careful way. We can see that the children were able to find a very nice way to think about two kinds of growth and change. “Increase-by” is very powerful idea because many of the changes in the physical world can be

modeled by one or two “increase by”s, and it is a representation that children can readily understand.


Now, what would be a really good representation to help children think about the Pythagorean Theorem?

Below left is Euclid’s proof of the Pythagorean Theorem for high school geometry students. It’s elegant and subtle, it casts light on other areas of geometry, but it is not a suitable first proof for most young minds.

Below right is a very different kind of proof: perhaps the original one by Pythagoras. We have seen many elementary-age children actually find this proof by playing with triangle and square shapes. Show the arrangement, surround the C square with 3 more triangles, copy the larger square, remove the C square, rotate the two triangles, note there is room for the A and B squares, move them, and Bingo! This proof has a visceral approach and feel, a *powerful simplicity* that is perfect for beginners’ minds and provides a solid foundation for later more abstract and subtle looks at the idea.



This?




Or This?

The rule of thumb here is to find ideas and representations that allow “beginners to act as intermediates”, that is, for learners to immediately start doing the actual activity in some real form.

The way calculus looks at many ideas is so powerful and so important that we want to start children learning to think along these lines much earlier in their lives. So we have made a form of real calculus that is thinkable by young minds and that the computer brings alive in many delightful ways.

A project that nine, ten and eleven year old children all over the world love is to design and make a car they would like to learn how to drive. They first draw their car (and often put big off-road tires on them like this).



So far this is just a picture. But then they can look “inside” their drawing to see its properties (for example where the car is located and heading) and behaviors (the ability to go forward in the direction it is heading, or change its heading by turning). These behaviors can be pulled out and dropped on the “world” to make a script – *without the need for typing* – which can be set “ticking” by clicking on the clock. The car starts moving in accordance with the script.

If we drop the car's pen on the world, it will leave a track (in this case a circle), and we



see that this is Papert's LOGO turtle in disguise – a turtle with a “costume” and easy ways to view, script and control it.


To drive the car, the children find that changing the number after **car turn by** will change its direction.



Then they draw a steering wheel (the very same kind of object as the car, but with a different costume) and see that if they could put **steer's heading** right after **car turn by ...** this might allow the steering wheel to influence the car.



They can pick up **steer's heading** (the name for the heading numbers that the steering wheel is putting out) and drop it into the script. Now they can steer the car with the wheel!




The children have just learned what a variable is and how it works. Our experience indicates that they learn it deeply from just this one example.




They quickly find that it is hard to control the car. They need to introduce a “gear” into the wheel's connection to the car. They can get the needed advice from a teacher, parent, friend, or from a child thousands of miles away via the mentoring interface over the Internet. They open the expression in the script, and divide the numbers coming

out of the steering wheel by 3. This *scaling* makes turns of the steering wheel have less influence. They have just learned what divide (and multiply) are really good for.




Quite a bit of doing is “just doing”, so it is a good idea to also reflect on what just happened. One way to do this is to have the objects leave trails that show what they were doing over time.

If the speed is constant then the trail of dots is evenly spaced, showing that the same distance was traveled in each little tick of time.



If we increase the speed each tick of the clock, we'll get a pattern that looks like the second picture. This is the visual pattern for uniform acceleration.



If we make the speed be random each time (in this case between 0 and 40), we will get an irregular pattern of distance traveled for each tick



It's fun to try using random in two dimensions. Here we can move a car forward randomly and turn it randomly. If we put the pen down, we get a series of “Drunkard's Walks”.



One way to look at this is that a little bit of randomness will cause a lot of territory to be visited given enough time and it changes the probability of collisions quite a bit (it is still low, but now “more possible”).

So we might guess that a random walk in a script that is using the principles of feedback could accomplish surprising things.

Up to this point our examples have been essentially mathematical, in that they are dealing with computer representations of relationships of ideas.

These ideas can be similar to the real world (the models of speed and acceleration), or different from the real world (the car in the speed and acceleration examples was supported by nothing, but didn't fall because there is no “gravity” in the computer world unless we model it). Sometimes we can make up a story that is similar to the real world, and even has a guess that works out. But for most of human history, the guesses about the physical world have been very far from the mark.


The physical sciences began in earnest when people started to do careful observations and measurements of the physical world, first to do accurate mapping for navigation, exploration and trade, and then to look more closely at more and more phenomena with better instruments and technique.

Another good example of “high-noticing low-cost” is the measuring of the circumference of the bicycle tire project for 5th graders. Much of the philosophical gold in science is to be found in this noticing activity. The students used different materials and got different answers, but were quite sure that there was an exact answer in centimeters (partly because schooling requires them to get *exact* rather than *real* answers).



String	155 -159 cm
Yarn	157-162 cm
Roll on Ruler	156-159 cm
Roll on Floor	153-164 cm
Cm cubes	157 cm
Manufacturer	159.6 ±1mm

What is the circumference of a bike tire?



What is the length of a shore line?

One of the teachers also thought this because on the side of the tire it said it was 20 inches in diameter. The teacher “knew” that the circumference was $\pi \times \text{diameter}$, that “ π is 3.14”, and “inches times 2.54 converts to ‘centimeters’”, etc., and multiplied it out to get the “exact circumference” of the tire = 159.512 cm. We suggested that they measure the diameter and they found it was actually more like $19 \frac{3}{4}$ inches (it was uninflated)! This was a shock, since they were all set up to believe pretty much anything that was written down, and the idea of doing an independent test on something written down had not occurred to them.

That led to questions of inflating to different pressures, etc. But still most thought that there was an exact circumference. Then one of us contacted the tire manufacturer (who happened to be Korean) and there were many interesting and entertaining exchanges of email until an engineer was found who wrote back that “We don’t actually know the circumference or diameter of the tire. We extrude them and cut them to a length that is $159.6 \text{ cm} \pm 1 \text{ millimeter tolerance!}”$

This really shocked and impressed the children – the maker of the tire doesn’t even know its diameter or circumference! – and it got them thinking much more powerful thoughts. Maybe you can’t measure things exactly. Aren’t there “atoms” down there? Don’t they jiggle? Aren’t atoms made of stuff that jiggles? And so forth. The analogy to “how long is a shoreline?” is a good one. The answer is partly due to the scale and tolerance of measurement. As Mandelbrot and others interested in fractals have shown, the length of a mathematical shoreline can be infinite, and physics shows us that the physical measurement could be “almost” as long (that is very long).

There are many ways to make use of the powerful idea of “tolerance”. For example, when the children do their gravity project and come up with a model for what gravity does to objects near the surface of the earth, (see the next project) it is very important for them to realize that they can only measure to within one pixel on their computer screens and that they can also make little slips. A totally literal take on the measurements can cause them to miss seeing that uniform acceleration is what’s going on. So they need to be tolerant of very small errors. On the other hand, they need to be quite vigilant about discrepancies that are outside of typical measuring errors. Historically, it was important for Galileo not to be able to measure really accurately how the balls rolled down the inclined plane, and for Newton not to know what the planet Mercury’s

orbit actually does when looked at closely.

Children Discover, Measure, and Mathematically Model Galilean Gravity

A nice “real science” example for 11 year olds is to investigate what happens when we drop objects of different weights.

The children think that the heavier weight will fall faster. And they think that a stopwatch will tell them what is going on.

But it is hard to tell when the weight is released, and just when it hits.

In every class, you’ll usually find one “Galileo child”. In this class it was a little girl who realized: well, you don’t really need the stop watches, just drop the heavy one and the light one and listen to see if they hit at the same time. This was the same insight that Galileo had 400 years ago, and apparently did not occur to any adult, (including the very smart Greeks) for our previous 80,000 years on this planet!

To really understand in more detail what is going on with gravity near the surface of the Earth, we can use a video camera to catch the dynamics of the dropping weight.


We can see the position of the ball frame by frame, $1/30^{\text{th}}$ of a second apart. To make this easier to see we can just pull out every fifth frame and put them side by side:



Dropping objects



The “Galileo Girl” explains a simple way to see if different weights fall at the same or different speeds.




The stacked up frames that reveal the acceleration pattern

Another good thing to do is to take each frame and paint out the non-essential parts and then stack them. When the children do this, most of them will immediately say “Acceleration!” because they recognize that the vertical spacing pattern is the same as the horizontal one they played with using their cars several months before.

But, what kind of acceleration? We need to measure.

Some children will measure directly on the spread frames, while others will prefer to measure the stacked frames.



The translucent rectangles help because the bottom of the balls can be seen more accurately. The height of a rectangle measures the speed of the ball at that time (speed is the distance traveled in a unit of time, in this case about $1/5^{\text{th}}$ of a second).

When we stack up the rectangles we can see that the difference in speed is represented by the little strips that are exposed, and the height of each of these strips appears to be the same!

These measurements reveal that the acceleration appears pretty constant, and the children made such scripts for their car months ago. Most quickly realize that since the ball is going vertically that they have to write the script so that it is the vertical speed that is increased and the vertical position y that is changed. They paint a small round shape to be the simulated ball, and write the script:



Now, how to show that this is a good model for what they have observed? 11 year old

Tyrone decided to do as he did with his car months before: to leave a dot copy behind to show that the path of his simulated ball hit the very same positions as the real ball in the video. Here is what he had to say when explaining what he did and how he did it:

And to make sure that I was doing it *just right*, I got a magnifier which would help me figure out if I had it - if the size was just right.

After I'd done that I would go and click on the little basic category button and then a little menu would pop up and one of the categories would be Geometry, so I clicked on that.


And here it has many things that have to do with the size and shape of the rectangle. So I would see what the height is... I kept going along the process until I had them all lined up with their height.

I subtracted the smaller one's height from the bigger one to see if there was a kind of pattern anywhere that could help me out. And my best guess worked: so in order to show that it was working, I decided to make – to leave – a dot copy (so that it would show that the ball was going at the exact right speed. And acceleration.)

An investigative work of beauty by an 11 year old!


In the United States, about 70 percent of the college kids who are taught about gravity near the surface of the Earth fail to understand it. It's not because the college kids are somehow stupider than the fifth graders, it's because the context and the mathematical approach most college kids are given to learn these ideas are not well suited to the ways they can think. We have found that more than 90% of 5th graders are able to understand by using this better context and representations for change.

Now that the children have “captured gravity”, they can use it to explore other physical situations and to make games. If the ball is repainted into a spaceship, and a moon is made that can be landed on, it is pretty easy for a 12-year old (and most 11 year olds)



to make the classic game “Lunar Lander”. The gravity script is standard and will accelerate the ship downwards to crash on the moon. The children can add a rocket motor script controlled by the joystick that accelerates the ship upwards. Note that in each case *ySpeed* is being increased on one direction or another.

The children put in some nice frills such as a crash if the spaceship touches the moon with its downward speed too high, and to show a flame when the spaceship rocket motor is on.



A lot of physical phenomena can be modeled by children using “increase-by” including: inertia, orbits, springs, etc. But now let’s take a look at a different powerful idea: one that allows progress to be made without enough information to make a complete plan.

Sometimes we have enough information to make a foolproof plan. But most of the time things don’t go quite as expected, (even with our “foolproof” plans), and we wind up having to find new information, make new corrections, sometimes new plans.


All animals and other mechanisms have limited abilities to gather information and very little ability to extrapolate into the future. For example, the simplest bacteria can be hurt or killed by too much acidity (or too little). They have evolved molecular machines that can help them detect when some dangerous substances are starting to affect them, and the ones that swim exhibit a tumbling behavior that radically (and randomly) changes the direction in which they are swimming. If things are “good” they don’t tumble, if “bad” then they tumble again.

This general strategy of sensing “good” and “bad” and doing something that might make things better is pervasive in biology, and is now in many of the mechanical and

electrical machines made by humans.

A fun activity for children is to be challenged to find their way around the outside of one of their school buildings blindfolded and just by touch. The simplest strategy works, follow the wall, and always turn in the same direction when it is lost.

We can get our cars to do that by painting a different color to be used as a “touch sensor” and writing a script that looks like:



Then we challenge the children to make a car and a road that will get the car to go down the center of the road instead of the outside. There are many solutions for this. Here is a nice one from two 11 year old girls who worked well together.



They figured out that if they made curbs on the road of two different colors then there would be only three cases: when the sensor is on the center, or one of the two sides. Their car, road and script look like:



We can see that this is a better script than the one we showed them. They decided that their robot car would only go forward when it is in the middle. This means that it can safely negotiate any turn (the first example can't always do this because the turns have a constant radius of 5).

Now let's model typical animal behavior that is used to follow chemical signals in the environment by being able to sense the relative concentration of the chemical and to be able to remember a past smell and concentration well enough to decide to keep going or to try a different path.

We'll pick a salmon swimming upstream to lay its eggs in the same part of the river in which it was born.



In our model we'll avoid the dramatic leaping backwards up the waterfalls, and concentrate on how the salmon might be able to guide itself by smelling a particular chemical from its spawning ground and being able to remember just the concentration of the last sniff it took.

To model the water with the chemical in it, we'll use a gradient of color, where the darker the color, the more concentrated is the chemical. Etoys lets us sense not just the color under an object, but also the brightness. So, for this simulation, less brightness means "getting closer".



salmon sniffDarker ! paused

salmon forward by 5

Test salmon's lastSeen > salmon's brightnessUnder

Yes

No salmon turn by random 90
salmon turn by random -90

salmon's lastSeen + salmon's brightnessUnder


This is the classic strategy used in most animals from bacteria on up to make progress with incomplete information from the environment.

1: keep moving
2: if things aren't better, do something random
3: remember last information from environment.


Below we see the salmon has successfully found the darkest corner, and to the right we see the path it took. The script under the path is a "turtle" following the salmon's posi-

tion and drawing a trail in a different playfield. The holder and script under the “river” animate the salmon’s body to wriggle as a fish does.


This simple scheme of “try something, test, keep going if OK, do something random if not” is found in most living creatures from bacteria on up, and a kind of abstraction of it is what makes evolution work.




Although there is a certain kind of uniqueness to each phenomenon in the world, most things are best understood by thinking of them as members of species which share most of their properties and behaviors. Since computers are very good at copying things quickly and inexpensively, we can use this to turn a model of an individual into one that has many actors in it. For example, we can introduce as many salmon as we wish into our model. This means that we can model population ecologies as well as individuals.




Ants are a wonderful example of an animal that can be studied and modeled by children. They use their ability to sense and follow gradients to communicate with each



Massively Parallel Particle System can handle tens of thousands of moving and background particles. This is a simulation of ants gathering food, and leaving a diffusing trail of pheromones to guide other ants.




This is the "master" script that fires off all the rest of the scripts for each ant. It is very simple.




If the ant has found food it will plot a course back to its nest and level a pheromone trail of diffusing evaporating odor. The diffusion is done similarly to the dye in the goldfish bowl.


other by laying scent trails to mark paths to help other ants find "interesting things" (usually food). Ants are a kind of "swarm animal" and often act like a larger organism whose "cells" can independently sense, think and do.



Massively Parallel Particle System can handle tens of thousands of moving and background particles. This is a simulation of ants gathering food, and leaving a diffusing trail of pheromones to guide other ants.



The wander script is the classic random walk that children use for many kinds of investigations.



If the ant has not found food, it just wanders around looking for food or for a pheromone trail left by other ants. Scientists are pretty sure that most food is found when the randomly stumbles into it.

A compelling example of “new powerful argumentation” is how self-authored computer simulations can help illuminate a very complex and hard to understand threat – for example, slow moving but deadly epidemics such as AIDS – in ways that transcend simple claims and assertions. One of the problems with the understanding of AIDS is that the incubation period for the virus is as long as five or more years. To normal human commonsense it appears that “nothing is happening” early on when the most positive actions can be taken to thwart the contagion. In many traditional societies around the world commonsense was heeded, no actions were taken, and the result some years later was and continues to be devastating. One of the simplest investigations that can be done with an authoring system which can manifest thousands of elements is to make a simulated village and try different scenarios with contagion, slow moving incubation and no cure. In the slow incubation cases they can see that “nothing appears to be happening” in the beginning, and yet all eventually die. Because the learners have set up the simulations themselves, they have created their own dynamic mathematics and modeling of something they care about. They get to choose the initial conditions and the emotional impact of the same disastrous endings has a deep effect on their view of epidemics.

There are now thousands of Etoys projects done in their native languages by children from many countries of the world, including: USA, Canada, Mexico, Argentina, Brazil, France, Germany, Spain, Japan, Korea, China, Nepal, and more. The low-cost OLPC laptop and others that it inspires will soon be spreading to millions of children. So, as with the advent of the book in the 15th century, the *potential* for completely changing learning – and, as McLuhan pointed out, of *changing change* itself – in most parts of the world has arrived.

We started our research about 40 years ago with the goal to help children – and thus humanity – to learn and absorb “science in the large”. We think of science as all the processes that can help “make the invisible more visible”. By *invisible* we mean what is invisible to we human beings for all reasons, including not just the usual scientific objects of interest that are too small or too far away or emanate in wavelengths we can’t perceive, but also ideas and objects that are invisible to us because our mental apparatus is not up to thinking about them, or has rejected them (because they “can’t possibly be true”), etc.

I include all the “serious arts” whose purpose is to wake us up, to get us to realize that what our consciousness presents to us is not reality but a *story* that may be very far from reality, sometimes dangerously far. What science does is not so much to change or fix our noisy mental apparatus but rather to add additional processes inside our heads (and outside in the society of scientists) to detect our many errors and try to reduce them in size and kind.

As Thomas Jefferson put it: “The moment a person forms a theory, his imagination sees, in every object, only the traits that favor that theory”. The larger society of science acts as a kind of “superscientist” – and far beyond just “knowing” more than any individual. In this superorganism are better, more skeptical debuggers of ideas than most individuals have in their own minds. The superorganism has more points of view on how the universe might work than any individual, and these are very useful (even if some of their motivation might have been less than scientific). So, without needlessly anthropomorphizing science, we are quite justified in saying that “science” is smarter, more knowledgeable, has stronger outlooks, and is “a better scientist” than any individual.

A larger society can also act smarter and be less prone to disastrous decisions and needlessly aggressive actions than most individuals. And it is the aim of education in democratic societies, particularly democratic republics whose representatives must be chosen by the whole society, to bring all the citizens into the strongest thinking processes, conversations and debates that have been invented. Jefferson again:

I know of no safe repository for the ultimate powers of society but the people themselves; and if we think them not enlightened enough to exercise their control with a wholesome discretion, the remedy is not to take it from them, but to increase their discretion by education.

H.G. Wells said that “Civilization is in a race between education and catastrophe”. Perhaps “education” is too vague a term here. I would replace it with “a race between *education of outlook* and catastrophe” because it is not knowledge per se that makes the biggest difference, but *outlook* or *point of view* which provides the context in which rational thinking actually matches up with the real world in the service of humanity. For example, a context in which some human beings were regarded as a kind of non-human vermin was allowed to be set up in the 20th century, and since we exterminate vermin it was logically decided in this horrendous context to exterminate the human beings in

question. This has not been at all rare in human history, was done more than once in the 20th century, and is happening today. Slavery is another outgrowth of horrendous contexts and expediency and is still with us today in many forms.

The first step in science is the startling realization that “the world is not as it seems” and many adults have never taken this step but instead take the world as it seems and their inner stories as reality with often disastrous consequences. The first step is a big one, and is best taken by children (and most who crossed this threshold of awareness did it early in their lives). From there, it is another giant step to include ourselves (that is, all humans) in the proper objects of study: to try to get past our stories about ourselves to understand better “what are we?” and ask “how can our flaws be mitigated?”.

Though the world today is far from peaceful, there are now examples of much larger groups of people living peacefully and prospering for many more generations than ever before in history. The enlightenment of some has led to communities of outlook, knowledge, wealth, commerce, and energy that help the less enlightened behave better. It is not at all a coincidence that the first part of this real revolution in society was powered by the printing press. The next revolutions in thought – such as whole systems thinking and planning leading to major new changes in outlook – will be powered by the real computer revolution – and it could come just in time to win over catastrophe.