MITCHEL RESNICK

THINKING LIKE A TREE (AND OTHER FORMS OF ECOLOGICAL THINKING)

1. INTRODUCTION: THE WALKING TREE

In the rain forests of Costa Rica, there is an unusual type of tree known as a “walking tree.” It is a strange looking tree. At the base of the tree is a tangle of roots, rising about a meter above the ground. It looks as if someone yanked the tree straight up out of the ground, leaving about a meter of its roots exposed above ground level.

According to rain-forest guides, the walking tree actually changes its location over time (although very slowly). How does the tree move? The roots act as a type of evaluation system, searching for good soil for the tree. If there is good soil on the north side of the tree, the roots on that side dig in deeply and hold firmly. If the soil on the south side isn’t as good, the roots on that side remain shallow and weak. As the roots on the north side become stronger and deeper, the whole tree gradually shifts toward the north, pulled by the strong roots in that direction. As the tree moves, new roots grow around the new location, some of them extending even further to the north. If the roots find even better soil there, the whole tree will, over time, shift even more to the north. Or, if there is better soil to the east, the tree will slowly shift to the east.

We might say that the walking tree follows a TREE strategy:

- **T**est **R**andomly (send out roots in all directions)
- **E**valuate (determine which roots find the best soil)
- **E**lect (choose which direction to move, based on the information from the roots)

The walking tree executes this strategy over and over; as it moves, it continually sends out new roots to search the area around its new location. Over time, it moves in the direction of better soil. Of course, the walking tree does not actually “choose” or “decide” which way to move, as a person would. But it is useful to think of the tree as executing a type of strategy or algorithm.[1]

International Journal of Computers for Mathematical Learning **8:** 43–62, 2003.

2. ECOLOGICAL THINKING

The TREE strategy is representative of a broader class of strategies that I call "ecological strategies." These strategies are very common in the biological world, used not only by walking trees but many other plants and animals as well. Ecological strategies share two common characteristics:

- *Responsive to local conditions.* In ecological strategies, decisions (e.g., which direction to grow roots) are based on local information, not centrally-planned solutions.
- *Adaptive to changing conditions.* As conditions change (e.g., deterioration of soil on one side of a tree), ecological strategies adjust and produce new solutions tuned to the new conditions. There is no pre-planned script; decisions and solutions change over time.

Ecological strategies might seem inefficient and indirect, but they tend to be simple, flexible, and robust. Many ecological strategies employ decentralized approaches, relying on small contributions by (and interactions among) many simple entities (e.g., the roots of the tree), rather than a single, sophisticated decision-making entity.

Although ecological strategies are most commonly associated with the biological world, they can be useful in a wide variety of other situations – for example, designing management and organizational structures, solving mathematics problems, coordinating communications systems. In the 1940s and 1950s, the field of cybernetics aimed to apply ecological-style strategies to many different types of systems – biological, social, scientific, and technological (Wiener, 1948; von Foerster, Mead and Teuber, 1949). The field attracted engineers, biologists, psychologists, anthropologists, all aiming to forge connections among their disciplines. Cybernetics never developed into a mainstream discipline. But cybernetic ideas have attracted renewed interest during the past decade, as part of research efforts in the fields of complex systems (e.g., Waldrop, 1992; Gell-Mann, 1994) and artificial life (e.g., Langton, 1989; Levy, 1992). In particular, researchers are using TREE-like evolutionary models to describe phenomena in a very broad range of domains (Dennett, 1995).

As some scientists have pointed out (e.g., Pagels, 1988), these new research initiatives can be viewed, in part, as a reaction against the metaphors of Newtonian physics that have dominated the world of science for the past 300 years. Newton offered an image of the universe as a machine, a clockwork mechanism ruled by linear cause and effect. Today, some researchers are shifting metaphors, viewing their objects of inquiry less as clockwork mechanisms and more as ecosystems. Ideas from ecology, ethology, and evolution are spreading beyond their disciplinary boundaries, influencing research in fields from economics to engineering to anthropology.[2]

Despite this growing interest within the scientific community, ecological strategies have made few inroads into pre-college curricula. It is rare enough for biology classes to explain or model biological phenomena (such as ant foraging or bird flocking) in terms of ecological strategies; it is far, far rarer for ecological ideas to be used as a basis for design or problem-solving strategies in other (non-biology) classes. Even as educational reform efforts have placed increased emphasis on "problem solving," there has been little emphasis on ecological-style problem solving. Some "systems thinking" approaches (e.g., Senge, 1990) incorporate ecological strategies, but few schools have embraced these approaches.

Part of the problem is that people seem to have a "centralized mindset" (Resnick, 1994, 1996a), tending to gravitate towards explanations and solutions with a single centralized cause or linear causal chain. When people see a flock of birds, they assume the bird in front is the leader; when people create a new organization, they assume that hierarchical control structures are needed. People also seem to resist strategies and explanations that require probabilistic reasoning (Wilensky, 1993). Adopting ecological strategies requires not just a change in lesson plans, but a change in the mindsets of students, teachers, and curriculum developers. New approaches are needed to help people become "ecological thinkers."

In this paper, I probe the nature of ecological thinking, with special emphasis on the use of ecological strategies in design and problem solving (not just as a framework for explaining phenomena in the natural world). Through a set of specific examples, I suggest several different categories of ecological strategies. The goal is to begin to develop a framework that can help people understand the nature and uses of ecological thinking. The paper does not aim to provide a rigorous analysis of the effectiveness and limitations of ecological thinking. Rather, it is more of an essay: intended to provoke thought and to draw attention to styles of thinking that have, too often, been overlooked and undervalued.

The paper focuses on very simple examples of ecological thinking. In the scientific community, there is growing interest in a class of ecological strategies known as genetic algorithms (Holland, 1975; Mitchell, 1996). Genetic algorithms start with a "population" of possible solutions to a problem, then they combine the best-so-far solutions (using genetic approaches like mutation and crossover) in an effort to "breed" even better solutions. Genetic algorithms can be used to solve a wide range of optimization problems. But they are not particularly well suited to my goal of helping people understand ecological strategies (and becoming better ecological thinkers). Genetic algorithms are “black boxes”: even when they come up with good solutions to a problem, it is very difficult to analyze how and why they did so. My hope is to move beyond black boxes, finding “transparent” examples that help people see the key ideas underlying ecological thinking.

Many of the examples in this paper involve the use of computers (and computer networks). That is no accident. Ecological strategies often require repeated application of simple rules and/or massively-parallel interactions among many entities. Computers (and computer networks) are particularly well-suited to such tasks. Indeed, it is fair to say that ecological strategies are “truly computational” (Resnick, 1997). Most uses of computers in math/science education involve a re-implementation of traditional strategies that were previously implemented (albeit less efficiently) with paper and pencil. For example, system dynamics programs like Stella (Roberts et al., 1983; Doerr, 1996) are based on the same differential-equation representations that have long been used by mathematicians and scientists for studying the behaviors of systems. Ecological strategies are different: they involve very different representations and approaches than were traditionally used in the paper-and-pencil era. The point is not that computer-based ecological strategies allow people to do things that they couldn’t do before (though that it certainly true); the point is that such strategies allow people to do things that they never even thought to do before.

This link between computers and ecological thinking is ironic. The popular perception of computers places them in opposition to the natural world, but it is possible that computers could, in fact, lead to a much more widespread application and appreciation of strategies from the natural world.

3. GETTING TO THE ROOTS OF THE PROBLEM

Inspired by the story of the walking tree in the Costa Rican rain forest, I decided to introduce the idea of ecological thinking at a workshop for high-school teachers in Costa Rica. There were about 30 teachers at the workshop, many of them with backgrounds in math and science. Only a few of them had heard of the walking tree, and none of them knew how it moved. After explaining the walking tree’s strategy, I suggested that we could use a similar strategy to solve math problems. I explained that each teacher could act like one of the roots of the tree, and collectively they could solve a math problem. The idea was for each teacher to do something quite simple (just as each root of the walking tree does something quite simple), but for the group (like the overall tree) to accomplish a meaningful goal.

I proposed the following algebra problem:

$$2x^2 - 7x + 29 = 3104$$

Typically, students learn to solve such problems by using the quadratic formula:

$$x = \frac{-b \pm \sqrt{b^2 - 4ac}}{2a}$$

for an equation in the standard form of: $ax^2 + bx + c = 0$.

Students plug coefficient values into the formula and calculate the solutions for x. In this case, $a = 2$, $b = -7$, and $c = -3075$ (that is, $29 - 3104$).

The quadratic formula certainly yields the correct solutions for this problem. But there is a different, more ecological approach for solving this problem. To illustrate this alternative approach, I asked each of the teachers to pick a random number between 0 and 100. Then, I told them to calculate the left side of the equation, using their randomly-chosen number as the value for x. For example, if a teacher chose the number 3, then the calculation would be $(2*3*3) - (7*3) + 29$, for a result of 26. Next, I told them to compare their result with the right-hand side of the equation (3104). Of course, there was very little chance of an exact match: after all, they had chosen their values for x randomly. For the teacher who chose the value of 3, there would be an "error" of 3078 (that is, $3104 - 26$).

After the teachers had done their calculations, I asked whether any of them had "errors" of less than 1000. Three of them raised their hands. One had chosen the number 44 for an error of 489. Another had chosen the number 35 for an error of 870. The third had chosen the number 40 for an error of just 155.

I explained that these three teachers represented the "strong roots" of our tree. They had found good soil, and they should dig in deeply, pulling the tree in their direction. The other teachers had weak roots. They should pull up their roots and pick new numbers close to the strong roots (but not exactly the same). Since the strong roots ranged from 35 to 44, the teachers should all pick numbers between 30 and 50, and repeat the activity.

This time, several teachers chose the number 41 and got an exact match (that is, an error of 0). So the number 41 is a "very strong root" of the equation $2x^2 - 7x + 29 = 3104$. In effect, the roots of the tree had become the roots of the equation.

Just like the walking tree, the teachers had used a TREE strategy: Test Randomly (each teacher chooses random numbers as possible values for x), Evaluate (each teacher plugs his/her number into the equation and calculates the "error"), and Elect (the group selects the numbers that generated the smallest "errors") – then repeat the whole process (using numbers close to those Elected in the previous round).

For each teacher, the strategy was quite simple to execute. Each teacher performed only simple arithmetic operations (multiplication, addition, and subtraction). But the whole group, working together, was able to solve an algebra problem.[3]

Of course, the TREE approach to this problem has some clear disadvantages and limitations. Plugging numbers into the quadratic formula is a more "efficient" strategy, producing a solution more quickly; it is hardly convenient to gather together 30 colleagues every time you want to solve an algebra problem. Also, the quadratic formula always gives an exact solution; the TREE approach does not. If I had given the teachers a problem with irrational roots, they never would have converged on an exact solution using the TREE strategy; and if the problem had imaginary roots, the TREE strategy would not work at all. Moreover, the TREE strategy is not fully specified: depending on how the participants "test randomly" (how do they choose their next number?), the strategy might not converge on an answer.

These are significant limitations. On the other hand, there are several reasons for introducing teachers (and students) to the TREE strategy:

- The TREE strategy can be applied to wide range of problems. What if I had given the teachers a problem involving a fifth-degree (or other higher-order) polynomial? The quadratic formula would no longer work. Indeed, they could no longer look up a closed-form solution. But the TREE strategy would still be useful. And, as I will discuss later, the TREE strategy (and other ecological strategies) can be useful in many other (non-algebra) situations. The TREE strategy is not just a "math trick" but a general problem-solving strategy.
- The TREE strategy tends to be more robust than the strategies traditionally taught in schools. Even if one of the participants makes a arithmetic mistake (or, in the case of an actual walking tree, if one of the roots breaks or wanders off in the "wrong" direction), the overall group would still reach the same result. Small errors generally don't cause big problems in the TREE strategy. The TREE strategy also tends to be robust in another sense: people are likely to forget the details of the quadratic formula, but they are likely to remember the TREE strategy.

- The TREE strategy is well-suited for computerization – so it is not really necessary to gather together 30 friends to use the strategy. The Costa Rican teachers, after enacting the TREE strategy themselves, created a computerized version of the strategy written in StarLogo (Resnick, 1991, 1994), a massively-parallel version of the Logo programming language (Papert, 1980; Harvey, 1985). They created 100 StarLogo turtles,[4] and programmed the turtles to follow rules similar to those the teachers had followed in solving the algebra problem. Each turtle started at a random x-position, and it used its x-position as the x-value for the equation. As the turtles executed the TREE strategy, they gradually converged to the same x-value, indicating a solution to the equation. One reason that the TREE strategy has been underused and overlooked in classrooms is that it involves application of the same rules over and over again – a tedious task for humans. Widespread availability of computers makes the TREE strategy much more accessible (both practically and conceptually).
- Regardless of whether the TREE strategy is "better" than other strategies for finding roots of a polynomial, it provides learners an additional way of understanding the idea of finding roots. Minsky (1987) and others have observed that you don't really understand an idea until you understand it several different ways. Each way of thinking about something strengthens and deepens each of the other ways of thinking about it. Thus, the TREE strategy, in conjunction with other strategies, provides a more robust understanding.
- Perhaps most important, experiences with the TREE strategy can fundamentally change the ways that people look at the world. As discussed earlier, students usually learn top-down, centralized strategies for solving problems and designing artifacts. The TREE strategy represents a more decentralized approach: lots of separate parts, each with very simple behaviors, work together to produce complex-seeming results.

It is common for mathematicians and scientists to represent problems as topological terrains, consisting of hills and valleys. Locations with higher elevation represent better solutions to the problem. To solve a problem, you need to find the highest hills in the terrain. For a quadratic equation, the "problem terrain" is quite simple, with two large hills representing the two roots to the problem. When you use the quadratic formula, you jump directly to the top of the hills. The TREE strategy is more incremental: it is an example of what is known as a hill-climbing strategy. In the Costa Rica workshop, each teacher began by choosing a random location in the problem terrain, then looked for colleagues at higher elevations and moved closer to them. Collectively, the teachers searched the hillsides for higher elevations, until they reached the summit.[5]

4. WHOSE HILL ARE WE CLIMBING?

The TREE approach in the teacher workshop used a “parallel” algorithm: multiple people (or, in the case of the StarLogo implementation, multiple turtles) worked in parallel to solve the problem. But parallelism is not absolutely necessary in this case. It is possible for a single person (or a single turtle) to use the TREE strategy to solve the algebra problem, using the following rules:

- choose two random numbers
- figure out which one is the “stronger root” (that is, which one has the smaller “error”)
- choose a new number close to the “stronger root”
- see how this new number compares to the stronger root
- and so on . . .

Using the “problem terrain” representation, you can think of an individual climbing the hill on their own. Instead of looking for colleagues at higher elevations, the individual takes tentative steps in various directions to figure out which way is uphill, and gradually moves up the incline. This single-person approach (sometimes known as “successive approximation”) does not have the same “feel” as the multi-person approach: the Costa Rican teachers reported that they found the TREE strategy exciting in part because in enabled them to work together as a group. The multi-person approach has some other advantages too. The group might solve the problem more quickly (since it has more “processing power”). And if the problem has multiple solutions (multiple hills), the multi-person approach might find several solutions simultaneously. But the approach is not “deeply parallel” – that is, the single-person approach uses the same underlying idea.[6]

With some ecological strategies, however, parallelism is fundamental to the solution. Consider the following problem, which arose while two high-school students were working on a StarLogo project. The students were using StarLogo to explore the formation of traffic patterns on a highway (Resnick, 1994). They created several dozen cars, gave them simple rules to follow (“if there’s a car close ahead of me, slow down; if not, speed up”), then observed the traffic patterns that formed. At one point, the students wanted to start the cars in a single lane, evenly spaced along the road. This was not a trivial problem for the students. Ultimately, they came up with the following solution. They created a horizontal road on the computer screen, and divided the width of the screen by the number of cars to get the desired spacing between the cars. Then they set the x-coordinates of each car to maintain the appropriate spacing between cars.

This strategy clearly works. But then one of my graduate students, Randy Sargent, suggested an alternative (more ecological) strategy that the high-school students liked much better. In this strategy, the cars start at random positions along the road (all within a single lane). Then each car repeatedly applies the following rule:

- calculate the distance to the car in front of me
- calculate the distance to the car behind me
- take a step towards whichever one is further away

When all of the cars follow this strategy, the resulting motion has an organic feel to it. The cars start jostling around. Some of the cars move one way, then the other, as if they are making up their minds. Clusters of cars gradually spread out, filling in the more sparsely populated portions of the road. Eventually, the cars settle down in an evenly-spaced pattern. It feels almost like a group of school children, spacing themselves out for a game of Ring-around-the-Rosie.

This strategy has an aesthetic appeal; the high-school students found it more interesting to watch than their own program. The strategy is also more effective in some ways. Like all ecological strategies, it is more adaptive to changing conditions. After the cars have reached an evenly-spaced pattern, you can move one of the cars, or even remove a car entirely, and all of the other cars will automatically adjust, with changes rippling down the highway, until the overall system "relaxes" a new equilibrium.

This type of strategy (sometimes known as "constraint propagation") is not uncommon among scientific researchers. But it is usually viewed as an "advanced technique," and is rarely if ever taught in pre-college curricula. New computational environments like StarLogo not only make these techniques more understandable, but they also provide contexts in which these techniques are much more useful than ever before, even for young students. Suddenly, these techniques do not seem so "advanced" or esoteric anymore.

As suggested earlier, the ecological car-spacing strategy is "deeply parallel" in nature. What do I mean by "deeply parallel"? To answer that, it is again useful to think about problem terrains. In the algebra example, the group succeeds as soon as any individual finds a summit. The group goal is the same as the individual goal. The car-spacing example is more complicated. Each individual car is trying stay equidistant from its two neighbors. But even if it succeeds perfectly, the group hasn't necessarily succeeded. For the group to succeed, *all* cars must be equidistant from their neighbors. It is as if the group is moving on a different problem terrain than the individuals. Many of the individuals might be at the hilltops in their terrains, but the overall group won't necessarily be at the hilltop of the group terrain – and it probably won't get there until some of the individuals are knocked off of their individual summits. We can call a strategy "deeply parallel" when the group terrain is different from the individual terrain. There are no single-person versions for deeply-parallel strategies; these strategies depend on the interactions among the participants, not on the performance of any individual.

For me, the car-spacing example brought back a childhood memory. As a child, I loved watching baseball games. At one point, I wondered why the bases were placed exactly 90 feet from one another. I noticed that when a batter hit a ground ball to the infield, the play at first base was almost always very close. I thought that the choice of 90 feet as the distance between bases was a great choice: it led to lots of close and exciting plays at first base. It wasn't until many years later that I realized the error in my thinking. In fact, the infielders adjust their positions to maximize the number of balls that they can reach while still being able to throw out runners at first. If they often threw out runners by a large margin, it would mean that they weren't playing far back enough (and thus not reaching as many ground balls as they might). So the infielders adjust themselves in a way that ensures lots of close plays at first base. There is nothing magical about 90 feet. If the bases were placed 100 feet apart, the infielders would (ecologically) adjust their locations in a way that would again lead to lots of close plays at first base.

5. ECOLOGIES ON THE NET

In recent years, education-research conferences have been full of sessions about the educational implications of the Internet. Many researchers focus on how the Net will provide students and teachers with easy access to huge libraries of information. Other researchers focus on how the Net will make possible new types of learning communities, connecting people with shared interests from all over the world. But little is heard about what, in my mind, is the most important implication of the Internet: how the Net can support and encourage new ways of thinking – in particular, ecological ways of thinking. To the common metaphors of Internet as library, highway, and marketplace (Stefik, 1996), we should add the metaphor of Internet as ecosystem.

As is well known, the Internet is based on a decentralized structure. New computers, new users, and new functionality can be added to the Net without any centralized decision-making. The Net makes possible new types of decentralized collaborations, enabling large numbers of people to work together on shared tasks, such as decoding ciphers (Leutwyler, 1994), building online help systems (Whitehead, 1994), or organizing virtual libraries. In this way, the Internet makes it possible to leverage the small efforts of the many, rather than the large efforts of the few (Whitehead, 1994; Kalil, 1996). Parts of the Internet seem to function as artificial ecologies. For example, Best (1996) analyzed the newsgroups of NetNews in terms of ecological interactions. In this view, ideas on NetNews compete with one another for the attention of human readers and posters; certain ideas reproduce and flourish (and even spread to other newsgroups), while others die out.

The ecologies of the Internet could be a particularly fertile ground for the development of ecological thinking because they can be designed, manipulated, and analyzed much more easily than "natural" ecologies. As Papert (1993) has argued, people learn with particular effectiveness when they are actively engaged in the design and construction of personally-meaningful artifacts. The Internet enables people to design and play with "ecological artifacts" to a far greater extent than ever before.

Engaging children in ecological-style thinking was one of the motivations behind MOOSE Crossing (Bruckman, 1997), an online community organized by Amy Bruckman (at the time, a graduate student in my research group). MOOSE Crossing is a multi-person, text-based virtual world in which children not only interact with one another but also collaboratively construct the virtual world in which they interact.[7] In MOOSE Crossing, children (mostly between the ages of 9 and 13) create new rooms and objects – and write programs to control the behaviors of those objects (using a scripting language called MOOSE). For example, a ten-year-old girl created a pet penguin that reacts when other people kiss it, hug it, or feed it. The penguin keeps track of how hungry it is, and it reacts differently to six different kinds of food. Another MOOSE Crossing member created a set of potatoes that obey Mendelian genetics; others built a mega-mall with specialty shops. Children can also take on new personas: for example, a child might decide to become a munchkin and help others in building a replica of Oz. Studies of MOOSE Crossing have documented the ways in which a community can provide strong support for design and construction activities – and, conversely, the ways in which ongoing design and construction activities can support the development of a stronger sense of community (Bruckman, 1997).

I have been particularly interested in the ways that children on MOOSE Crossing get ideas for new projects. In any design situation, it is often a good strategy to start by looking at previously-developed projects, then to consider variations on these model projects. For this reason, it is important for designers to have easy access to good sample projects. Suppliers of computer-based design tools (such as paint programs and programming languages) often include sample projects with their tools, to help users get a sense of what's possible. But there are some important limitations to these pre-packaged sample projects: they reflect the interests and ideas of only the supplier (not necessarily of the full range of interests of the user community) and they are static (changing only if the supplier ships an updated version).

The collection of sample projects on MOOSE Crossing has a very different (and more ecological) feel. The sample projects are created by the MOOSE Crossing members themselves. In fact, since each object created in MOOSE Crossing is fully inspectable and copyable, each object becomes a sample project for everyone else in the community. If a MOOSE Crossing member sees an interesting object, she can "look inside" the object to see the computer code underlying the behavior – and, perhaps, create a new version of the object with slightly modified code.

The collection of sample projects in MOOSE Crossing continually changes, always reflecting the current interests of the community – without any centralized control. As with other ecological processes, the collection of sample projects automatically adapts and self-adjusts to changing conditions. If a group of MOOSE Crossing members becomes interested in a particular type of project, the collection of relevant sample projects automatically increases. At one point, for example, some MOOSE Crossing members became interested in magic. One member created a magic wand (and wrote a set of programs to accompany it); another created a "generic magician" as a new player class; a third created a spell book. The spell book was full of simple programs that could "cast spells" on other people in the room. Many children in the community made copies of the spell book, and many added new spells (programs) to their personalized spell books. Someone even created a spell that ran all of the other spells in the spell book – a popular program that was soon copied by many others in the community. One child opened a "magic store" where members of the community could get copies of the latest magic-related objects. Eventually, people started to lose interest in magic, and the number of magic-related objects gradually decreased over time – again, automatically adjusting to the current needs and interests of the community.

It is again useful to think about the problem-terrain representation. In this case, we can think of the problem terrain as a representation of how well the MOOSE Crossing sample projects match the current interests of the community. In the previous examples in this paper, there were clear, pre-defined optimal solutions: the exact roots of the quadratic equation, an even spacing among all of the cars. So the group terrains were fixed ahead of time; the challenge was to find the fixed hilltops. The MOOSE Crossing case is different. Since the interests of the community are always changing (as new members join and as existing members develop new interests), the problem terrain is always shifting. If an "outsider" tried to supply sample projects to meeting these shifting interests, it would be a very difficult challenge. The ecological strategy of letting the participants themselves create the sample projects does a much better job of staying close to the ever-shifting hilltops. Ecological strategies are especially useful in situations like this, where the problem terrain is constantly changing. In these situations, responsiveness to local conditions and adaptiveness to changing conditions are particularly important.

The hope (and this, admittedly, has yet to be proven) is that children who actively participate in artificial ecologies like MOOSE Crossing will be better prepared to use (and understand) ecological-style strategies in other situations. And with the rise of the Internet, there will be more opportunities to use such strategies. For example, in MediaMOO (Bruckman and Resnick, 1995), a networked virtual world intended for media researchers, one participant proposed an ecological strategy for automatically creating new (and useful) paths within the virtual world. The idea was to write a program that kept track of people's movements. If the program noticed that people often went from one particular room to another particular room (moving through several other rooms in between), it could automatically construct "short cuts" to allow people to jump directly from initial room to the final destination. Several universities tell stories of how an architect used a similar strategy to decide on the placement of walkways around the university library: the architect surrounded the library with grass and waited a year to see where the paths developed, then installed the permanent walkways. This approach allowed the community itself to decide on the placement of the walkways (in an informal, decentralized way). As the Internet allows more people to become "architects" of multi-user spaces, it will become increasingly important for people to learn how and when to use such ecological strategies.

6. ECOLOGICAL LEARNING ENVIRONMENTS

Ecological thinking can apply to education at several different levels. The previous sections focused on the need to help students develop as ecological thinkers – helping them learn how and when to use ecological strategies (especially in the context of new computational media). This section focuses at a different level, discussing how the ideas of ecological thinking can be applied to the design of learning environments themselves. Just as students too often adopt centralized strategies in trying to solve problems, educators too often adopt centralized strategies in designing learning environments.

In creating The Computer Clubhouse[8] (Resnick and Rusk, 1996), we explicitly tried to apply ecological ideas to the design of a learning environment. The Clubhouse is an after-school learning center for youth (ages 10–16) from under-served communities. At the Clubhouse, youth use new technological tools to work on projects related to their own interests and experiences. Clubhouse members, with support from volunteer adult mentors, become designers and creators (not just users) of technological artifacts. They create their own animated stories, robotic constructions, interactive newsletters, musical compositions, Web sites, and simulations. The goal is not simply to help youth develop new technological skills, but rather to help them develop new ways of thinking about the processes of thinking, designing, and learning, and also to help them explore mathematical and scientific ideas. For example, when Clubhouse youth use “programmable bricks” (Martin, 1994; Resnick, Martin, Sargent and Silverman, 1996) to create their own “robotic creatures,” they explore the differences and similarities between animals and machines, and they begin to develop intuitions about engineering concepts such as feedback – concepts that are traditionally taught at the university level.

The activity structures at the Clubhouse are based on ecological principles. Adult mentors play a very important role at the Clubhouse, but they do not plan activities in a centralized way. When Clubhouse members want to start on a new project, they often begin by looking at samples of previous Clubhouse projects (which are kept on display throughout the Clubhouse), then thinking about variations or extensions that they can work on. After that, their projects continue to evolve through various interactions – interactions with other Clubhouse members, interactions with mentors, and interactions with the media and materials on hand at the Clubhouse. Many projects involve groups of Clubhouse members working together, but we do not explicitly organize members into assigned teams, as is often done in classroom-based collaborative activities. Rather, we try to create an environment in which collaborative groups emerge as a natural part of ongoing activities. To a large extent, that has happened. At the Clubhouse, projects and project teams are not fixed entities; they grow and evolve over time. A member or mentor might start with one idea, a few others will join for a while, then some others will start working on a related project.

One Clubhouse project, for example, started with two Boston University graduate students who volunteered as mentors. The two students were both enthusiastic about robotics. Initially, they wanted to organize a workshop to teach Clubhouse members about robots. We discouraged that approach, and instead encouraged them to start by building their own robot at the Clubhouse. Our hope was that Clubhouse members would view the graduate students as fellow learners, not as traditional teachers.

For several days, the two graduate students worked on their own; none of the youth seemed particularly interested. But as the project began to take shape, a few youth took notice. One decided to build a new structure to fit on top of the robot, another saw the project as an opportunity to learn about programming. After a month or so, there was a small team of people working on several robots. Some youth were integrally involved, working on the project every day. Others chipped in from time to time, moving in and out of the project team. The process allowed for what Lave and Wenger (1991) have called "legitimate peripheral participation" – different youth were able to contribute to different degrees, at different times. In general, design teams at the Clubhouse form informally, coalescing around common interests. Communities are dynamic and flexible, adapting to the ever-changing needs of the project and interests of the participants – much as the collection of sample projects adapts in MOOSE Crossing.

As in natural ecologies, diversity is important to this process. At the Clubhouse, we have tried to attract a community of adult mentors with diverse professional and cultural backgrounds. One reason is obvious: a diverse mix of mentors can better match the diverse backgrounds and interests of the Clubhouse youth. But that is only one reason. There is also an evolutionary argument in favor of diversity within a learning environment. New projects at the Clubhouse emerge through a process related to Darwinian variation and selection. The "selection" of new projects works best when there is a rich "variation" in the combinations of youth, mentors, tools, and ideas. As in natural ecologies, diversity at the Clubhouse leads to a greater robustness and adaptiveness in the types of activities and projects that evolve.

Designing an ecological learning environment requires a shift in traditional ways of thinking about "control." Learning experiences can not be directly controlled or planned in a top-down way. Indeed, the specific experiences at the Clubhouse have been quite different from what we (as developers) expected. Educational designers can not control exactly what (or when or how) students learn. In some ways, the design of a new learning environment is like the design of a StarLogo simulation. To create a StarLogo simulation, you write rules for individual objects, then observe the large-scale patterns that emerge. You do not program the patterns directly. So too with the design of learning environments. Developers of learning environments can not "program" learning experiences directly. The challenge, instead, is to create fertile environments in which interesting activities and ideas are likely to grow and evolve.

7. SEEING THE FOREST

This paper has presented the idea of "ecological thinking" through a collection of metaphors and examples, not precise definitions or models. This approach runs the risk of misinterpretation. At a recent educational conference, for example, I described the foraging strategy of an ant colony as an example of an ecological strategy: the colony as a whole accomplishes complex tasks (bringing food back to the nest) and adapts well to changing conditions, even though each individual ant follows very simple rules and reacts only to local stimuli. In the same presentation, I described the uses of ecological thinking in the design of learning environments like the Computer Clubhouse. At the end of my talk, someone in the audience asked whether an ant colony is really a good model for a learning environment. Do we really want to think of students as ants, each following the same simple rules in an almost mindless fashion?

Certainly not. Students in a classroom are, of course, much different than ants in a colony – or roots on a walking tree. But in all of these cases, the behaviors of the overall system arise, often in unexpected ways, from interactions among the parts of the system. The defining characteristic of an ecological strategy is not the simplicity of the component parts, but the nature of the interactions among those parts. Developing as an "ecological thinker" requires a sensitivity to the role of interactions in a system – and an understanding that effective solutions are not always prescribed in a centralized way but rather arise indirectly from many interactions.

Another possible misinterpretation involves the role of non-ecological strategies. This paper aims to highlight the fact that ecological strategies have been overlooked and undervalued in the past. But it would be a mistake to read the paper as a rejection of all other (more "centralized") approaches. An exclusive focus on ecological strategies is no better than an exclusive focus on traditional centralized strategies – just as, in economics, an unyielding commitment to market mechanisms can cause just as many problems as an unyielding commitment to centralized planning. If you focus only on TREE-like strategies, you might not see the forest. There is no doubt that the quadratic formula is still a very useful strategy to learn, and teachers still need to play a centralized role in some classroom situations.

An important research challenge for the future is the development of a more systematic framework of how, when, and why ecological strategies are useful. The goal should not be to ignore or replace traditional strategies, but to expand the repertoire of strategies that people have at their disposal – and to help people learn which strategies (or which mixtures of strategies) are best suited to which situations. Indeed, one of the most important benefits of introducing ecological thinking in classrooms is to help students learn that there are, in fact, multiple ways of thinking about problems.

ACKNOWLEDGMENTS

Discussions with Brian Silverman, Seymour Papert, Alan Kay, Uri Wilensky, and Rick Borovoy have been helpful in developing these ideas. Alan Kay has also used the phrase “ecological thinking” to refer to a similar set of ideas. Eleonora Badilla organized the Costa Rican teacher workshop where I developed the “walking tree” example. Brian Silverman and Andy Begel have played major roles in the design and implementation of StarLogo. Amy Bruckman coordinated the development of MOOSE Crossing; she and Austina De Bonte provided me with useful insights into children’s activities on MOOSE Crossing. The section on ecological learning environments grew out of discussions with Natalie Rusk. I would like to thank Andy diSessa and two anonymous reviewers for their insightful comments on an earlier version of this paper. This research has been supported by the LEGO Group, the National Science Foundation (grants 9358519-RED, CCR-0122419, and CDA-9616444), the center for Bits and Atoms, and the MIT Media Laboratory’s Things That Think, Digital Life, and News in the Future consortia.

NOTES

[1] The story of the walking tree, which I heard from a rain-forest guide in Costa Rica, is actually a mixture of fact and fiction. The tree (genus Socratea, species S.exorrhiza) does move several centimeters during its lifetime, but its movement is not actually guided by soil quality. Nevertheless, the story of the walking tree provides a nice metaphor.

[2] This physics vs. biology dichotomy is, of course, overly simplistic. Some areas of physics (most notably statistical mechanics, which focuses on the patterns that arise from

local interactions among large numbers of elements) are very related to ideas underlying ecological thinking. And some areas of biology are based on centralized strategies. But the core ideas of ecological thinking (especially those related to adaptation and change) are most deeply rooted in the biological fields of ecology, ethology, and evolution.

[3] Note that the given equation has two roots: 41 and -37.5. The way the activity was organized, the teachers found only one of the two roots. But with slight modifications to the activity, allowing negative numbers and non-integers, the teachers could have found both roots, though the process would have taken longer.

[4] Turtles are computational objects with position and heading. With StarLogo, users can write programs for thousands of turtles, then observe the patterns that form from all of the interactions. You can download StarLogo from http://www.media.mit.edu/starlogo/

[5] Hill-climbing strategies have some well-known limitations: on certain terrains, they can get trapped on small hills (local maxima) and never reach higher elevations (global maxima). Depending on the exact details of the algorithm (e.g., how to make the “next guess”) and on the nature of the terrain, some hill-climbing strategies never converge. But it is not a goal of this paper to analyze the limitations of the TREE strategy (or other hill-climbing strategies) in detail.

[6] Papert (1996) makes a similar point about the so-called “rugby problem” (Resnick, 1996b; Wilensky, 1996), noting that the problem can be solved with a single turtle instead of multiple turtles.

[7] MOOSE Crossing is an example of a “MUD” (Curtis, 1993). The first MUDs were created to support online versions of the game “dungeons and dragons”; the acronym stands for “multi-user dungeons.” MUDs today are used for many different purposes, not just adventure games.

[8] The Computer Clubhouse is a joint project of the Boston Museum of Science and the MIT Media Lab. It was co-founded by Natalie Rusk and myself. The first Computer Clubhouse opened in 1993; there is now a network of more than 75 Clubhouses. For more information, see http://www.computerclubhouse.org

REFERENCES

Best, M. (1996). *An Ecology of the Net: Message Morphology and Evolution in NetNews*. Masters thesis, Cambridge, MA: MIT Media Laboratory.

Bruckman, A. (1997). *MOOSE Crossing: Construction, Community, and Learning in a Networked Virtual World for Kids*. Doctoral dissertation, Cambridge, MA: MIT Media Laboratory.

Bruckman, A. and Resnick, M. (1995). The MediaMOO project: Constructionism and professional community. *Convergence* 1(1): 94–109.

Curtis, P. (1993). Mudding: Social phenomena in text-based virtual realities. Presented at the *Directions and Implications of Advanced Computing (DIAC) Conference*, sponsored by Computer Professionals for Social Responsibility. Berkeley, CA.

Dawkins, R. (1976). *The Selfish Gene*. New York: Oxford University Press.

Dennett, D. (1995). *Darwin’s Dangerous Idea*. New York: Simon & Schuster.

Doerr, H. (1996). Stella ten years later: A review of the literature. *International Journal of Computers for Mathematical Learning* 1(2): 201–224.

Gell-Mann, M. (1994). *The Quark and the Jaguar: Adventures in the Simple and the Complex*. New York: W.H. Freeman and Co.

Harvey, B. (1985). *Computer Science Logo Style*. Cambridge, MA: MIT Press.

Holland, J. (1975). *Adaptation in Natural and Artificial Systems*. Ann Arbor: University of Michigan Press.

Kalil, T. (1996). Leveraging cyberspace. *IEEE Communications Magazine*, July, pp. 82–86.

Langton, C. (Ed.) (1989). *Artificial Life*. Redwood City, CA: Addison Wesley.

Lave, J. and Wenger, E. (1991). *Situated Learning: Legitimate Peripheral Participation*. Cambridge, UK: Cambridge University Press.

Leutwyler, K. (1994). Superhack. *Scientific American*, July, pp. 17–20.

Levy, S. (1992). *Artificial Life*. New York: Pantheon Books.

Martin, F. (1994). *Circuits to Control: Learning Engineering by Designing LEGO Robots*. Doctoral dissertation, Cambridge, MA: MIT Media Laboratory.

Minsky, M. (1987). *The Society of Mind*. New York: Simon & Schuster.

Mitchell, M. (1996). *An Introduction to Genetic Algorithms*. Cambridge, MA: MIT Press.

Pagels, H. (1988). *The Dreams of Reason*. New York: Simon & Schuster.

Papert, S. (1980). *Minstorms: Children, Computers, and Powerful Ideas*. New York: Basic Books.

Papert, S. (1993). *The Children's Machine*. New York: Basic Books.

Papert, S. (1996). An exploration in the space of mathematical educations. *International Journal of Computers for Mathematical Learning* 1(1): 95–123.

Resnick, M. (1991). Animal simulations with *Logo: Massive parallelism for the masses. In J. Meyer and S. Wilson (Eds), *From Animals to Animats* (pp. 534–539). Cambridge, MA: MIT Press.

Resnick, M. (1994). *Turtles, Termites, and Traffic Jams: Explorations in Massively Parallel Microworlds*. Cambridge, MA: MIT Press.

Resnick, M. (1996a). Beyond the centralized mindset. *Journal of the Learning Sciences* 5(1): 1–22.

Resnick, M. (1996b). New paradigms for computing, new paradigms for thinking. In A. diSessa, C. Hoyles, R. Noss and L. Edwards (Eds), *The Design of Computational Media to Support Exploratory Learning* (pp. 31–43). Springer-Verlag.

Resnick, M. (1997). Learning through computational modeling. *Computers in the Schools*.

Resnick, M., Martin, F., Sargent, R. and Silverman, B. (1996). Programmable bricks: Toys to think with. *IBM Systems Journal* 35(3): 443–452.

Resnick, M. and Rusk, N. (1996). The computer clubhouse: Preparing for life in a digital world. *IBM Systems Journal* 35(3&4): 431–439.

Roberts, N., Anderson, D., Deal, R., Garet, M. and Shaffer, W. (1983). *Introduction to Computer Simulation: A System Dynamics Modeling Approach*. Reading, MA: Addison-Wesley.

Senge, P. (1990). *The Fifth Discipline: The Art and Practice of The Learning Organization*. New York: Doubleday/Currency.

Stefik, M. (1996). *Internet Dreams: Archetypes, Myths, and Metaphors*. Cambridge, MA: MIT Press.

von Foerster, H., Mead, M. and Teuber, H., eds. (1949). *Cybernetics: Circular, Causal, and Feedback Mechanisms in Biological and Social Systems*. New York: Josiah Macy Jr. Foundation.

Waldrop, M. (1992). *Complexity: The Emerging Science at the Edge of Order and Chaos*. New York: Simon and Schuster.

Whitehead, S. (1994). Auto-FAQ: An Experiment in Population Leveraging. *Proceedings of the Second International WWW Conference* (pp. 25–38).

Wiener, N. (1948). *Cybernetics, or Control and Communication in the Animal and the Machine*. New York: John Wiley.

Wilensky, U. (1993). *Connected Mathematics: Building Concrete Relationships with Mathematical Knowledge*. Doctoral dissertation, Cambridge, MA: MIT Media Laboratory.

Wilensky, U. (1996). Modeling rugby: Kick first, generalize later? *International Journal of Computers for Mathematical Learning* 1(1).

MIT Media Laboratory
E-mail: mres@media.mit.edu
http://www.media.mit.edu/˜mres/