MASSACHUSETTS INSTITUTE OF TECHNOLOGY
MEDIA LABORATORY

EPISTEMOLOGY AND LEARNING GROUP

E&L Memo No. 10
November 1990

Children and Artificial Life

Mitchel Resnick
Fred Martin

Abstract

Artificial Life is a new field in which researchers study living systems by trying to build artificial versions of them. In this paper, we argue that ideas from Artificial Life research can and should be shared with children. We describe various computational tools (including LEGO/Logo and Electronic Bricks) that students can use to build artificial creatures. By building and programming artificial creatures (and discussing and thinking about how the creatures behave), children can explore some of the central ideas of Artificial Life—ideas like feedback, levels of organization, and emergence.

Children and Artificial Life

Mitchel Resnick and Fred Martin
MIT Media Laboratory
20 Ames Street
Cambridge, MA 02139
mres@media-lab.media.mit.edu
fredm@media-lab.media.mit.edu

Abstract

Artificial Life is a new field in which researchers study living systems by trying to build artificial versions of them. In this paper, we argue that ideas from Artificial Life research can and should be shared with children. We describe various computational tools (including LEGO/Logo and Electronic Bricks) that students can use to build artificial creatures. By building and programming artificial creatures (and discussing and thinking about how the creatures behave), children can explore some of the central ideas of Artificial Life—ideas like feedback, levels of organization, and emergence.

Introduction

Scientific ideas have their greatest force when they begin to seep out of the scientific community and into the general culture. This has happened with computer-science ideas during the past decade. As computers have proliferated throughout society, so have computational ideas and computational metaphors. Mental mistakes that were once seen as "Freudian slips" are now viewed as "information processing errors" [Turkle 1984]. Indeed, computational metaphors are influencing the way people think about things in nearly all disciplines, in nearly all parts of society.

Ideas from the emerging field of Artificial Life could have a similar influence in the next decade. Like biology, Artificial Life aims to explain how animals behave and how they evolve. But Artificial Life (or ALife, as it is sometimes called) tries to achieve this goal through a radically different approach. Rather than observing real animals, ALife researchers try to create artificial animals, typically in the form of robots and computer simulations. Some ALife researchers [Deneubourg et al. 1989; Travers 1989] are creating "artificial ants" on the computer screen, hoping to gain a better understanding of how real ants build nests and find food. Other ALife researchers [Reynolds 1987] are creating "artificial birds," hoping to gain a better understanding of how real birds move together in flocks. In short, ALife researchers think that the best way to understand Life is to build it. (For an overview of ALife research, see [Langton 1989].)

Artificial Life is not merely a field of research, but a set of powerful intellectual ideas—in particular, ideas about how complex systems (like living creatures) develop, interact, and behave. As ideas from Artificial Life infiltrate the culture, it is important to develop ways to share these ideas with children. Giving children access to ideas from the forefront of scientific research is, in general, a powerful pedagogic approach. When children sense that they are sharing in a new and dynamic enterprise, they are more likely to invest themselves in the process of learning. Thus, children can both benefit from and contribute to the dissemination of new scientific ideas.

In certain fields, such as high-energy physics, it is difficult to make cutting-edge ideas accessible and relevant to children. But that is certainly not the case with Artificial Life. There is something compelling and exciting about the idea of Artificial Life. The study of Artificial Life is full of objects that are, in Sherry Turkle's phrase, "betwixt and between"—objects that straddle the fuzzy boundary between the living and the inert [Turkle 1984]. Such objects are the source of both tension and fascination. And that makes them rich terrain for learning.

Through ALife projects, students can explore a set of important scientific ideas that are generally ignored in pre-college classrooms. In particular, students can explore ideas about *systems*. For example: How does the complex behavior of an ant colony emerge from the simple actions of individual ants? In working on ALife projects, students focus particularly on biological and ecological systems. But many of the ideas they learn (ideas like *feedback* and *levels of organization* and *emergence*) apply equally well to physical systems and social systems. Thus, ALife projects hold the promise of providing students with a new framework for thinking about a wide range of phenomena in the world.

In this paper, we discuss how (and why) to make Artificial Life ideas and activities accessible to children. We draw on several years of experiences in both classrooms and informal-education settings (see, for example, [Martin 1988] and [Resnick et al. 1988]). Some of the examples involve computer-based systems that are already available commercially (like Logo and LEGO/Logo). Others involve systems that are in the research stage at our laboratory at MIT. With these tools, children can explore how (artificial) creatures behave in different environments, and how they interact with other (artificial) creatures. In the course of such projects, children explore some of the central themes of Artificial Life research.

Logo and LEGO/Logo

One could argue that Artificial Life came to the classroom a long time ago, in the form of the Logo turtle [Papert 1980]. After all, isn't the Logo turtle a type of artificial creature? Well, yes and no. For the most part, students (and teachers) have tended to focus on the Logo turtle as a "drawing turtle": they use the turtle to draw geometric shapes, patterns, and pictures. But for the purposes of Artificial Life, it is more useful to think of the turtle in a different way—not as a drawing turtle, but as a "behavioral turtle."

Imagine, for example, a Logo program in which the turtle continually moves forward, but turns 160 degrees whenever it "sees" a red pixel under it. The turtle starts to have a "personality": one could say that the turtle is "scared" of the color red. In this program, the Logo primitive `colorunder` acts as a crude "sensor" for the turtle, providing the turtle with information about the "world" in which it lives.

In versions of Logo with multiple turtles, more complex examples are possible. Imagine a Logo program in which one turtle runs around randomly, while a second turtle "chases" the first (always setting its heading towards the first). The first turtle can be viewed as the "prey," the second turtle as the "predator." This simple microworld suggests many interesting experiments. If the prey is faster than the predator, will the predator ever catch the prey? If so, how long will it take?

This type of project begins to capture the spirit of Artificial Life. ALife projects typically involve a cycle of activities: build a new simulation, analyze what happens, develop a new theory, then modify the simulation to test the new theory. Every simulation leads to new questions and new simulations. In this way, ALife projects blur the traditional distinction between "theoretical" and "experimental" science. ALife projects blend theory and experiment, analysis and synthesis.

The integration of LEGO bricks with Logo programming (known as LEGO/Logo) opens new possibilities for Artificial Life activities in the classroom. With LEGO/Logo, children can write Logo programs to control machines that they construct out of LEGO building pieces (including not only the traditional LEGO bricks, but also LEGO gears and motors and sensors). Students have used LEGO/Logo to build (and program) all sorts of things: merry-go-rounds, conveyor belts, automatic doors, moving sculptures (see, for example, [Resnick et al. 1988]). Most LEGO/Logo projects in schools today seem to focus more on "machines" than "creatures." Students often build machines that they have seen in the world—or fantasy machines that they would like to see in the world. But there is great potential for using LEGO/Logo as a construction set for artificial creatures.

In building LEGO/Logo creatures, *sensors* are particularly important. Using sensors, a LEGO/Logo creature can sense the world around itself—then change its behavior depending on what it senses. For example, a LEGO/Logo creature with a touch sensor can be programmed to change direction whenever it bumps into an obstacle. The motion of the creature is not fixed: it depends on what obstacles are in the room. With sensors, LEGO/Logo creatures don't just *act in* the world, they *interact with* the world.

Electronic Bricks

Although LEGO/Logo is a rich construction environment, it does have some limitations. One problem is that LEGO/Logo machines must be connected with wires to a desktop computer. Wires are a practical nuisance, particularly when children use LEGO/Logo to create mobile "creatures." Wires get tangled with other objects in the environment, they get twisted in knots as the creature rotates, and they restrict the overall range of the creature. More important, wires are a conceptual nuisance. It is difficult to think of a LEGO/Logo machine as an autonomous creature as long as it is attached by umbilical cord to a computer.

The wiring problem could be eliminated by using an infrared link, as in television remote control. But we decided to make a more serious modification: we began to build the electronics *inside* the LEGO bricks [Martin 1988]. With these "Electronic Bricks," students no longer need to connect their LEGO creatures to a personal computer. Rather, they can build computational power directly into their LEGO machines and creatures. No more umbilical cord.

We have created about a dozen types of Electronic Bricks, falling into four categories. There are Action Bricks (motors, lights), Sensor Bricks (light sensors, sound sensors), Logic Bricks (and-gates, flip-flops, timers), and a Programmable Brick. The Programmable Brick is a fully-programmable computer with the power of an Apple IIe, built inside a LEGO brick the size of a deck of cards. Students can create different behaviors for their creatures by wiring Electronic Bricks together in different ways, much in the spirit of Valentino Braitenberg's *Vehicles* [Braitenberg 1984].

Consider, for example, a simple creature with a Light-Sensor Brick (pointing upward) and a Motor Brick. The output from the Light-Sensor Brick is connected to the Speed input of the Motor Brick. When the Light-Sensor Brick detects light above a certain threshold, it sends a signal to the Motor Brick, and the creature moves. When creature is in the dark, it

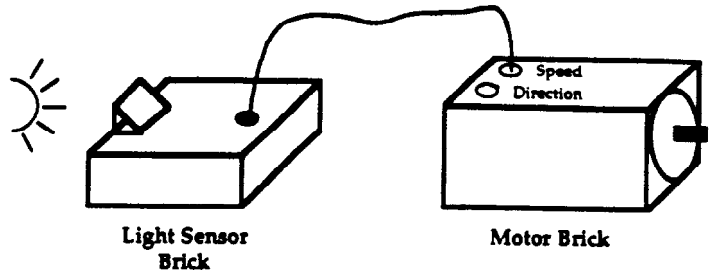

Figure 1: Wary

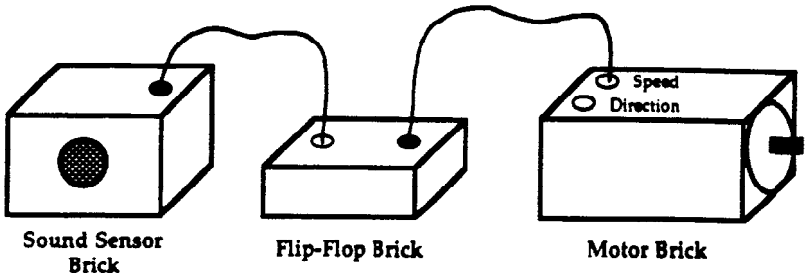

Figure 2: Obedient

stops. We dubbed this creature *Wary* (Figure 1). If you then replace the Light-Sensor with a Sound-Sensor and a Flip-Flop Brick, the creature might be called *Obedient* (Figure 2). Every sharp sound "flips" the state of the Flip-Flop Brick. So the creature stops moving when you clap your hands, then starts again when you clap again.

Students have worked with Electronic Bricks in several ways. Some have observed pre-built creatures, trying to figure out how the creatures work—much as ethologists observe animals [Granott Farber 1990]. Others have built creatures of their own, trying to create new behaviors from simple behavioral building blocks [Martin 1988; Bourgoin 1990].

Thinking About Systems

Viewing the world in terms of *systems* is a fundamental, organizing idea throughout the natural sciences, social sciences, and engineering. Yet pre-college education rarely addresses systems-oriented concepts. Part of the problem is that students (and teachers) have not had appropriate tools to "mess about" with systems. New tools like LEGO/Logo and Electronic Bricks can change that. Through our observations of children building artificial creatures with LEGO/Logo and Electronic Bricks, we have found that they explore a wide range of systems ideas—including interaction, feedback, levels, and emergence.

Interaction. Systems are defined not only by their components or elements, but by the ways in which those components interact with one another. In LEGO/Logo and Electronic Brick projects, interaction occurs at several levels: interactions among the components that make up a creature, interactions between the creature and the environment, and interactions between the creature and other creatures.[1]

In even the simplest activities, children begin to gain a sense of the interaction between creatures and the environment. Imagine a LEGO creature with a built-in motor to power its motion. When a child connects an Electronic Brick light sensor to the motor, the creature immediately begins reacting to changes in light. Moreover, children quickly discover that a creature's behavior can be changed not only by modifying the creature itself, but by changing the environment in which the creature lives. For example, children can add more lights to the environment, or make the lights flash. In many cases, the creature, placed in a new environment, reacts in unexpected ways, revealing new aspects of the creature's behavior—and driving home the importance of thinking about creatures in the context of their environments.

Feedback. Real creatures are not pre-programmed with an exact set of actions. Rather, they constantly modify their behaviors based on interactions with the environment. Many engineered systems use the same strategy. A rocket, for instance, continually makes mid-course adjustments to stay on its desired course.

Children can build LEGO creatures that work on the same *feedback* principle. One popular project is to build a creature that moves towards the light. A feedback-based approach works as follows: if the creature senses that the light is brighter to the left, turn a bit to the right; if it senses that the light is brighter to the right, turn a bit to the left. With this simple strategy, the creature will (with a bit of weaving) move toward a bright light.

Of course, there are limitations to this strategy. For example, the creature will not necessarily move towards the brightest light in the environment. Imagine that there are two light bulbs near the creature: one bright bulb and one dim bulb. If the creature happens to "see" the dim bulb first, it might move toward that bulb, ignoring the brighter bulb. This problem is often referrred to as the "local maximum" problem. Such problems are typically not discussed until college engineering courses. But LEGO creatures provide a meaningful context for children to start exploring such problems as early as elementary school.

---

[1] Of course, interactions between the creature and other creatures can be viewed as a special case of interactions between the creature and the environment.

Levels. As students play with artificial creatures, we are particularly interested in how the students think about the creatures. Do they think of the LEGO creatures as machines, or as animals? In fact, we have found that students (and adults) regard the creatures in many different ways. Sometimes students view their creatures on a *mechanistic* level, examining how one LEGO piece makes another move. At other times, they might shift to the *information* level, exploring how information flows from one Electronic Brick to another. At still other times, students view the creatures on a *psychological* level, attributing intentionality or personality to the creatures. One creature "wants" to get to the light. Another creature "likes" the dark. A third is "scared" of loud noises.

Sometimes, students will shift rapidly between levels of description. Consider, for example, the comments of Sara, a fifth-grader [Martin 1988]. Sara was considering whether her creature would sound a signal when its touch sensor was pushed:

> *"It depends on whether the machine wants to tell... if* we *want the machine to tell us... if we* tell *the machine to tell us."*

Within a span of ten seconds, Sara described the situation in three different ways. First she viewed the machine on a psychological level, focusing on what the machine "wants." Then she shifted intentionality to the programmer, and viewed the programmer on a psychological level. Finally, she shifted to a mechanistic explanation, in which the programmer explicitly told the machine what to do.

Which is the correct level? That is a natural, but misleading, question. Complex systems can be meaningfully described at many different levels. Which level is "best" depends on the context: on what you already understand and on what you hope to learn. In certain situations, for certain questions, the mechanistic level is the best. In other situations, for other questions, the psychological level is best. By playing with artificial creatures, students can learn to shift between levels, learning which levels are best for which situations.

Emergence. In many systems, certain behaviors or patterns *emerge* from interactions among system components. For example, the patterns in a bird flock emerge from the interactions among individual birds. Similarly, the behavior of a market economy emerges from the interactions among millions of economic actors. These behaviors or patterns are not explicitly programmed or represented in the system components. Rather, these emergent behaviors represent a higher "level" in the system. This higher level must

described with different parameters than those used to describe the lower-level components.

In many animal systems, there are two types of emergence. First, the behavior of each individual creature emerges from interactions among the "agents" that make up the creature's mind [Minsky 1988]. At the same time, the behavior of the entire animal colony or society emerges from the interactions among the individual creatures. In short: the colony level emerges from the creature level, which in turn emerges from the agent level.

With LEGO/Logo and Electronic Brick creatures, children can begin to observe and experiment with emergent behaviors. Consider a simple LEGO creature with a light sensor pointing upward. Imagine that the creature is programmed with two rules: (1) move forward when you detect light, (2) move backward when you are in the dark. When this creature is released in the environment, it exhibits a type of emergent behavior: it seeks out the edge of a shadow, then happily oscillates around the shadow edge. The creature can be viewed as an "Edge-Finding Creature." This edge-finding capability is not explicitly represented in the creature's two rules. Rather, it emerges from those rules.

Here is another example. When we started to develop LEGO/Logo, one of our first projects was to program a LEGO "turtle" follow a line on the floor. The basic strategy was to make the turtle weave back and forth across the line, making a little forward progress on each swing. First, the turtle veered ahead and to the right, until it lost sight of the line. Then it veered ahead and to the left, until it again lost sight of the line. Then it started back to the right, and so on. This behavior can be represented by two simple rules: (1) if you are veering to the left and you lose sight of the line, begin to veer right, (2) if you are veering to the right and you lose sight of the line, begin to veer left.

We tried the program, and the turtle followed the line perfectly. But as the turtle approached the end of the line, we realized that we hadn't "programmed in" any rules for what to do at the end of the line. We didn't know what the turtle would do. We were pleased with the behavior that emerged: the turtle turned all the way around and started heading back down the line in the other direction. This "end-of-line" behavior was not explicitly programmed into the turtle. Rather it emerged from the interactions between the turtle's rules and the unfamiliar environment at the end of the line.

Of course, these examples represent rather simple cases of emergence. But emergence is an ill-defined and poorly understood concept. Simple cases are probably the best place to start—perhaps not just for children and novices, but for scientists as well.

Thinking About Themselves

Playing with Artificial Life systems can influence the way children think about systems of all kinds: physical systems, political systems, economic systems. But perhaps most important is how it influences the way children think about *themselves* as systems.

While building and programming artificial creatures, children naturally ask questions about themselves. In some cases, Artificial Life activities encourage children to reexamine things that once seemed obvious. As children try to make their creatures move, they reexamine how they themselves walk. And they reflect on the advantages and disadvantages of walking versus rolling, of legs versus wheels. (These activities exemplify of what Papert [1980] calls "body syntonic" learning.)

As children add light sensors to their creatures, they wonder: what would it be like to have only one eye? Or three? As children work with sensors, they also gain a deeper understanding of how their own senses work. At first, many children think that LEGO sensors will work "by themselves." They put a touch sensor on the front of a creature, for example, and they expect the creature to automatically change direction when it bumps into something. They don't realize that they must *program* that behavior into the creature (by writing a Logo program, or making the appropriate connections among Electronic Bricks).

This mistake is not all that surprising. Johnson and Wellman [1982] showed that fifth-grade children generally view the brain as the organ of mental life. Thus, children see the brain as essential for thinking, dreaming, and remembering, but they deem it irrelevant to walking, sneezing, or sensing. After building and programming artificial creatures, children generally have a much different view. They recognize the important role that the brain plays in processing sensory inputs and orchestrating actions by the body.

Through our work with Artificial Life and children, we hope to gain a sharper sense of how children think about their own minds. As children build and program artificial creatures, how will their images of the mind change? Karen Thimmesch, a teacher at Galtier Magnet School in St. Paul, Minnesota, has begun to investigate this question. She asked her fourth-grade students to draw pictures showing how information is transmitted and processed inside their own heads.

The children's pictures indicate a wide range of metaphors—in some ways mirroring the metaphors of mind that scientific researchers have used through history. Some students used transportation metaphors, showing cars and bicycles carrying information around the brain (Figure 3). Others used communications metaphors, showing telephone wires (Figure 4) or walkie-talkies throughout the brain (Figure 5). Some children relied on

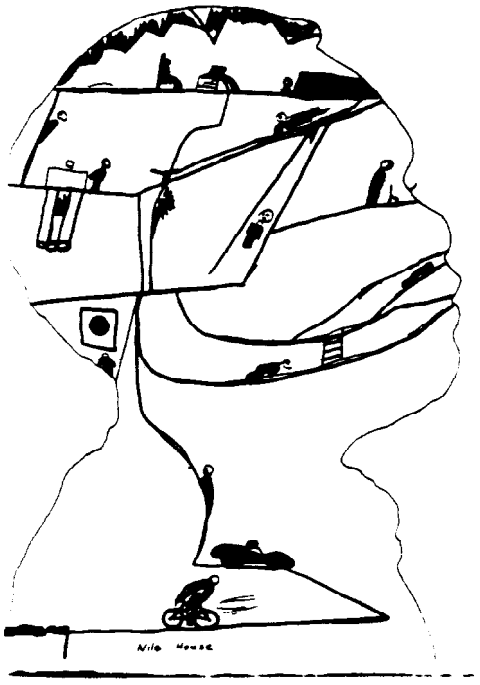

Figure 3

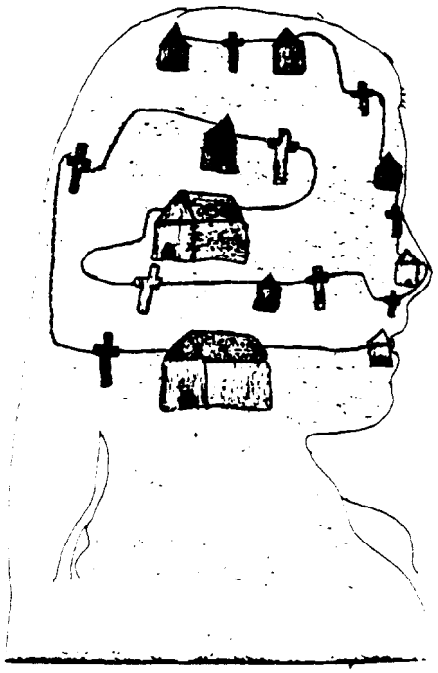

Figure 4

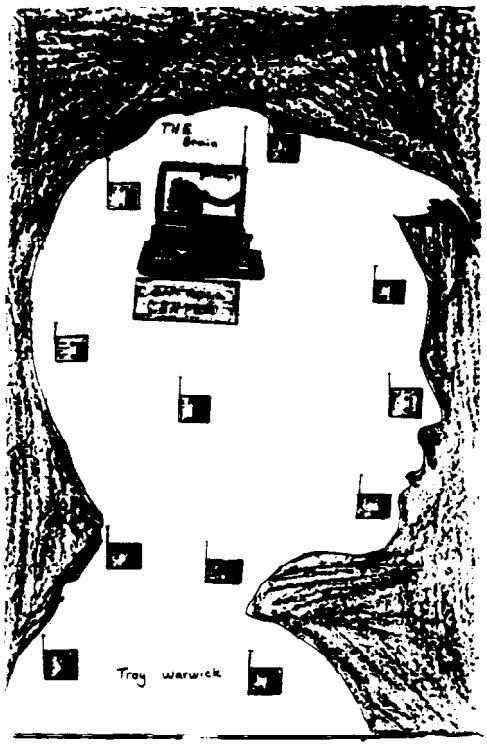

Figure 5

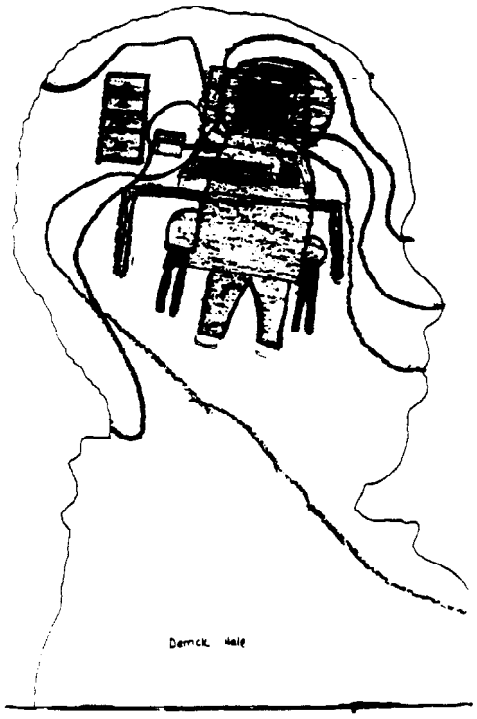

Figure 6

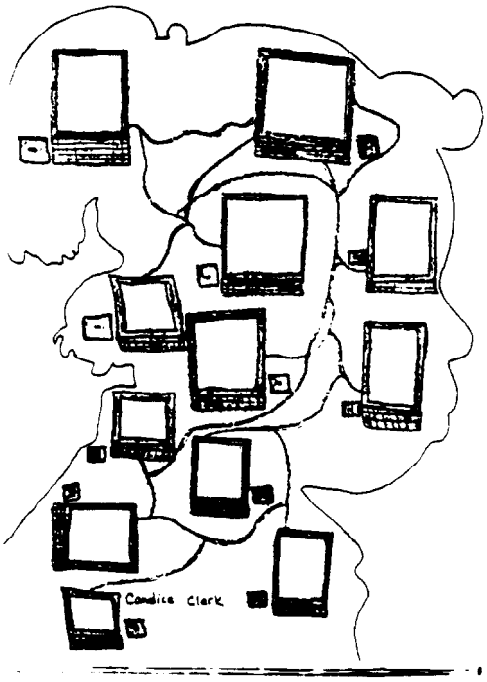

Figure 7

Figure 8

computational metaphors, showing a little person at a computer terminal inside the head (Figure 6). Of course, scientists these days prefer distributed (not centralized) computational metaphors of the mind, with "societies" of agents interacting with one another [Minsky 1988]. Some children had similar ideas (Figure 7). The most unusual picture, perhaps, was an ant colony in the brain (Figure 8). This might seem strange at first, but many ALife researchers and cognitive scientists view ant colonies as one of the best metaphors for the brain/mind [e.g. Hofstadter 1979]. In this view, each individual (ant or neuron) is capable of very little, but collectively they are far more than the sum of the parts.

How will children's images of their minds evolve as they spend more time with artificial creatures? So far, we do not know; we are just starting to study this question. But we expect that Artificial Life activities will influence children's ideas in a fundamental way. No one expects children to come up with answers to ages-old questions about the mind, but it is important for them to think about these questions, to play with new ideas and try them out. At the least, LEGO/Logo and Electronic Bricks will give children a better set of tools for thinking about an important set of questions.

Is It Too Soon?

Some people might wonder whether it is appropriate (or feasible) to introduce ideas from Artificial Life at a pre-college level. After all, scientists themselves do not have a very good understanding of the mechanisms (or even the meaning) of Artificial Life. Shouldn't we wait until scientists figure it out before introducing it in schools?

Our answer is an emphatic *no*. Although ideas like emergence are still poorly understood, certain core ideas of Artificial Life are clear. And new computational tools (like LEGO/Logo and Electronic Bricks) make it possible for pre-college students to explore these ideas, at least qualitatively, if not formally. By building and programming artificial creatures, and by discussing and thinking about the results, children begin to view the world through a systems-oriented lens. The fact that scientists are still struggling to construct the intellectual framework of Artificial Life can, in fact, be turned to advantage. Too often, school science insulates students from the excitement that drives real scientific research. By exploring Artificial Life, students can share in the excitement.

Scientists, too, can learn a great deal from the introduction of Artificial Life into the classroom. In many cases, the best way to clarify a new (and still-fuzzy) idea is to try to make it accessible to children. In developing explanations and tools for children, researchers are forced to confront the core issues in a domain. They must strip away levels of detail and formalism, and focus on the central concepts of the domain. As a result of this process, not only do students get a better sense of the domain, but so do the researchers.

Acknowledgments

Many of the ideas in this paper are based on collaboration with Seymour Papert, Stephen Ocko, Edith Ackermann, and Brian Silverman. In addition, Mike Eisenberg and Franklyn Turbak provided helpful comments on a draft of the paper. We are particularly grateful to Karen Thimmesch of Galtier Magnet School of the St. Paul Public Schools for sharing her experiences and students' drawings with us.

Our research on LEGO/Logo is supported by grants from the LEGO Group and the National Science Foundation (Grants 851031-0195, MDR-8751190, and TPE-8850449). Some of the ideas in this paper were previously presented in a paper titled "LEGO, Logo, and Life" [Resnick 1989].

References

Bourgoin, M. (1990). "Children Using LEGO Robots to Explore Dynamics," in *Constructionist Learning*, MIT Media Laboratory. Cambridge, MA.

Braitenberg, V. (1984). *Vehicles* (Cambridge: MIT Press).

Deneubourg, J.L., S. Goss, N. Franks, and J.M. Pasteels (1989). "The Blind Leading the Blind: Modeling Chemically Mediated Army Ant Raid Patterns," *Journal of Insect Behavior* (vol. 2, no. 5).

Granott Farber, N. (1990). "Puzzled Minds and Weird Creatures: Spontaneous Inquiry and Phases in Knowledge Construction," in *Constructionist Learning*, MIT Media Laboratory. Cambridge, MA.

Hofstadter, D. (1979). *Godel, Escher, Bach: An Eternal Golden Braid* (New York: Basic Books).

Johnson, C.N., and H.M. Wellman (1982). "Children's developing conceptions of the mind and the brain," *Child Development* (vol. 53), pp. 222-234.

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

Martin, F. (1988). "Children, Cybernetics, and Programmable Turtles," unpublished Master's Thesis, MIT Media Laboratory.

Minsky, M. (1988). *The Society of Mind* (New York: Simon & Schuster).

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

Resnick, M., S. Ocko, and S. Papert (1988). "LEGO, Logo, and Design," *Children's Environments Quarterly* (vol. 5, no. 4).

Resnick, M. (1989). "LEGO, Logo, and Life," in *Artificial Life,* edited by C. Langton (Redwood City, CA: Addison-Wesley).

Reynolds, C. (1987). "Flocks, Herds, and Schools: A Distributed Behavioral Model," *Computer Graphics* (vol. 21, no. 4).

Travers, M. (1989). "Animal Construction Kits," in *Artificial Life,* edited by C. Langton (Redwood City, CA: Addison-Wesley).

Turkle, S. (1984). *The Second Self: Computers and the Human Spirit* (New York: Simon and Schuster).