

SFI@30



# INFORMATION, ADAPTATION, AND EVOLUTION **IN SILICO**

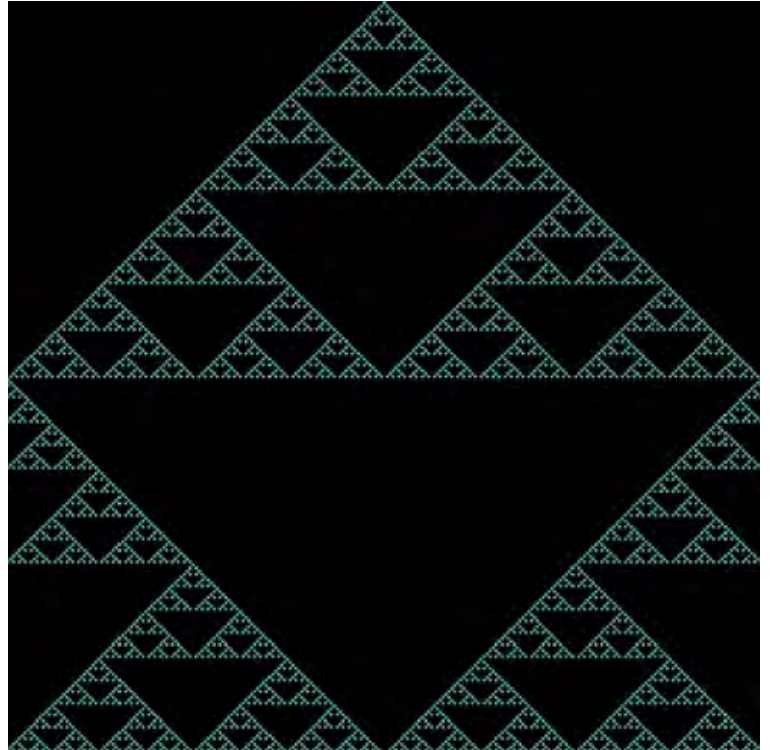
BY MELANIE MITCHELL

Professor, Computer Science,  
Portland State University;  
External Professor, Santa Fe Institute

In 1984 the nascent Santa Fe Institute sponsored two workshops on “Emerging Syntheses in Science,” at which the Institute’s founders brainstormed their plans for the future. At the time I was a beginning graduate student in computer science and had never heard of SFI, but reading the workshop proceedings a few years later, I was very excited by the Institute’s goal to “pursue research on a large number of highly complex and interactive systems which can be properly studied only in an interdisciplinary environment.”

The founders planned to define particular themes or programs that would benefit from the kind of intensive cross-disciplinary interaction offered by the new institute. SFI’s first official program, formed in 1987, was Economics. Before long, several influential players in the field took note of SFI’s novel interdisciplinary approach to economics, and the program grew quickly, in fact threatening to take over the fledgling organization.

Founder and first SFI President George Cowan wanted to make sure economics did not come to dominate. Cowan wrote: “We had to start somewhere, but we also had to make sure from the beginning that economics didn’t become the one



Left: Melanie Mitchell. Above: SFI researchers explored the use of genetic algorithms to automatically evolve computational structures, including cellular automata – a collection of “colored” cells on a grid that iterates through a number of time steps according to simple rules and often resulting in surprising patterns.

interest of the institute...I pushed hard to support at least one other program that would be equal in size to the economics program. We needed to broaden our academic agenda, and spread our bets.” Cowan’s push was to start a program in “adaptive computation.”


## What is adaptive computation?

Adaptive computation is a broad area that covers three major threads of computing in the context of complex adaptive systems:

**Agent-based simulation:** In science, computers have traditionally been used to perform “numerical simulation” – using mathematical equations to simulate behavior in physical systems, such as turbulence in fluids and missile trajectories in space. But in most systems of interest to SFI – economics and other social systems being prime examples – it’s hard, if not impossible, to find a set of equations that describes these multi-agent, open-ended, evolving systems. Instead, the study of complex systems has pioneered computational agent-based simulation, in which the behaviors of individual system components (“agents”) and their interactions are explicitly simulated, rather than simulating the behavior of equations describing the system as a whole. Agent-based simulation has become a mainstay of research in complex systems ranging from gene networks to financial markets.

**Nature-inspired computation:** Computers themselves can be made more adaptive and lifelike by adopting ideas from natural adaptive systems. This idea goes back to the beginnings of the computer age, when trailblazers such as Alan Turing and John von Neumann were thinking hard about the connections between programmable computers and living systems. Since that time, theories about the brain have inspired computational neural networks that learn on their own; the theory of Darwinian evolution has inspired genetic algorithms in which programs evolve via computational “natural selection” rather than being engineered by humans; and immunology has inspired software-based artificial immune systems for computers to fight computer viruses and other so-called malware.

Other inspirations for new computing methods have come from natural adaptive systems as diverse as ant colonies, slime molds, economic markets, and social networks.



Patterns created via adaptive computational models often resemble those found in nature.

**Computation as a framework for understanding nature:** The idea of computation goes beyond what we traditionally call “computers.” As former SFI scientist Chris Langton eloquently put it: “The proper domain of computer science is information processing writ large across all of nature.” A key property of complex adaptive systems is their ability to process information – to compute – in order to adapt and thrive in an environment. Thus, the concepts and theories of computer science can themselves be adapted to provide a scientific language for understanding information processing in the natural world.

### SFI’s Adaptive Computation Program

SFI’s Adaptive Computation (AC) program originated from the work of John Holland at the University of Michigan in the 1960s and 70s. Holland developed cross-disciplinary theories of adaptation, as well as computer models of evolution and learning. Genetic algorithms (methods for “evolving” programs and other computer structures) were Holland’s invention, as were classifier systems (evolving rule-based learning systems) and the “Echo” model (an agent-based model of an evolving ecosystem).

Holland was originally recruited to be SFI’s first resident faculty member and to lead the AC program. However, after much thought, he decided to stay at Michigan, albeit with frequent visits to SFI (which continued for 30 years). John Miller, who was Holland’s collaborator and SFI’s first postdoc, took over direction of the program and organized its 1992 founding workshop, which brought together many of the leading thinkers in areas related to adaptive computation. Miller, however, was soon leaving for a faculty position at Carnegie Mellon University, so another director had to be found.

I had come to the University of Michigan in 1984 to work on artificial intelligence (AI) with Douglas Hofstadter. During my first year there I took John Holland’s course, “Adaptation in Natural and Artificial Systems.” I was enchanted by the beautiful theory developed by Holland and

LEFT: NICK ZANTOP; CREATIVE COMMONS/SHAREALIKE 3.0



amazed by the abilities of genetic algorithms to evolve sophisticated programs and designs. My focus was still on my AI work with Hofstadter, but I made time to work with Holland and his students (particularly Stephanie Forrest) on genetic algorithms. Holland became my co-advisor. After graduating with a Ph.D. in 1990, I was offered a postdoctoral fellowship in the Michigan Society of Fellows, to work primarily with Holland on genetic algorithms.


In early 1992, Holland asked me if I might want to come to SFI for the next academic year to direct the AC program. I jumped at the chance. For a young postdoc, it was the opportunity of a lifetime, and I somehow managed to stretch a single year as a visitor into six years as a resident professor at SFI. Over the six years of its existence, the AC program's resident and visiting researchers made significant impacts in all three of the threads described above.

Here I summarize some prominent examples of results from this program:

**Computer models of adaptive systems:** SFI's AC program promoted sustained collaborations between computer scientists and evolutionary biologists in building models of biological evolution at different scales. In addition, scientists in the program built extensively on Holland's work on genetic algorithms and classifier systems; his thoughts about energy flow in ecologies were formalized in the "Echo" model (developed by Holland, Terry Jones, and Stephanie Forrest).

Stephanie Forrest and Alan Perelson pioneered work on agent-based modeling of the adaptive immune system. AC program researchers also collaborated with economists, creating agent-based models of economic markets in which the agents (like real-world decision makers in an economy) had limited rationality and knowledge, but were able to adapt and learn. This notion of adaptive economic agents has become a central aspect of what is now called complexity economics.

Finally, the AC program included several visitors who worked on agent-based models of human and animal learning and on the



CREATED BY MELANIE MITCHELL USING NETLOGO

perform “emergent computation.” Invented in the 1940s by computing pioneers John von Neumann and Stanislaw Ulam, a cellular automaton is a grid of simple (simulated) “cells.” In our formulation, each cell can be in one of two “states” – black or white – and is connected to a small number of neighboring cells. At each time step, each cell updates its state (either staying the same color or changing color) depending on the states of its neighbors.

In a given cellular automaton, each cell obeys the same “rule” in updating its state, but different cellular automata can have different rules. Our group’s work was on using genetic algorithms to evolve cellular automata rules that would enable a cellular automaton to act as a specialized sort of computer. We showed that this was not only possible, but that the genetic

Modeling the dynamics of “swarm intelligence” using principles from insect colonies was among the early research of SFI’s Adaptive Computation program.

interaction between learning and evolution. These modeling efforts were variously based on genetic algorithms, classifier systems, neural networks, and dynamical systems.

**Nature-inspired computation:** A few examples of the program’s wide-ranging research included development of computational “immune systems” that used immunological principles to recognize computer viruses and network attacks and set up a defense; the use of genetic algorithms to automatically evolve computational structures such as computer programs, neural networks, and cellular automata; the implementation of “swarm intelligence” using principles from insect colonies; and the simulation of host-parasite co-evolution in order to improve the performance of learning programs (“hosts”) on the data they learn from (“parasites”).

I’ll expand a bit on one of these examples. SFI’s Jim Crutchfield and I led the “Evolving Cellular Automata” group. Our research focused on using genetic algorithms to evolve cellular automata to

algorithm was able, in many cases, to discover more sophisticated “solutions” than had been created by humans working on the same problem.

**Computation as a framework for understanding nature:** The SFI AC program included several computer scientists (including myself) who were intensely interested in what their field could offer to the study of natural systems beyond the development of modeling algorithms and tools. A major focus was to understand how the complex dynamics we observe in many natural systems give rise to information processing.

As one example, we found that cellular automata produced by our genetic algorithm exhibited dynamical patterns that roughly resembled physical particles moving through a medium, colliding, and producing new particles. We were able to show that these cellular-automata-based “particles” were the locus of information storage: their movements through the grid affected information transfer, and their collisions were the sites of information processing.

In this way, we were able to make sense of the emergent computational properties of a cellular automaton in terms of its underlying dynamics. This is a novel approach to understanding computation that has had impacts ranging from new ideas for the design of nanoscale computers to the understanding of how plants process information to regulate the balance between their water intake and carbon dioxide output.

### Impact of SFI's AC program

The examples of research I've described above, and many others, put SFI on the map as a key, well-respected player in adaptive computation. In addition to the many publications written and invited lectures given by program members, there was one particularly impressive honor.

In 1997, five years after the AC program's founding, *Business Week* polled a large group of researchers for the answer to this question: "If you were 35 and had just won the first Nobel Prize for Information Technology, triggering invitations to the lab of your choice, which one would you pick?" Once the votes were counted, *Business Week* published the top-ten list, which included Stanford, Berkeley, MIT, Bell Labs, Microsoft, and similar institutions. Of these, tiny SFI was tied for 5th. And when the question was restricted to labs focusing on biologically inspired computation, SFI moved up to first place.

SFI's small size and limited resources means that for it to prosper intellectually, it must keep bringing in new ideas and "new blood." Thus, the Institute has no tenure and no permanent faculty positions. It also has no permanent programs. In 1999, with my SFI faculty term coming to an end, I left the Institute for an academic job in Oregon. The Adaptive Computation program also came to an end (as did the Economics program). More generally, the whole idea of official, broad "programs" at the Institute was restructured into the notion of ever-changing interdisciplinary research themes.

However, the spirit of adaptive computation remains strong at SFI, and many people incorporate

AC into their research, on topics ranging from agent-based models of ancient state formation to genetic algorithms for automatically finding software bugs. Furthermore, the ideas of adaptive computation have spread far and wide into universities' computer science curricula, and into many other departments.

Indeed, our program's results impressed not only computer scientists, but people in many areas of science. Perhaps my favorite example of this was the well-known (and famously skeptical) evolutionary theorist Richard Lewontin, who, after hearing about SFI's Adaptive Computation research, announced, "I don't believe in adaptation. But I sure as heck believe in computation!"

Whether it is through adaptation or computation, or both, I'm proud to have been part of the birth and development of this foundational program for understanding complex adaptive systems. ◀

*Melanie Mitchell is a Professor of Computer Science at Portland State University and an External Professor and member of the Science Board at the Santa Fe Institute. She has held faculty positions at the University of Michigan, Los Alamos National Laboratory, and the Oregon Graduate Institute School of Science and Engineering. She is author or editor of five books and more than 70 scholarly papers in artificial intelligence, cognitive science, and complex systems. Her most recent book, Complexity: A Guided Tour (Oxford University Press, 2009) won the 2010 Phi Beta Kappa Science Book Award and was among Amazon.com's ten best science books of 2009.*



MINESH BACCRANIA PHOTOGRAPHY