Ramps-AntiRamps And the Red Queen

An Early Genetic Algorithm

Danny Hillis, 1991, 'Co-evolving Parasites Improve Simulated Evolution as an Optimization Procedure': in C. Langton et al. (eds) Artificial Life II.

From Charles Ofria
http://www.krl.caltech.edu/~charles/stories/alife.html

Ramps is a genetic algorithm evolving to reach a fitness peak at solving a mathematical problem - the ability to sort a random number list. Fitness is measured by the shortest number of steps evolved to solve the various problems present in the environment.

Antiramps is a genetic algorithm evolving to reach a fitness peak at creating test cases the Ramps can not solve well with the strategies evolved to date. That is, the most fit Antiramps are those which resist being sorted easily or well.

The excerpts below come from Steven Levy's book "Artificial Life [1]." It reveals how Hillis came up with the ramps/antiramps idea, and relates it to related ideas in computer programming, and biology.

Synopsis

An interesting addition to the Genetic Algorithms was discovered by Danny Hillis which he called the Red Queen Behavior. He realized that if you want to evolve one algorithm, it can often be done much faster if a competitive algorithm was run against it (called a parasite by him). The example he used was Ramps which was a program of his design which created algorithms to sort numbers. Initially (before he added in parasites) he was able to develop algorithms (in about a day of running) to sort numbers that would only require 65 exchanges to complete. He then created "anti-ramps" which had as their fitness function the number of Ramps that they were able to stump with the data they created. So effectively, the anti-ramps were causing the Ramps to work much harder in order to create a proper algorithm. In the end, the anti-ramps enabled Hillis to find an algorithm that only required 61 exchanges (which took the same period of time to run.) The best algorithm a human has been able to find requires 60 changes, so 61 is a significant achievement. When Hillis saw that his Ramps weren't able to beat the human best, however, he gave up on the Ramp project (though still does use the principles behind the Red Queen Behavior since they obviously were proven to work.)

As Hillis pondered the notion, he recognized that the most complex ensembles were living systems life-forms, including human ones. When one asked a biologist how we and our living cousins attained this level of complexity, the short answer was the single word, "evolution." And indeed, evolution was something based on simple rules that yielded wondrously complicated results. So Hillis decided to look closer at evolution to better understand emergence. As he read his Darwin, his Fisher, his Haldane, and his Gould, however, he realized that the theory of evolution was far from complete. There were gaps in our knowledge and spirited disagreement about some components of general evolutionary theory.

---

[1] Steven Levy, 1993, Artificial Life: A Report from the Frontier Where Computers Meet Biology: Vintage Books; Reprint edition (August 1993) ISBN: 0679743898

The Problem

Hillis was drawn to one particular controversy concerning the nature of evolutionary progress. One biological theorem, postulated by R. A. Fisher, conjectured that evolution proceeded by steady improvements in fitness. Hillis saw this as equivalent to the so-called hill climbing technique used by certain computer optimization procedures, such as learning in neural network simulations. Because each generation was supposedly slightly fitter than the previous one, a graph illustrating this progress would show a line angling upward, as though the fitness of the species were engaged in scaling a peak.

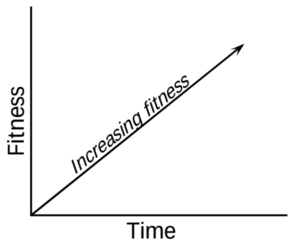

R.A. Fisher's Fitness Concept

Biologists studying the problem had created a more complicated, multidimensional map of the way that a species might evolve. This "adaptive landscape," first postulated by Sewall Wright, represented the space of all possible genetic combinations. It was filled with bumps, peaks, valleys, and spikes. The gene pool of an entire population resided at a single area on this landscape. The higher the ground, the fitter the population would be if it found its way there. When the terrain was fairly level, a population theoretically engaged in a "random walk," with the effects of crossover and mutation moving its genetic composition to different places, until it found an ascending plane. From that point, the more fit individuals within the population would push fitness higher, and the rest of the population would follow.

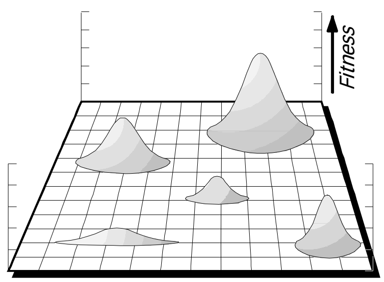

Sewall Wright's Fitness Landscape

But, if hill climbing was indeed the method nature used to achieve higher fitness, the discovery of the highest ground could not always be assumed. Once the population scaled a medium-size peak, it tended to get stuck. This was due to the built-in reluctance of a population to decrease its fitness, which would be necessary in order to search the landscape for an even higher peak. The population would remain fat and happy on its hill but miss out on the mountains that lay somewhere else on the chart. The population was then "stuck on a local maximum" with no incentive to make the giant evolutionary leaps that push life toward more complexity.

Hillis suspected that something other than hill climbing must have been required for the rich properties of life to emerge. "If evolution is a hill-climbing technique, why doesn't it seem to have problems that we know hill-climbing techniques suffer from?" he asks, referring to the local maxima problem. "Why is evolution so much more powerful than any other hill-climbing technique? Why is it able to evolve much more complicated things?"

Advantages of Alife for Solving Problems

In 1986 Hillis began seeking answers to these questions by simulating evolution on his multimillion-dollar machine. Combining his sixty-four thousand parallel processors, an experimental chemist's careful methodology, and the observational acumen of a biologist, Hillis had the power to generate hitherto-unsuspected insights from a sophisticated form of artificial evolution.

Physicist Niels Bohr once observed that "we should doubtless kill an animal if we tried to carry on the investigation of its organs so far that we could describe the role played by single atoms in vital functions.... The minimum freedom which we must allow organisms in this respect is just large enough to permit it, so to say, to hide its ultimate secrets from us." This limitation did not exist with the artificial organisms that Hillis studied. Ultimate secrets could now be exposed. "I can run a population for 100,000 generations," Hillis explains. "And then I can look at the same thing geologists look at, the fossil record. Then I can look at the actual genes themselves not just the phenotypes of the individuals but the actual genetic material. And I don't have to look at them with just my own eyes, I've got a big parallel computer to look at them. I can see things that are going on there that biologists don't have the data to see."

The Evolutionary Landscape of Ramps

Hillis worked with artificial organisms strings of numbers that represented genes, which in turn expressed themselves as phenotypes by performing computational tasks. It was very much in the spirit of John Holland's genetic algorithm. At first he set relatively simple problems for his organisms. He would require them to arrange their bits in a certain order, rewarding the ones that came closest to the desired arrangement after each generation. Fitness would be determined, of course, by criteria set by Hillis on what was the perfect order, and his clever use of that complication prevented the system from getting stuck in local maxima. It was easy to demonstrate how. Begin with a string of twelve numbers. Hillis might decide to regard a perfect sequence as one in which each number has a higher number to its right. The perfect sequence, then, would begin with numeral 1 and proceed, in order, to 12:

1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12.

In evolving this order, hill-climbing techniques would be effective only to a point because the problem might not be solved purely by increments. For example, consider an arrangement that began with 2 and ascended by even numbers up to 12; then, at the seventh digit dropped to 1 and counted by odd numbers to 11:

2, 4, 6, 8, 10, 12, 1, 3, 5, 7, 9, 11.

Only one step from perfection, yet obviously flawed, this organism would find it difficult to evolve further. It was stuck in a local maximum. Because of the ramplike slopes this scheme would draw on a landscape mapwhere hill climbing would lead to a precipice Hillis called his organisms "Ramps." He posed increasingly difficult problems to them so he could learn about evolution from their behavior.

Sex as an Evolutionary Strategy

When biologist Charles Taylor saw the system, he proposed that Hillis try experiments using both sexual reproduction and asexual reproduction. Hillis could then address an enduring problem in biology: why was there sex? Some theorists believed that it was only by strange happenstance -a frozen accident that organisms mated, because the short-term advantage would be for an individual to pass on as many of its genes as possible. Sex diluted the number of genes one passed to offspring, and common sense indicated that the fittest organisms would be better off reproducing asexually. In the long run, sexual reproduction did strengthen fitness in the population, but evolution, which proceeded generation by generation, had no consciousness and worked from no billion-year blueprint. Evolution proceeded like a cellular automaton experiment in that decisions were local, although consequences resonated globally. So it was difficult to see how even sound long-term behaviors could overrule decisions that made sense in the short run.

Hillis implemented sex in his system, effectively giving organisms the choice whether to reproduce sexually. He accomplished this by introducing a gene that controlled the percentage of the time that the organism reproduced sexually. The percentages that helped increase fitness in the phenotypes of the organisms would, of course, dominate. In the first experiments Hillis ran with this parameter the organisms found it in their interest to practice asexual behavior. Later, he posed some different tasks, and in those cases the system took the far-sighted decision and sought mates. Hillis could not isolate what, if any, rule determined whether populations chose or rejected sexual reproduction. But the experiment did indicate that sexual reproduction drove the Ramp population away from local maxima and freed it to seek higher peaks. This was a logical consequence of sex, which had a built-in risk of reducing fitness when relatively optimal genes were "watered down" when combined with a less fit set of genes. This result made Hillis feel he was on to something. He suspected that his experiments could unearth clues that biologists were missing.

Pitting Ramps Against Human Intelligence

At the time, Hillis was training his Ramps to tackle a thorny problem. In this experiment, the measure of a Ramp's fitness was the ability of the organism to sort numbers. The degree to which each Ramp sorted a list of sixteen numbers in descending order determined how successful that individual was. Sorting-network problems were a familiar challenge to computer hackers, who tried to build systems that arrange numbers using the fewest steps, or exchanges the fewer exchanges, the more wizardry required. Using his modified version of the genetic algorithm, Hillis would seed his next generation with the organisms of the current generation who sorted the list in the fewest exchanges.

The particular sorting network problem Hillis chose for his Ramps had long been used both as a benchmark of programming skill and as a proving ground for theoretical approaches to data manipulation. Essentially, he was pitting his Ramps, with evolution as a cornerman, against the cream of human endeavor. In 1962 a pair of computer scientists had published an article claiming the best possible solution to sorting networks, a system that would sort the sixteen integers in sixty-five exchanges. Two years later, the guru of code crunchers, Donald Knuth, created a system requiring only sixty-three exchanges. In 1969 the computer world was astonished when someone did it in one fewer exchange, and later that year the amazement was "tripled," to quote Knuth, when even that solution was eclipsed. This ultimate sorting program, written by a man named Milton Green, was elegant enough to arouse suspicion of Faustian dealings. It involved only sixty exchanges.

Hillis was running a population of 64 K, or 65,536, Ramps to evolve themselves to a state where they could become computer programs that solved the problem. (Because each individual could be assigned its own processor in the Connection Machine, the experiment could be run with lightning speed.) He generally ran his populations for five thousand or more generations, a time-consuming process for biologists breeding fruit flies but a day's work in the Cambridge headquarters of Thinking Machines. As was typical in Hillis's experiments, relatively capable individuals emerged in early runs and spread their highly fit genes throughout the population; by the end of the run the population had long found itself at a local maximum. The best of these sorted all the numbers successfully, requiring sixty-five exchanges.

This was a fairly impressive performance by a set of initially random numbers who found their way to a solution without human intervention; after all, two very satisfied computer scientists in 1962 had published their results in a paper, to some acclaim, after accomplishing the task in the same number of steps. But Hillis wanted his system to find the higher peaks in the landscape. He introduced conditions that would drag the population from its stagnant perch and force it to seek the higher ground. First, he tried increasing the rate of mutations. Although this drove the population off the hill, the mutations triggered dire lapses in fitness from which the population rarely recovered.

AntiRamps and the Red Queen

The breakthrough came when Hillis heard about the Red Queen hypothesis. The appellation was borrowed from Alice in Wonderland, wherein the Red Queen goaded the young protagonist into running furiously, although it seemed to advance the young girl not an inch. When she complained to the queen, Her Majesty informed her that constant running was required to remain in the same place. Biologists used this anecdote in describing "evolutionary arms races," when two populations of differing species were set against each other, in predatorprey or host-parasite relationships. Regarding the latter, if a host population evolved strategic traits to foil the parasite, the parasite would in turn evolve a strategy to compensate. William Hamilton, among others, had suggested that the presence of parasites might have been integral in accelerating the pace of evolution to a rate capable of yielding its present diversity and complexity; he had even run his own computer simulations, which indicated that organisms might have adopted sexual reproduction to thwart parasitic invasions on their offspring. Hillis decided to introduce parasites in his system.

Hillis called his parasites "anti-Ramps." Like their rivals, they were rewarded according to a fitness function the degree to which they harassed their digital cousins. It was a classic evolutionary arms race in which both species would coevolve and discover improvements in response to their opponent's evolutionary improvements. Hillis arranged an ingenious method of attack: the anti-Ramps literally provided test cases to gauge the Ramps' solutions to the sorting problems. As the simulation progressed, and as the Ramps came up with better solutions, the anti-Ramps would evolve increasingly challenging test cases. If one thought of the Ramps as chess players, the anti-Ramps were chess impresarios, who produced a series of opponents. They first ushered in fumbling novices, provided experienced players when the beginners were consistently vanquished, and eventually flew in cunning grand masters.

Under continual attack from these demanding challengers, the Ramps were forced to devise evolutionary strategies that would maintain and even improve the quality of their sorting. The first strategy was to settle on a fortified arrangement of integers so that the bite of the anti-Ramp would not be fatal. The second was to proceed with a high genetic variation from one generation to the next in order to assure that the parents' defects would not always be passed to the offspring and that the predators would be kept off balance.

Without anti-Ramps to keep them honest, previous populations of Ramps required tens of thousands of generations and long nights of Connection Machine time to find their local maxima. When coevolving parasites invaded the evolutionary landscape, however, a different story unfolded. Hillis liked to show videotapes of the screen display from the Connection Machine that illustrated this genetic drama. Each pixel on the screen represented an individual Ramp, and the fitness of each Ramp (how well it sorted the numbers) was represented by an arbitrarily assigned color. (The anti-Ramps were not depicted, but their effects were apparent.) As individuals were eliminated, selected, and mated in each time step, a new generation would replace the current one on the screen, and the color of the new population reflected these offspring. Clusters formed, and sometimes waves of similar Ramps pulsated in apparent synchrony. At first, the Ramps began to improve their ability to sort the numbers; this was reflected by localized changes in color from blue to green. Each time pockets of Ramps stabilized, however, a grim apocalyptic wave swept over them. The anti-Ramps obviously had devised test cases that broke the sorting schemes. But some Ramps evolved solutions that both met the demands of the test cases and developed immunities to that breed of anti-Ramp. Clusters of these improved Ramps, now in greenish yellow, roiled and spread on the screen, attained stability, and were besieged once more. Each time they reappeared, the newly immunized Ramps bore a color indicating higher fitness. Soon, some Ramps appeared in bright red, which indicated a fitness that Hillis's previous Ramps could not have hoped to attain. They had landed on the evolutionary landscape equivalent of the Andes.

The very first time Danny Hillis tried this, the entire epochal struggle described above occurred in fifteen minutes. The Ramp population was harshly dislodged from its comfy maximum. Racing around the evolutionary landscape as though pursued by hellhounds, the Ramps found and scaled the elusive higher peak. Although their solution did not equal Green's championship sorting network of sixty, they did match the second-best total of sixty-two, a significant improvement from previous runs. Hillis ran the experiment a few times more, and the population, again hounded by coevolving anti-Ramps, managed to construct a sorting network of sixty-one exchanges. Hillis was ecstatic, especially because he had been under the misimpression that the best- ever total had been sixty-two exchanges; he thought that his Ramps had outstripped the apex of human achievement. When he rechecked the literature and realized that his system fell short, he was so disappointed that he soon shelved the sorting-network problem indefinitely.

The concept, however, of accelerated evolution through coevolving parasites stuck with Hillis. He was convinced that Red Queen behavior yielded effective results because coevolving rivals forced the evolutionary system away from stability, on the cusp of phase transitions such as the ones Langton postulated in his work on the (Cellular Automata) λ parameter. Thus the Ramps and anti-Ramps combined to drive the system to the edge of chaos, and a rich complexity of novel genotype strategies arose from that fertile computational region.