NATURAL COMPUTING

.....

DNA, Quantum Bits, and the Future of Smart Machines

DENNIS SHASHA
and CATHY LAZERE

W. W. NORTON & COMPANY
NEW YORK LONDON

PN / PT

Part II

HARNESSING LIFESTUFF

MODERN MANUFACTURING BEGAN WITH THE NOTION OF INTERchangeable parts, dating back to Gutenberg's movable type in the fifteenth century. By the eighteenth century, manufacturers became more concerned about the precision of the parts. Eli Whitney's interchangeable musket parts had a precision of 1/30th of an inch (about 1 millimeter). Machine tolerances today are typically 10 microns,100 times more precise than Whitney's. Optical tolerances are now measured in the nanometer range—one million times more precise than Whitney's. Designers now have the opportunity to build machines of exquisite precision for the task at hand.

Mainstream computer science is built upon algorithms. An algorithm is a method that is guaranteed to produce a correct response for a large class of stimuli with a specified efficiency. Think of algorithms as recipes—to produce a chocolate layer cake, combine specified ingredients in a recommended order to obtain a desired result.

For example, a "mergesort" algorithm puts items in order no matter what kinds of items are presented to it. While algorithms

Part Opener
Text

will always play a central role in computing, some problems are fundamentally not algorithmic.

Consider the following problem. You are to survive in Antarctica and keep equipment operating at any temperature down to -60° C (-76 ° F). You know that your shelter and clothes may suffer any one of many possible mishaps. An algorithmically-oriented computer scientist would complain that the problem is ill-posed. If the mishaps are great enough, it may not be possible to survive. But what if you had to design at least a decent solution to the problem?

As early as 1954, the mathematician Nils Aall Barricelli, working at the Institute for Advanced Studies at Princeton, tried to simulate evolution using a computer. In the natural world, evolution applies to organisms. In the mputational world, evolution applies to designs. In both cases, evolution can lead to beautiful results *without the benefit of a conscious designer*. In 1975, John Holland of the University of Michigan wrote a landmark book *Adaptation in Natural and Artificial Systems* where he showed the commonalities among the different approaches to evolutionary design and improved them through a uniform mathematical framework.

Holland's framework became the basis for modern genetic (sometimes called evolutionary) algorithms. It consists of repeated applications of the following procedure:

1. Start with a population of candidate designs.
2. Evaluate each one to give a "fitness" score perhaps based on monetary cost or energy consumption;
3. Remember the design receiving the highest fitness score.

4. Create a new population by selecting the fittest candidate designs and changing them slightly in a random way or changing them greatly by combining different designs together.

Suppose you are designing a car through this method. If a good design proposes a composite chassis with a six cylinder engine and another good design proposes an aluminum chassis with an electric engine, the combined design might be a composite chassis with an electric engine. Parent designs beget children having some characteristics of each.

While evolution can lead to better designs, small adaptations require lesser efforts. For example, you can learn to ride a bicycle or juggle without evolving. Adaptations at that level may entail trial and error, but the organism doesn't need to change. **Rodney Brooks** does adaptation in motion. Since the 1980s he has designed robots that move intelligently by adaptation. Since starting his pioneering work, Brooks has been inspired by insects, elephants, and geckos. In the process, he has redefined what it means for robots to be smart.

Suppose you are designing software for a robot that must navigate the surface of another planet. You don't know what the exact task will be. You do know the ground is rough. You also know the environment is extremely hostile, but you don't know the particulars. You are faced, in fact, not only with unknowns but with what **Glenn Reeves** of the Jet Propulsion Laboratory calls "unknown unknowns" – unknowns you can't even characterize. It won't work to design a Rover, send it, and then hope for the best. In fact, the current state of the art is to diagnose from afar—100 million miles away.

Num. List

TEXT

Chap. Frontis

Caption

Caption to come for opener images?

Chap. Quote

* * *

> You're all assuming there's a logical representation inside the robot. What if there is no logical representation? I had been watching insects and how they do stuff . . . do they really have a three-dimensional rendering of the world around them—a computer graphics model inside that puny little head with 50,000 neurons. Is that what those neurons are doing?
>
> —*Rodney Brooks*

CN / CT

Chapter 1

RODNEY BROOKS

Sub-CT

Animals Rule

WHEN ARTIFICIAL INTELLIGENCE WAS BORN IN THE 1950s, doing well at IQ tests or chess seemed to be a good indication of intelligence. After all, that's what schools measured. Since then, a slew of other definitions have been added to the mix, including emotional IQ and Howard Gardner's interpersonal and kinesthetic measures of intelligence. But if we define intelligence as the ability to survive in the world, we need to look at more fundamental skills. How is it that we can walk, recognize objects and navigate around obstacles? You may say "Animals can do that!" To which, Rodney Brooks might respond, "Exactly!" In fact, robots might do better if they didn't copy humans in all aspects of our behavior. For example, who walks better over rough terrain— humans or insects? If you've ever seen insects scramble out of impossible holes, you might vote for the insects.

In a seminal paper from 1990 entitled "Elephants don't play chess," Rodney Brooks presented an evolutionary argument for the relative insignificance of human "higher" intelligence. Life arose on earth 3.5 billion years ago, he noted; vertebrates and insects the last 10% of that time, approximately 450 million years ago. The great apes emerged in the last 0.5% of that period, about 18 million years ago. Agriculture was created only 19,000 years ago, 0.0005% of life's time on earth. Expert knowledge has appeared only in the last few hundred years.

Computers are most successful at rapidly performing the skills learned in the last few hundred years of human history, perhaps because we are most conscious of those skills and they take the most conscious effort. But our unconscious acts pose a greater computing challenge. In the brief history of space travel it has been easier to build a computer program to guide a spacecraft to Mars than to build a robot that could navigate over rough terrain with anything like the skill of a billygoat. *Evolution required billions of years to arrive at the billygoat, but only a few million more to arrive at human intelligence.* The factor of a thousand in relative time scales should give us a certain humility before these "primitive" intelligences.

When Brooks wrote his Elephant paper, the field of Artificial Intelligence (AI) modeled intelligence as symbol manipulation. The scientific goal was to design sensor modules such as vision systems. These would abstract the world into symbols and pass those into an intelligent core, a kind of electronic monarch. The monarch would manipulate the symbols and then instruct actuators (normally wheels) to move. In many ways, this mirrored the idealized hierarchy of a large corporation or the military – "brains" on top, eyes and limbs on the bottom. Brooks objected to this paradigm on philosophical as well as pragmatic grounds.

Brooks was born in 1954 in Adelaide, Australia. While not exactly the Outback, this was a long way from the centers of computer science research. This may have been an advantage. Nobody told him the right way to approach the field. At age eight, on his own, he began designing computers to play games. At 12, he built one out of old telephone relays to play tic-tac-toe. He resolved to pursue a career in game design.

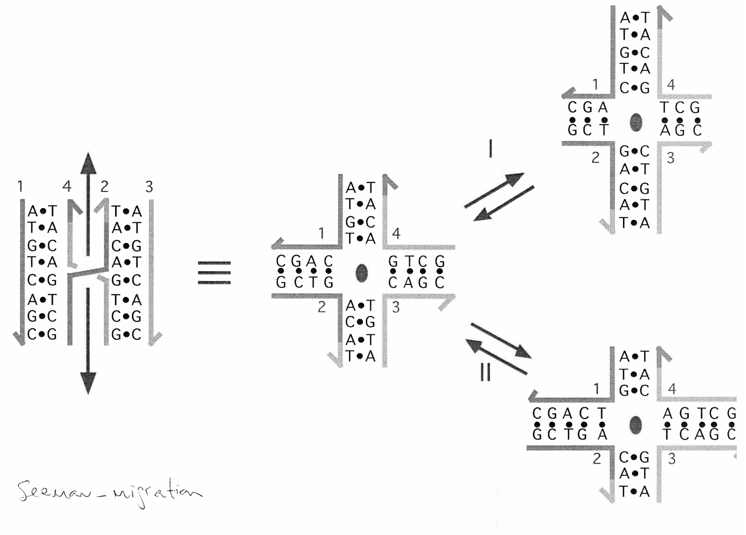

Some of the ways strands of DNA may escape linearity. Watson-Crick pairing (A binds with T and C binds with G) can create lots of different shapes. Credit: Ned Seeman

In 1972, Brooks began studies at Flinders University of South Australia. On the weekends he was permitted to use the lone university computer with its 16 kilobytes (roughly 16,000 bytes) of memory and a one megabyte disk. Its million bytes were only one millionth of the memory capacity of a contemporary desktop computer. Still, a megabyte was a lot more than 16 kilobytes. Brooks figured out how to program the computer as if it had the full one megabyte of memory by moving data from the disk

RHs:
L: Part
R: Chapter

Art

Caption

when necessary. He used an innovation called “virtual memory” that had been realized only a few years before. Brooks didn’t look for papers describing how to do it. He just did it. “Someone had described to me the idea,” says Brooks. “It sounded pretty good so I implemented virtual memory on this computer. “

Brooks’s doctoral thesis was on machine vision, a classic AI approach in which the camera would feed a computer an image. It would then translate the visual scene into symbols to be processed by a hypothetical “intelligence”—the symbol-manipulating monarch. “The basic idea that nobody was questioning was that you’ve got a camera, you’ve got pixels and you just change the pixels into a logical description of the world. “

Space Break

ON VACATION IN THAILAND in 1988, he visited his first wife’s family home which stood on stilts by a river. No one spoke English so Brooks sat by himself, watching the insects.

The more he watched, the more he began to question the symbolic AI paradigm. He just couldn’t believe that insects were capable of forming logical descriptions inside “that puny little head with 50,000 neurons.”

In 1990, his paper “Elephants Don’t Play Chess” explained what he playfully called “Nouvelle AI.” His hypothesis was that an intelligent system had to have its representations grounded in the physical world. “The world is the best model of itself,” as he put it. The world is up to date and contain all necessary details. This meant that Brooks’s robots would dispense with the hierarchical structure of “classical” AI with its symbolic representation of the world. Instead “nouvelle AI” robots would possess a set of independently designed skills.

Priority Inversion: Three Tasks and a Lock

Box Title

Priority inversion occurs because of a turf war between three tasks and a “lock.” The low priority task L acquires the lock and begins to use a resource. The high priority task H stops L from running the resource and then H runs itself. But then H stops when it needs the lock and wants to use the resource. At this point, if you’re unlucky, a middle priority task M begins to run. M doesn’t need the lock so it executes without allowing L to complete and release the lock. In effect, M prevents H from executing, thus “inverting” priority.

Box Text

The best known solution to priority inversion, proposed by Lui Sha, Ragunathan Rajkumar, and John P. Lehoczky of Carnegie Mellon University, is to raise the priority of L to the same level as H’s priority when H finds itself in need of the lock that L has. That is, L “inherits” the priority of H. The net effect is to give priority of L over M until L releases the lock.*

* * * *

*Aidan Daly, our illustrator, suggests the following good consequence of priority inversion from Tolkien’s fictional Middle Earth:

Box FN

The Hobbits (low priority) have a lock on the One Ring

Sauron (high priority) cannot be resurrected and rule without the One Ring

Thus, Men (medium priority) can rule in peace

Just as a human plays basketball and walks using the same limbs and eyes, a robot shares sensors and actuators for different skills. The skills, however, are independent —some of them, especially the highest level ones,. may fail without disturbing others. This approach follows the paradigm:

Equation

question → model → program → calculate → answer

Following this philosophy, the Brooks lab at MIT built an early robot called Allen in 1985, which Brooks puckishly named after Allen Newell, one of the early proponents of symbolic AI. Allen had three skills: avoid collisions, wander around randomly, and go to distant objects. "Allen would happily sit in the middle of a room until approached, then scurry away, avoiding collisions as it went," remembers Brooks. "The internal representation was that every sonar return represented a repulsive force."

C Head

Principles of Analog Programming

It is a tradition in computer science when getting familiar with a language or new computer to write a program that prints the sentence "Hello, world." When Mills asks his students to create an "H" at first they think digitally like a dot matrix—points in the form of an H. "An analog way of thinking is to ask whether you can put something in there to create hills and valleys. If you would look at them as a topographic map, you could create an H." That's the "Hello, world!" program for analog.

So, at the lowest level, Allen acted like a frightened mouse, following a primal rule—avoid hitting or being hit. What would keep Allen from hiding in a corner? Every ten seconds, Allen would be told to wander randomly. Note that wandering also requires moving wheels, in accordance with the Brooks strategy of having the different skills use the same actuators. Yet here is Feynman talking about some possibilities:

Extract

> *Biology is not simply writing information; it is doing something about it. A biological system can be exceedingly small. Many of the cells are very tiny, but they are very active.*

At the third layer, the robot used its sonar to look for distant places and try to go to them. It would measure distance using an odometer. Like a runner trying to complete a mountain race while avoiding a slip off the edge, the robot combined goal-seeking with underlying survival skills. The question is where might this all lead? Let's start with the technology.

D Head (run in)

Computing will leave its self-imposed digital electronic prison. Industrial design schools discovered the principle "form follows function" in the mid-twentieth century. It seemed almost too easy to be research. "I argue for simplicity," says Brooks. "Get away from hairy equations."

Computers will be expected to fend for themselves. He contrasts that viewpoint with that some of his colleagues and critics who think the hairier the better. To Brooks, if you need to explain something with a lot of convoluted mathematics, that means the solution will be "pretty unstable." "I'm interested in building something that can't fail to work," he says.

Rolling robots were one thing. What about walking robots? All that thinking in Thailand would be put into play. Working with Colin Angle and Grinnell Moore, a high school student, Brooks built a six- legged walking robot named Genghis.

BM Head

NATURAL COMPUTING TIMELINE

Timeline

1673	Gottfried Wilhelm von Leibniz invents a machine to do multiplication
1805	Joseph Marie Jacquard's makes weaving device based on holes punched in cards
1821	Charles Babbage designs his Analytical Engine to do calculations
1859	Charles Darwin proposes the theory of natural selection in his *On the Origin of Species*
1864	Herbert Spencer publishes *Principles of Biology* which applies key concepts of evolution to the social sciences
1866	Gregor Mendel publishes his paper *Experiments in Plant Hybridization* which becomes the basis of modern genetics
1921	Word "robot" coined by Czech playwright Karel Capek in production of RUR (Rossums Universal Robots)
1927	Vannevar Bush and MIT team begin design of differential analyzer, a sophisticated mechanical analog computer

BM Head

REFERENCES

C-1 Head

Chapter 1: **RODNEY BROOKS**

Notes

1. "Fast, Cheap and Out of Control: a robot invasion of the solar system" by Rodney A. Brooks, Anita M. Flynn, Journal of the British Interplanetary Society, 1989, pp. 478-485.
2. "Elephants Don't Play Chess" Rodney A. Brooks, Robots and Autonomous System, pp. 3-15
3. "Biologically inspired climbing with a hexapedal robot" M. J. Spenko, G. C. Haynes, J. A. Saunders, A. A. Rizzi, R. J. Full, D. E. Koditschek, Journal of Field Robotics, April 2008, pp. 223-242.

Chapter 2: **GLENN REEVES AND ADRIAN STOICA**

1. "What really happened on Mars?" by Glenn Reeves, The Risks Digest Volume 19: Issue 49.
2. An overview of the Mars exploration rovers' flight software, Reeves, G.E. Snyder, J.F., Systems, Man and Cybernetics, 2005 IEEE International Conference, 10-12 Oct. 20053) *Postcards from Mars* (Hardcover).
4. *Genetic Programming III—Darwinian Invention and Problem Solving* by John R. Koza, F. H. Bennett III, D. Andre, and M. A. Keane, San Francisco: Morgan Kaufman, 1999, ISBN 1-55860-543-6.

Caption to come for opener images?

. . . .

The market goes up and then goes back down—for no apparent reason. You see it all the time. It seems repeatable, at least partially. At a micro market level, the finer the time series, the more repeatable it appears. Next year, Microsoft's price might be double or half of its current value, but in the next microsecond it's only going up a penny or down a penny.

—*Jake Loveless*

Chapter 4

JAKE LOVELESS

. . . .

Riding the Big One

WHEN ARTIFICIAL INTELLIGENCE WAS BORN IN THE 1950s, doing well at IQ tests or chess seemed to be a good indication of intelligence. After all, that's what schools measured. Since then, a slew of other definitions have been added to the mix, including emotional IQ and Howard Gardner's interpersonal and kinesthetic measures of intelligence. But if we define intelligence as the ability to survive in the world, we need to look at more fundamental skills. How is it that we can walk, recognize objects and navigate around obstacles? You may say "Animals can do that!" To which, Rodney Brooks might respond, "Exactly!" In fact, robots might do better if they didn't copy humans in all aspects of our behavior. For example, who walks better over rough terrain— humans or insects? If you've ever seen insects scramble out of impossible holes, you might vote for the insects.

In a seminal paper from 1990 entitled "Elephants don't play chess," Rodney Brooks presented an evolutionary argument for the relative insignificance of human "higher" intelligence. Life arose on earth 3.5 billion years ago, he noted; vertebrates and in