



Meet Shaky, the first electronic person

The fascinating
and fearsome reality
of a machine with
a mind of its own

Computer scientist Charles Rosen communes with Shaky, the intelligent machine he helped create.





by Brad Darrach

It looked at first glance like a Good Humor wagon sadly in need of a spring paint job. But instead of a tinkly little bell on top of its box-shaped body there was this big metallic whangdoodle that came rearing up, full of lenses and cables, like a junk-sculpture gargoyle.

"Meet Shaky," said the young scientist who was showing me through the Stanford Research Institute. "The first electronic person."

I looked for a twinkle in the scientist's eye. There wasn't any. Sober as an equation, he sat down at an input terminal and typed out a terse instruction which was fed into Shaky's "brain," a computer set up in a nearby room: PUSH THE BLOCK OFF THE PLATFORM.

Something inside Shaky began to hum. A large glass prism shaped like a thick slice of pie and set in the middle of what passed for his face spun faster and faster till it dissolved into a glare. Then his superstructure made a slow 360° turn and his face leaned forward and seemed to be staring at the floor. As the hum rose to a whir, Shaky rolled slowly out of the room, rotated his superstructure again and turned left down the corridor at about four miles an hour, still staring at the floor.

"Guides himself by watching the baseboards," the scientist explained as we hurried to keep up. At every open door Shaky stopped, turned his head, inspected the room, turned away and rolled on to the next open door. In the fourth room he saw what he was looking for: a platform one foot high and eight feet long with a large wooden block sitting on it. He went in, then stopped short in the middle of the room and stared for about five seconds at the platform. I stared at it too.

"He'll never make it," I found myself thinking. "His wheels are too small." All at once I got gooseflesh. "Shaky," I realized, "is thinking the same thing I am thinking!"

Shaky was also thinking faster. He rotated his head slowly till his eye came to rest on a wide shallow ramp that was lying on the floor on the other side of the room. Whirring briskly, he crossed to the ramp, semicircled it and then pushed it straight across the floor till the high end of the ramp hit the platform. Rolling back a few feet, he cased the situation again and discovered that only one corner of the ramp was touching the platform. Rolling quickly to the far side of the ramp, he nudged it till the gap closed. Then he swung around, charged up the slope, located the block and gently pushed it off the platform.

Compared to the glamorous electronic elves who trundle across television screens, Shaky may not seem like much. No death-ray eyes, no secret transistorized lust for nubile lab technicians. But in fact he is a historic achievement. The task I saw him perform would tax the talents of a lively 4-year-old child, and the men who over the last two years have headed up the Shaky project—Charles Rosen, Nils Nilsson and Bert Raphael—say he is capable of far more sophisticated routines. Armed with the right devices and

CONTINUED

Computers will be playing office politics

CONTINUED

programmed in advance with basic instructions, Shaky could travel about the moon for months at a time and, without a single beep of direction from the earth, could gather rocks, drill cores, make surveys and photographs and even decide to lay plank bridges over crevices he had made up his mind to cross.

The center of all this intricate activity is Shaky's "brain," a remarkably programmed computer with a capacity of more than 7 million "bits" of information. In defiance of the soothing conventional view that the computer is just a glorified abacus that cannot possibly challenge the human monopoly of reason, Shaky's brain demonstrates that machines can think. Variouslly defined, thinking includes such processes as "exercising the powers of judgment" and "reflecting for the purpose of reaching a conclusion." In some of these respects—among them powers of recall and mathematical agility—Shaky's brain can think better than the human mind.

Marvin Minsky of MIT's Project Mac, a 42-year-old polymath who has made major contributions to Artificial Intelligence, recently told me with quiet certitude: "In from three to eight years we will have a machine with the general intelligence of an average human being. I mean a machine that will be able to read Shakespeare, grease a car, play office politics, tell a joke, have a fight. At that point the machine will begin to educate itself with fantastic speed. In a few months it will be at genius level and a few months after that its powers will be incalculable."

I had to smile at my instant credulity—the nervous sort of smile that comes when you realize you've been taken in by a clever piece of science fiction. When I checked Minsky's prophecy with other people working on Artificial Intelligence, however, many of them said that Minsky's timetable might be somewhat wishful—"give us 15 years," was a common remark—but all agreed that there would be such a machine and that it could precipitate the third Industrial Revolution, wipe out war and poverty and roll up centuries of growth in science, education and the arts. At the same time a number of computer scientists fear that the godsend may become a Golem. "Man's limited mind," says Minsky, "may not be able to control such immense mentalities."

Intelligence in machines has developed with surprising speed. It was only 33 years ago that a mathematician named Ronald Turing proved that a computer, like a brain, can process any kind of information—words as well as numbers, ideas as easily as facts; and now there is Shaky, with an inner core resembling the central nervous system of human beings. He is made up of five major systems of circuitry that correspond quite closely to basic human faculties—sensation, reason, language, memory, ego—and these faculties cooperate harmoniously to produce something that actually does behave very much like a rudimentary person.

Shaky's memory faculty, constructed after a model de-

CONTINUED

Shaky could learn to recognize a familiar face

CONTINUED

veloped at MIT, takes input from Shaky's video eye, optical range finder, telemetering equipment and touch-sensitive antennae; taste and hearing are the only senses Shaky so far doesn't have. This input is then routed through a "mental process" that recognizes patterns and tells Shaky what he is seeing. A dot-by-dot impression of the video input, much like the image on a TV screen, is constructed in Shaky's brain according to the laws of analytical geometry. Dark areas are separated from light areas, and if two of these contrasting areas happen to meet along a sharp enough line, the line is recognized as an edge. With a few edges for clues, Shaky can usually guess what he's looking at (just as people can) without bothering to fill in all the features on the hidden side of the object. In fact, the art of recognizing patterns is now so far advanced that merely by adding a few equations Shaky's creators could teach him to recognize a familiar human face every time he sees it.

Once it is identified, what Shaky sees is passed on to be processed by the rational faculty—the cluster of circuits that actually does his thinking. The forerunners of Shaky's rational faculty include a checker-playing computer program that can beat all but a few of the world's best players, and Mac Hack, a chess-playing program that can already outplay some gifted amateurs and in four or five years will probably master the masters. Like these programs, Shaky thinks in mathematical formulas that tell him what's going on in each of his faculties and in as much of the world as he can sense. For instance, when the space between the wall and the desk is too small to ease through, Shaky is smart enough to know it and to work out another way to get where he is going.

Shaky is not limited to thinking in strictly logical forms. He is also learning to think by analogy—that is, to make himself at home in a new situation, much the way human beings do, by finding in it something that resembles a situation he already knows, and on the basis of this resemblance to make and carry out decisions. For example, knowing how to roll up a ramp onto a platform, a slightly more advanced Shaky equipped with legs instead of wheels and given a similar problem could very quickly figure out how to use steps in order to reach the platform.

But as Shaky grows and his decisions become more complicated, more like decisions in real life, he will need a way of thinking that is more flexible than either logic or analogy. He will need a way to do the sort of ingenious, practical "soft thinking" that can stop gaps, chop knots, make the best of bad situations and even, when time is short, solve a problem by making a shrewd guess.

The route toward "soft thinking" has been charted by the founding fathers of Artificial Intelligence, Allen Newell and Herbert Simon of Carnegie-Mellon University. Before Newell and Simon, computers solved (or failed to solve) nonmathematical problems by a hopelessly tedious process of trial and error. "It was like looking up a name in a big-city telephone book that nobody has bothered to ar-

CONTINUED

The best machines can't quite tell words from noise

CONTINUED

range in alphabetical order," says one computer scientist. Newell and Simon figured out a simple scheme—modeled, says Minsky, on "the way Herb Simon's mind works." Using the Newell-Simon method, a computer does not immediately search for answers, but is programmed to sort through general categories first, trying to locate the one where the problem and solution would most likely fit. When the correct category is found, the computer then works within it, but does not rummage endlessly for an absolutely perfect solution, which often does not exist. Instead, it accepts (as people do) a good solution, which for most non-numerical problems is good enough. Using this type of programming, an MIT professor wrote into a computer the criteria a certain banker used to pick stocks for his trust accounts. In a test, the program picked the same stock the banker did in 21 of 25 cases. In the other four cases the stocks the program picked were so much like the ones the banker picked that he said they would have suited the portfolio just as well.

Shaky can understand about 100 words of written English, translate these words into a simple verbal code and then translate the code into the mathematical formulas in which his actual thinking is done. For Shaky, as for most computer systems, natural language is still a considerable barrier. There are literally hundreds of "machine languages" and "program languages" in current use, and computers manipulate them handily, but when it comes to ordinary language they're still in nursery school. They are not very good at translation, for instance, and no program so far created can cope with a large vocabulary, much less converse with ease on a broad range of subjects. To do this, Shaky and his kind must get better at working with symbols and ambiguities (the dog in the window had hair but it fell out). It would also be useful if they learned to follow spoken English and talk back, but so far the machines have a hard time telling words from noise.

Language has a lot to do with learning, and Shaky's ability to acquire knowledge is limited by his vocabulary. He can learn a fact when he is told a fact, he can learn by solving problems, he can learn from exploration and discovery. But up to now neither Shaky nor any other computer program can browse through a book or watch a TV program and grow as he goes, as a human being does. This fall, Minsky and a colleague named Seymour Papert opened a two-year crash attack on the learning problem by trying to teach a computer to understand nursery rhymes. "It takes a page of instructions," says Papert, "to tell the machine that when Mary had a little lamb she didn't have it for lunch."

Shaky's ego, or executive faculty, monitors the other faculties and makes sure they work together. It starts them, stops them, assigns and erases problems; and when a course of action has been worked out by the rational faculty, the ego sends instructions to any or all of Shaky's six small on-board motors—and away he goes. All these separate systems merge smoothly in a totality more intricate than many forms of sentient life and they work together with wonderful agility and resourcefulness. When, for example, it turns out that the platform isn't there because somebody has moved it, Shaky spins his superstructure, finds the platform again and keeps pushing the ramp till he gets it where he wants it—and if you happen to be the somebody who

CONTINUED

Psychoanalyzing a computer with a Freudian complex

CONTINUED

has been moving the platform, says one SRI scientist, "you get a strange prickling at the back of your neck as you realize that you are being hunted by an intelligent machine."

With very little change in program and equipment, Shaky now could do work in a number of limited environments: warehouses, libraries, assembly lines. To operate successfully in more loosely structured scenes, he will need far more extensive, more nearly human abilities to remember and to think. His memory, which supplies the rest of his system with a massive and continuous flow of essential information, is already large, but at the next step of progress it will probably become monstrous. Big memories are essential to complex intelligence. The largest standard computer now on the market can store about 36 million "bits" of information in a six-foot cube, and a computer already planned will be able to store more than a trillion "bits" (one estimate of the capacity of a human brain) in the same space.

Size and efficiency of hardware are less important, though, than sophistication in programming. In a dozen universities, psychologists are trying to create computers with well-defined humanoid personalities. Aldous, developed at the University of Texas by a psychologist named John Loehlin, is the first attempt to endow a computer with emotion. Aldous is programmed with three emotions and three responses, which he signals. Love makes him signal approach, fear makes him signal withdrawal, anger makes him signal attack. By varying the intensity and probability of these three responses, the personality of Aldous can be drastically changed. In addition, two or more different Aldouses can be programmed into a computer and made to interact. They go through rituals of getting acquainted, making friends, having fights.

Even more peculiarly human is the program created by Stanford psychoanalyst Kenneth M. Colby. Colby has developed a Freudian complex in his computer by setting up conflicts between beliefs (I must love Father, I hate Father). He has also created a computer psychiatrist and when he lets the two programs interact, the "patient" resolves its conflicts just as a human being does—by forgetting about them, lying about them or talking truthfully about them with the "psychiatrist." Such a large store of possible reactions has been programmed into the computer—and there are so many possible sequences of question and answer—that Colby can never be exactly sure what the "patient" will decide to do.

Colby is currently attempting to broaden the range of emotional reactions his computer can experience. "But so far," one of his assistants says, "we have not achieved computer orgasm."

Knowledge that comes out of these experiments in "sophistication" is helping to lead toward the ultimate sophistication—the autonomous computer that will be able to write its own programs and then use them in an approximation of the independent, imaginative way a human being dreams up projects and carries them out. Such a machine is now

CONTINUED

CONTINUED

being developed at Stanford by Joshua Lederberg (the Nobel Prize-winning geneticist) and Edward Feigenbaum. In using a computer to solve a series of problems in chemistry, Lederberg and Feigenbaum realized their progress was being held back by the long, tedious job of programming their computer for each new problem. "That started me wondering," says Lederberg. "Couldn't we save ourselves work by teaching the computer how we write these programs, and then let it program itself?"

Basically, a computer program is nothing more than a set of instructions (or rules of procedure) applicable to a particular problem at hand. A computer can tell you that $1 + 1 = 2$ —not because it has that fact stored away and then finds it, but because it has been programmed with the rules for simple addition. Lederberg decided you could give a computer some general rules for programming; and now, based on his initial success in teaching a computer to write programs in chemistry, he is convinced that computers can do this in any field—that they will be able in the reasonably near future to write programs that write programs that write programs . . .

This prospect raises a haunting question: won't computers then be just as independent as human beings are? Peter Ossorio, a philosopher at the University of Colorado who has pondered the psychology of computers, says that autonomy is part of the computer's inherent nature. "Free will," Ossorio says, "is a characteristic of serial processors—of all systems that do one thing after another and there-

fore have more options than they are able to use. Serial systems naturally have to make choices among alternatives. People are serial systems and so are computers."

Many computer scientists believe that people who talk about computer autonomy are indulging in a lot of cybernetic hoopla. Most of these skeptics are engineers who work mainly with technical problems in computer hardware and who are preoccupied with the mechanical operations of these machines. Other computer experts seriously doubt that the finer psychic processes of the human mind will ever be brought within the scope of circuitry, but they see autonomy as a prospect and are persuaded that the social impact will be immense.

Up to a point, says Minsky, the impact will be positive—"The machine dehumanized man, but it could rehumanize him." By automating all routine work and even tedious low-grade thinking, computers could free billions of people to spend most of their time doing pretty much as they damn please. But such progress could also produce quite different results. "It might happen," says Herbert Simon, "that the Puritan work ethic would crumble too fast and masses of people would succumb to the diseases of leisure." An even greater danger may lie in man's increasing and by now irreversible dependency upon the computer. The electronic circuit has already replaced the dynamo at the center of technological civilization. Many U.S. industries and businesses, the telephone and power grids, the airlines and the mail service, the systems for distributing food and, not least,

CONTINUED

There need be no question of computer malice'

CONTINUED

the big government bureaucracies would be instantly disrupted and threatened with complete breakdown if the computers they depend on were disconnected. The disorder in Western Europe and the Soviet Union would be almost as severe.

What's more, our dependency on computers seems certain to increase at a rapid rate. Doctors are already beginning to rely on computer diagnosis and computer-administered postoperative care. Artificial Intelligence experts believe that fiscal planners in both industry and government, caught up in deepening economic complexities, will gradually delegate to computers nearly complete control of the national (and even the global) economy. In the interests of efficiency, cost-cutting and speed of reaction, the Department of Defense may well be forced more and more to surrender human direction of military policies to machines that plan strategy and tactics. In time, say the scientists, diplomats will abdicate judgment to computers that predict, say, Russian policy by analyzing their own simulations of the entire Soviet state and of the personalities—or the computers—in power there.

Man, in short, is coming to depend on thinking machines to make decisions that involve his vital interests and even his survival as a species. What guarantee do we have that in making these decisions the machines will always consider our best interests? There is no guarantee unless we provide it, says Minsky, and it will not be easy to provide—after all, man has not been able to guarantee that his own decisions are made in his own best interests. Any supercomputer could be programmed to test important decisions for their value to human beings, but such a computer, being autonomous, could also presumably write a program that countermanded these “ethical” instructions. There need be no question of computer malice here, merely a matter of computer creativity overcoming external restraints.

The men at Project MAC foresee an even more unsettling possibility. A computer that can program a computer, they reason, will be followed in fairly short order by a computer that can design and build a computer vastly more complex and intelligent than itself—and so on indefinitely.

“I’m afraid the spiral could get out of control,” says Minsky. It is possible, of course, to monitor computers, to make an occasional check on what they are doing in there; but even now it is difficult to monitor the larger computers, and the computers of the future may be far too complex to keep track of.

Why not just unplug the thing if it got out of hand? “Switching off a system that defends a country or runs its entire economy,” says Minsky, “is like cutting off its food supply. Also, the Russians are only about three years behind us in A-I work. With our system switched off, they would have us at their mercy.”

The problem of computer control will have to be solved, Minsky and Papert believe, before computers are put in charge of systems essential to society's survival. If a com-

CONTINUED

If we're lucky, they might decide to keep us as pets'

CONTINUED

puter directing the nation's economy or its nuclear defenses ever rated its own efficiency above its ethical obligation, it could destroy man's social order—or destroy man. "Once the computers got control," says Minsky, "we might never get it back. We would survive at their sufferance. If we're lucky, they might decide to keep us as pets."

But even if no such catastrophe were to occur, say the people at Project MAC, the development of a machine more intelligent than man will surely deal a severe shock to man's sense of his own worth. Even Shaky is disturbing, and a creature that deposed man from the pinnacle of creation might tempt us to ask ourselves: Is the human brain outmoded? Has evolution in protoplasm been replaced by evolution in circuitry?

"And why not?" Minsky replied when I recently asked him these questions. "After all, the human brain is just a computer that happens to be made out of meat."

I stared at him—he was smiling. This man, I thought, has lived too long in a subtle tangle of ideas and circuits. And yet men like Minsky are admirable, even heroic. They have struck out on a Promethean adventure and you can tell by a kind of afterthought in their eyes that they are haunted by what they have done. It is the others who depress me, the lesser figures in the world of Artificial Intelligence, men who contemplate infinitesimal riddles of circuitry and never once look up from their work to wonder what effect it might have upon the world they scarcely live in. And what of the people in the Pentagon who are footing most of the bill in Artificial Intelligence research? "I have warned them again and again," says Minsky, "that we are getting into very dangerous country. They don't seem to understand."

I thought of Shaky growing up in the care of these careless people—growing up to be what? No way to tell. Confused, concerned, unable to affirm or deny the warnings I had heard at Project MAC, I took my questions to computer-memory expert Ross Quillian, a nice warm guy with a house full of dogs and children, who seemed to me one of the best-balanced men in the field. I hoped he would cheer me up. Instead, he said, "I hope that man and these ultimate machines will be able to collaborate without conflict. But if they can't, we may be forced to choose sides. And if it comes to a choice, I know what mine will be." He looked me straight in the eye. "My loyalties go to intelligent life, no matter in what medium it may arise." ■

Shaky and one of the computer scientists who helped create him eye each other in an office at Stanford. The scientists offhandedly refer to Shaky as "he," and, one says, "we have enough problems already without creating a female Shaky. But we've discussed it."

