

# Rise of the Robots

Originally published in December 1999

by Hans Moravec

**By 2050 robot “brains” based on computers that execute 100 trillion instructions per second will start rivaling human intelligence**

In recent years the mushrooming power, functionality and ubiquity of computers and the Internet have outstripped early forecasts about technology's rate of advancement and usefulness in everyday life. Alert pundits now foresee a world saturated with powerful computer chips, which will increasingly insinuate themselves into our gadgets, dwellings, apparel and even our bodies.

Yet a closely related goal has remained stubbornly elusive. In stark contrast to the largely unanticipated explosion of computers into the mainstream, the entire endeavor of robotics has failed rather completely to live up to the predictions of the 1950s. In those days, experts who were dazzled by the seemingly miraculous calculational ability of computers thought that if only the right software were written, computers could become the artificial brains of sophisticated autonomous robots. Within a decade or two, they believed, such robots would be cleaning our floors, mowing our lawns and, in general, eliminating drudgery from our lives.

Obviously, it hasn't turned out that way. It is true that industrial robots have transformed the manufacture of automobiles, among other products. But that kind of automation is a far cry from the versatile, mobile, autonomous creations that so many scientists and engineers have hoped for. In pursuit of such robots, waves of researchers have grown disheartened and scores of start-up companies have gone out of business.

It is not the mechanical “body” that is unattainable; articulated arms and other moving mechanisms adequate for manual work already exist, as the industrial robots attest. Rather it is the computer-based artificial brain that is still well below the level of sophistication needed to build a humanlike robot.

Nevertheless, I am convinced that the decades-old dream of a useful, general-purpose autonomous robot will be realized in the not too distant future. By 2010 we will see mobile robots as big as people but with cognitive abilities similar in many respects to those of a lizard. The machines will be capable of carrying out simple chores, such as vacuuming, dusting, delivering packages and taking out the garbage. By 2040, I believe, we will finally achieve the original goal of robotics and a thematic mainstay of science fiction: a freely moving machine with the intellectual capabilities of a human being.

## Reasons for Optimism

In light of what I have just described as a history of largely unfulfilled goals in robotics, why do I believe that rapid progress and stunning accomplishments are in the offing? My confidence is based on recent developments in electronics and software, as well as on my own observations of robots, computers and even insects, reptiles and other living things over the past 30 years.

The single best reason for optimism is the soaring performance in recent years of mass-produced computers. Through the 1970s and 1980s, the computers readily available to robotics researchers were capable of executing about one million instructions per second (MIPS). Each of these instructions represented a very basic task, like adding two 10-digit numbers or storing the result in a specified location in memory.

In the 1990s computer power suitable for controlling a research robot shot through 10 MIPS, 100 MIPS and has lately reached 1,000 in high-end desktop machines. Apple's new iBook laptop computer, with a retail price at the time of this writing of \$1,600, achieves more than 500 MIPS. Thus, functions far beyond the capabilities of robots in the 1970s and 1980s are now coming close to commercial viability.

For example, in October 1995 an experimental vehicle called Navlab V crossed the U.S. from Washington, D.C., to San Diego, driving itself more than 95 percent of the time. The vehicle's self-driving and navigational system was built around a 25-MIPS laptop based on a microprocessor by Sun Microsystems. The Navlab V was built by the Robotics Institute at Carnegie Mellon University, of which I am a member. Similar robotic vehicles, built by researchers elsewhere in the U.S. and in Germany, have logged thousands of highway kilometers under all kinds of weather and driving conditions.

In other experiments within the past few years, mobile robots mapped and navigated unfamiliar office suites, and computer vision systems located textured objects and tracked and analyzed faces in real time. Meanwhile personal computers became much more adept at recognizing text and speech.

Still, computers are no match today for humans in such functions as recognition and navigation. This puzzled experts for many years, because computers are far superior to us in calculation. The explanation of this apparent paradox follows from the fact that the human brain, in its entirety, is not a true programmable, general-purpose computer (what computer scientists refer to as a universal machine; almost all computers nowadays are examples of such machines).

To understand why this requires an evolutionary perspective. To survive, our early ancestors had to do several things repeatedly and very well: locate food, escape predators, mate and protect offspring. Those tasks depended strongly on the brain's ability to recognize and navigate. Honed by hundreds of millions of years of evolution, the brain became a kind of ultrasophisticated—but special-purpose—computer.

The ability to do mathematical calculations, of course, was irrelevant for survival. Nevertheless, as language transformed human culture, at least a small part of our brains evolved into a universal machine of sorts. One of the hallmarks of such a machine is its ability to follow an arbitrary set of instructions, and with language, such instructions could be transmitted and carried out. But because we visualize numbers as complex shapes, write them down and perform other such functions, we process digits in a monumentally awkward and inefficient way. We use hundreds of billions of neurons to do in minutes what hundreds of them, specially “rewired” and arranged for calculation, could do in milliseconds.

A tiny minority of people are born with the ability to do seemingly amazing mental calculations. In absolute terms, it's not so amazing: they calculate at a rate perhaps 100 times that of the average person. Computers, by comparison, are millions or billions of times faster.

### Can Hardware Simulate Wetware?

The challenge facing roboticists is to take general-purpose computers and program them to match the largely special-purpose human brain, with its ultraoptimized perceptual inheritance and other peculiar evolutionary traits. Today's robot-controlling computers are much too feeble to be applied successfully in that role, but it is only a matter of time before they are up to the task.

Implicit in my assertion that computers will eventually be capable of the same kind of perception, cognition and thought as humans is the idea that a sufficiently advanced and

sophisticated artificial system—for example, an electronic one—can be made and programmed to do the same thing as the human nervous system, including the brain. This issue is controversial in some circles right now, and there is room for brilliant people to disagree.

At the crux of the matter is the question of whether biological structure and behavior arise entirely from physical law and whether, moreover, physical law is computable—that is to say, amenable to computer simulation. My view is that there is no good scientific evidence to negate either of these propositions. On the contrary, there are compelling indications that both are true.

Molecular biology and neuroscience are steadily uncovering the physical mechanisms underlying life and mind but so far have addressed mainly the simpler mechanisms. Evidence that simple functions can be composed to produce the higher capabilities of nervous systems comes from programs that read, recognize speech, guide robot arms to assemble tight components by feel, classify chemicals by artificial smell and taste, reason about abstract matters and so on. Of course, computers and robots today fall far short of broad human or even animal competence. But that situation is understandable in light of an analysis, summarized in the next section, that concludes that today's computers are only powerful enough to function like insect nervous systems. And, in my experience, robots do indeed perform like insects on simple tasks.

Ants, for instance, can follow scent trails but become disoriented when the trail is interrupted. Moths follow pheromone trails and also use the moon for guidance. Similarly, many commercial robots today follow guide wires installed beneath the surface they move over, and some orient themselves using lasers that read bar codes on walls.

If my assumption that greater computer power will eventually lead to human-level mental capabilities is true, we can expect robots to match and surpass the capacity of various animals and then finally humans as computer-processing rates rise sufficiently high. If on the other hand the assumption is wrong, we will someday find specific animal or human skills that elude implementation in robots even after they have enough computer power to match the whole brain. That would set the stage for a fascinating scientific challenge—to somehow isolate and identify the fundamental ability that brains have and that computers lack. But there is no evidence yet for such a missing principle.

The second proposition, that physical law is amenable to computer simulation, is increasingly beyond dispute. Scientists and engineers have already produced countless useful simulations, at various levels of abstraction and approximation, of everything from automobile crashes to the “color” forces that hold quarks and gluons together to make up protons and neutrons.

### Nervous Tissue and Computation

If we accept that computers will eventually become powerful enough to simulate the mind, the question that naturally arises is: What processing rate will be necessary to yield performance on a par with the human brain? To explore this issue, I have considered the capabilities of the vertebrate retina, which is understood well enough to serve as a Rosetta stone roughly relating nervous tissue to computation. By comparing

how fast the neural circuits in the retina perform image-processing operations with how many instructions per second it takes a computer to accomplish similar work, I believe it is possible to at least coarsely estimate the information-processing power of nervous tissue—and by extrapolation, that of the entire human nervous system.

The human retina is a patch of nervous tissue in the back of the eyeball half a millimeter thick and approximately two centimeters across. It consists mostly of light-sensing cells, but one tenth of a millimeter of its thickness is populated by image-processing circuitry that is capable of detecting edges (boundaries between light and dark) and motion for about a million tiny image regions. Each of these regions is associated with its

full human-scale brainpower.

Commercial and research experiences convince me that the mental power of a guppy—about 1,000 MIPS—will suffice to guide mobile utility robots reliably through unfamiliar surroundings, suiting them for jobs in hundreds of thousands of industrial locations and eventually hundreds of millions of homes. Such machines are less than a decade away but have been elusive for so long that only a few dozen small research groups are now pursuing them.

Commercial mobile robots—the smartest to date, barely insectlike at 10 MIPS—have found few jobs. A paltry 10,000 work worldwide, and the companies that made them are struggling or defunct. (Makers of robot manipulators are not doing

## Asked why there are candles on the table, a third-generation robot might reply that it put them there because its owner likes candlelit dinners and it likes to please its owner.

own fiber in the optic nerve, and each performs about 10 detections of an edge or a motion each second. The results flow deeper into the brain along the associated fiber.

From long experience working on robot vision systems, I know that similar edge or motion detection, if performed by efficient software, requires the execution of at least 100 computer instructions. Thus, to accomplish the retina's 10 million detections per second would require at least 1,000 MIPS.

The entire human brain is about 75,000 times heavier than the 0.02 gram of processing circuitry in the retina, which implies that it would take, in round numbers, 100 million MIPS (100 trillion instructions per second) to emulate the 1,500-gram human brain. Personal computers in 1999 beat certain insects but lose to the human retina and even to the 0.1-gram brain of a goldfish. A typical PC would have to be at least a million times more powerful to perform like a human brain.

### Brainpower and Utility

Though dispiriting to artificial-intelligence experts, the huge deficit does not mean that the goal of a humanlike artificial brain is unreachable. Computer power for a given price doubled each year in the 1990s, after doubling every 18 months in the 1980s, and every two years before that. Prior to 1990 this progress made possible a great decrease in the cost and size of robot-controlling computers. Cost went from many millions of dollars to a few thousand, and size went from room-filling to handheld. Power, meanwhile, held steady at about 1 MIPS. Since 1990 cost and size reductions have abated, but power has risen to near 1,000 MIPS per computer. At the present pace, only about 30 or 40 years will be needed to close the millionfold gap. Better yet, useful robots don't need

much better.) The largest class of commercial mobile robots, known as Automatic Guided Vehicles (AGVs), transport materials in factories and warehouses. Most follow buried signal-emitting wires and detect end points and collisions with switches, a technique developed in the 1960s.

It costs hundreds of thousands of dollars to install guide wires under concrete floors, and the routes are then fixed, making the robots economical only for large, exceptionally stable factories. Some robots made possible by the advent of microprocessors in the 1980s track softer cues, like magnets or optical patterns in tiled floors, and use ultrasonics and infrared proximity sensors to detect and negotiate their way around obstacles.

The most advanced industrial mobile robots, developed since the late 1980s, are guided by occasional navigational markers—for instance, laser-sensed bar codes—and by preexisting features such as walls, corners and doorways. The costly labor of laying guide wires is replaced by custom software that is carefully tuned for each route segment. The small companies that developed the robots discovered many industrial customers eager to automate transport, floor cleaning, security patrol and other routine jobs. Alas, most buyers lost interest as they realized that installation and route changing required time-consuming and expensive work by experienced route programmers of inconsistent availability. Technically successful, the robots fizzled commercially.

In failure, however, they revealed the essentials for success. First, the physical vehicles for various jobs must be reasonably priced. Fortunately, existing AGVs, forklift trucks, floor scrubbers and other industrial machines designed for human riders or for following guide wires can be adapted for autonomy. Second, the customer should not have to call in specialists to

put a robot to work or to change its routine; floor cleaning and other mundane tasks cannot bear the cost, time and uncertainty of expert installation. Third, the robots must work reliably for at least six months before encountering a problem or a situation requiring downtime for reprogramming or other alterations. Customers routinely rejected robots that after a month of flawless operation wedged themselves in corners, wandered away lost, rolled over employees' feet or fell down stairs. Six months, though, earned the machines a sick day.

Robots exist that have worked faultlessly for years, perfected by an iterative process that fixes the most frequent failures, revealing successively rarer problems that are corrected in turn. Unfortunately, that kind of reliability has been achieved only for prearranged routes. An insectlike 10 MIPS is just enough to track a few handpicked landmarks on each segment of a robot's path. Such robots are easily confused by minor surprises such as shifted bar codes or blocked corridors (not unlike ants thrown off a scent trail or a moth that has mistaken a streetlight for the moon).

### A Sense of Space

Robots that chart their own routes emerged from laboratories worldwide in the mid-1990s, as microprocessors reached 100 MIPS. Most build two-dimensional maps from sonar or laser rangefinder scans to locate and route themselves, and the best seem able to navigate office hallways for days before becoming disoriented. Of course, they still fall far short of the six-month commercial criterion. Too often different locations in the coarse maps resemble one another. Conversely, the same location, scanned at different heights, looks different, or small obstacles or awkward protrusions are overlooked. But sensors, computers and techniques are improving, and success is in sight.

My small laboratory is in the race. In the 1980s we devised a way to distill large amounts of noisy sensor data into reliable maps by accumulating statistical evidence of emptiness or occupancy in each cell of a grid representing the surroundings. The approach worked well in two dimensions and guides many of the robots described above.

Three-dimensional maps, 1,000 times richer, promised to be much better but for years seemed computationally out of reach. In 1992 we used economies of scale and other tricks to reduce the costs of three-dimensional maps 100-fold. We now have a test program that accumulates thousands of measurements from stereoscopic camera glimpses to map a room's volume down to centimeter-scale. With 1,000 MIPS, the program digests over a glimpse per second, adequate for slow indoor travel.

### Robot, Version 1.0

One thousand MIPS is only now appearing in high-end desktop PCs. In a few years it will be found in laptops and similar smaller, cheaper computers fit for robots. To prepare for that day, we recently began an intensive three-year project to develop a prototype for commercial products based on such a computer. We plan to automate learning processes to optimize hundreds of evidence-weighting parameters and to write programs to find clear paths, locations, floors, walls, doors and other objects in the three-dimensional maps. We

will also test programs that orchestrate the basic capabilities into larger tasks, such as delivery, floor cleaning and security patrol.

The initial testbed will be a small camera-studded mobile robot. Its intelligence will come from two computers: an Apple iBook laptop on board the robot, and an off-board Apple G4-based machine with about 1,000 MIPS that will communicate wirelessly with the iBook. Tiny mass-produced digital camera chips promise to be the cheapest way to get the millions of measurements needed for dense maps.

As a first commercial product, we plan a basketball-size "navigation head" for retrofit onto existing industrial vehicles. It would have multiple stereoscopic cameras, generic software for mapping, recognition and control, a different program for its specific application (such as floor cleaning), and a hardware connection to vehicle power, controls and sensors. Head-equipped vehicles with transport or patrol programs could be taught new routes simply by leading them through once. Floor-cleaning programs would be shown the boundaries of their work area.

Introduced to a job location, the vehicles would understand their changing surroundings competently enough to work at least six months without debilitating mistakes. Ten thousand AGVs, 100,000 cleaning machines and, possibly, a million forklift trucks are candidates for retrofit, and robotization may greatly expand those markets.

### Fast Replay

Income and experience from spatially aware industrial robots would set the stage for smarter yet cheaper (\$1,000 rather than \$10,000) consumer products, starting probably with small robot vacuum cleaners that automatically learn their way around a home, explore unoccupied rooms and clean whenever needed. I imagine a machine low enough to fit under some furniture, with an even lower extendable brush, that returns to a docking station to recharge and disgorge its dust load. Such machines could open a true mass market for robots.

Commercial success will provoke competition and accelerate investment in manufacturing, engineering and research. Vacuuming robots ought to beget smarter cleaning robots with dusting, scrubbing and picking-up arms, followed by larger multifunction utility robots with stronger, more dexterous arms and better sensors. Programs will be written to make such machines pick up clutter, store, retrieve and deliver things, take inventory, guard homes, open doors, mow lawns, play games and so on. New applications will expand the market and spur further advances when robots fall short in acuity, precision, strength, reach, dexterity, skill or processing power. Capability, numbers sold, engineering and manufacturing quality, and cost-effectiveness will increase in a mutually reinforcing spiral. Perhaps by 2010 the process will have produced the first broadly competent "universal robots," as big as people but with lizardlike 5,000-MIPS minds that can be programmed for almost any simple chore.

Like competent but instinct-ruled reptiles, first-generation universal robots will handle only contingencies explicitly covered in their application programs. Unable to adapt to changing circumstances, they will often perform inefficiently or not at all. Still, so much physical work awaits them in businesses,



HANS MORAVEC

Robot vision would convey the key elements of a scene in a map useful for navigation. The latest maps, such as the one at the right, are three-dimensional and convey details down to centimeter-scale. The map was made from 20 stereoscopic images.

streets, fields and homes that robotics could begin to overtake pure information technology commercially.

A second generation of universal robot with a mouselike 100,000 MIPS will adapt as the first generation does not and will even be trainable. Besides application programs, such robots would host a suite of software “conditioning modules” that would generate positive and negative reinforcement signals in predefined circumstances. For example, doing jobs fast and keeping its batteries charged will be positive; hitting or breaking something will be negative. There will be other ways to accomplish each stage of an application program, from the minutely specific (grasp the handle underhand or overhand) to the broadly general (work indoors or outdoors). As jobs are repeated, alternatives that result in positive reinforcement will be favored, those with negative outcomes shunned. Slowly but surely, second-generation robots will work increasingly well.

A monkeylike five million MIPS will permit a third generation of robots to learn very quickly from mental rehearsals in simulations that model physical, cultural and psychological factors. Physical properties include shape, weight, strength, texture and appearance of things, and how to handle them. Cultural aspects include a thing’s name, value, proper location and purpose. Psychological factors, applied to humans and robots alike, include goals, beliefs, feelings and preferences. Developing the simulators will be a huge undertaking involving thousands of programmers and experience-gathering robots. The simulation would track external events and tune its models to keep them faithful to reality. It would let a robot learn a skill by imitation and afford a kind of consciousness. Asked why there are candles on the table, a third-generation robot might consult its simulation of house, owner and self to

reply that it put them there because its owner likes candlelit dinners and it likes to please its owner. Further queries would elicit more details about a simple inner mental life concerned only with concrete situations and people in its work area.

Fourth-generation universal robots with a humanlike 100 million MIPS will be able to abstract and generalize. They will result from melding powerful reasoning programs to third-generation machines. These reasoning programs will be the far more sophisticated descendants of today’s theorem provers and expert systems, which mimic human reasoning to make medical diagnoses, schedule routes, make financial decisions, configure computer systems, analyze seismic data to locate oil deposits and so on.

Properly educated, the resulting robots will become quite formidable. In fact, I am sure they will outperform us in any conceivable area of endeavor, intellectual or physical. Inevitably, such a development will lead to a fundamental restructuring of our society. Entire corporations will exist without any human employees or investors at all. Humans will play a pivotal role in formulating the intricate complex of laws that will govern corporate behavior. Ultimately, though, it is likely that our descendants will cease to work in the sense that we do now. They will probably occupy their days with a variety of social, recreational and artistic pursuits, not unlike today’s comfortable retirees or the wealthy leisure classes.

The path I’ve outlined roughly recapitulates the evolution of human intelligence—but 10 million times more rapidly. It suggests that robot intelligence will surpass our own well before 2050. In that case, mass-produced, fully educated robot scientists working diligently, cheaply, rapidly and increasingly effectively will ensure that most of what science knows in 2050 will have been discovered by our artificial progeny!

## THE AUTHOR

HANS MORAVEC is a principal research scientist at the Robotics Institute at Carnegie Mellon University. Over the past 40 years he has worked on eight mobile robots, the first of which—an assemblage of tin cans, batteries, lights and a motor—he constructed at age 10. His current work focuses on enabling robots to determine their position and to navigate by a three-dimensional awareness of their surroundings.

## MORE TO EXPLORE

VEHICLES: EXPERIMENTS IN SYNTHETIC PSYCHOLOGY. Valentino Braitenberg. MIT Press, 1984.  
AI: THE TUMULTUOUS HISTORY OF THE SEARCH FOR ARTIFICIAL INTELLIGENCE. Daniel Crevier. Basic Books, 1993.  
ROBOT: MERE MACHINE TO TRANSCENDENT MIND. Hans Moravec. Oxford University Press, 1999.