Research News

Fast, Cheap, and Out of Control

The visionaries at the MIT Insect Lab are out to revolutionize artificial intelligence, reform the space program, and mow your lawn with swarms of microminiaturized "gnat robots"

Cambridge, Massachusetts

GENGHIS DOESN'T LOOK LIKE A REVOLUtionary. Skittering across the floor and gamely climbing over books that have been stacked in its path, the little six-legged robot actually does a pretty fair imitation of a foot-long cockroach.

Nor does Squirt seem especially subversive. Hiding under chairs and in other dark places, and occasionally venturing out on its tiny wheels to investigate a handclap, the 1-inch mite seems more like—well, a bug.

But in the world of artificial intelligence and robotics research, Genghis, Squirt, and all the other electronic beasties that have lately been creeping around the ninth floor of the Massachusetts Institute of Technology's artificial intelligence (AI) laboratory are radical indeed. Their bodies push the state of the art in robot miniaturization, holding out the promise of 3-kilogram Mars rovers and tiny "gnat" robots, which might be deployed in swarms to gather military intelligence, or assemble fiber optics networks, or assist in delicate neurosurgery. Their microprocessor brains challenge some fundamental assumptions about the nature of reasoning and intelligence, producing startlingly life-like behavior with simple stimulus-response reflexes—and virtually none of the elaborate symbol processing used in expert systems and other traditional AI programs. Their performance has made other robot researchers nearly unanimous in using words such as "intriguing," "exciting," and "impressive."

And their antics, not incidentally, are a lot of fun to watch. "It's a real problem," laments graduate student Colin Angle, creator of Genghis and one of the two dozen students and postdocs making up what is informally known as the Insect Lab. "You show these robots to people, like grant examiners, and they say, 'Gee, that's—*cute*.' It makes it hard to realize just how much they can do."

Of course, given the ample supply of cockroaches in the world already, a visitor could also be forgiven for asking why anyone would want to build more.

"I'm *not* trying to build insects," retorts Rodney Brooks, the Australian-born AI researcher who is founder and guru of the group. "But I am trying to build robots that can get around in the real world and do useful work. And I've been inspired by the ability of insects to do a lot in the world without much neural circuitry."

That inspiration first struck about 5 years ago, says Brooks. He had just joined the faculty at the MIT and was looking around for a research project different from everyone else's, when he hit upon mobile robots. "The existing mobile robots would just sit there, computing for hours on end before they would make a move," he says. (He exaggerates only slightly: the Autonomous Land Vehicle, a state-of-the-art wheeled robot developed in the mid-1980s for the Defense Advanced Research Agency's Strategic Computing program, was the size of a van and was barely able to roll along a roadway at 20 kilometers per hour—and then only if the road were perfectly smooth and had perfectly even illumination.) Yet, Brooks thought, look what insects can do with hardly any brains at all. Clearly, something was wrong.

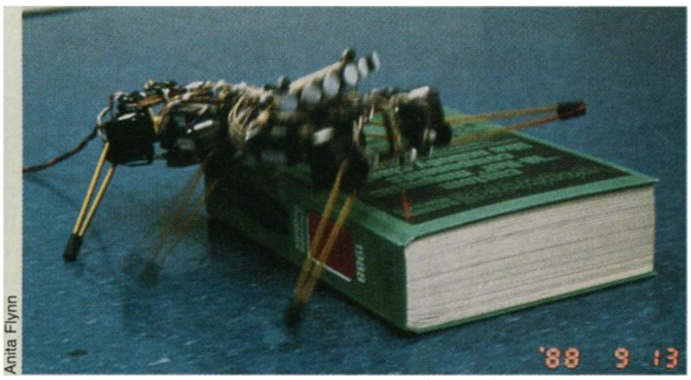

Anita Flynn

Genghis on the move. *With a brain based entirely on simple reflexes, the six-legged robot manages some surprisingly life-like behavior.*

The problem, he decided, lay with the conventional approach to robot control, which is derived from the standard AI model of human cognition first articulated in the mid-1950s by Allen Newell and Herbert Simon of Carnegie-Mellon University. A robot made according to that model first had to process the data its sensors picked up from the environment so it could identify objects in its field of view. Next, it had to construct some kind of internal data structure to represent the scene as a whole; then reason about that structure to construct a plan for accomplishing its goals; and then figure out how to execute those plans as a specific sequence of motor commands. Only after all that was accomplished could the robot actually *do* anything.

In the abstract, each of these steps seems essential. But in practice they seemed to represent an intolerable computational bottleneck. So Brooks decided simply to eliminate the bottleneck by eliminating cognition. His robots would neither reason, nor plan, nor make internal models of the world. Instead, they would have brains organized around coherent behaviors such as "avoid obstacles," "wander around," or "explore."

Each behavior, in turn, would be programmed as a kind of reflex, a direct link between perception and action: "Whenever *this* pattern of sensor readings occurs, fire *that* set of motor commands." If a wheeled robot were equipped with sonar, for example, and if the sonar readings suddenly showed that something was looming up on the right as it rolled along, the "avoid obstacle" behavior would kick in and tell the wheels to swerve to the left.

However, Brooks strongly emphasized that the simplicity of these individual behaviors would *not* necessarily imply simple-minded behavior for the robot as a whole. The separate circuits would be coupled in a rich web of interactions, with A triggering B, B suppressing C and enhancing the effect of D, and so on down the line. Moreover, each behavior circuit would constantly be altering its responses according to sensor input from an ever-changing environment. Such a robot could behave in exceedingly complex and surprising ways, he maintained.

Brooks called this arrangement of interacting perception-action links "the subsumption architecture" and argued that it was actually a much better model of human cognition than the Newell-Simon view. When you consider what a vast amount of sensory-motor coordination is involved in such a simple-seeming act as walking down the sidewalk, you have to wonder how "self-aware" we really are, he says. "My hypothesis is that large chunks of us are like this [directly linking perception with action]. There's only some thin veneer on top that rationalizes what goes on below."

Be that as it may, Brooks also argued that robots organized along the lines he proposed would have some major practical advantages. Robustness, for example: even if a sensor or a motor or a behavior circuit were to break down completely, the dense network of interconnections would still allow the robot to do *something* useful. Furthermore, the individual behavior patterns would be so simple that they could be etched directly into the robot's microcircuitry, thus allowing for extremely fast response.

In 1986, Brooks and his students demonstrated the feasibility of his approach with their first and simplest robot, Allen, which looked a bit like a 2-foot-high footstool on wheels. Allen had three behaviors implemented in a layered fashion. The lowest layer was "avoid obstacles": Allen would happily sit in the middle of a room until someone approached, and then scurry away, using its sonar distance sensors to avoid collisions as it went.

The robot's second layer varied according to who was programming it, but one version was "follow walls": this layer led Allen to do just that, while the previous layer forced it to maintain a safe distance from the wall. A third layer, "find doors," would cause Allen to swerve toward any gap it found in the wall while "avoid obstacles" kept it away from the doorjamb. As a result, without having any internal data structure that would allow it to recognize the concept of "door," Allen was able to roll right through the middle every time.

A more ambitious demonstration was another wheeled robot named Herbert, which

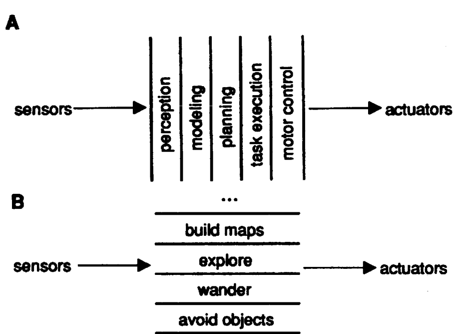

A better way? *The standard model calls for cognition to be sequential (A), but Brooks wants behaviors to work in parallel (B).*

Robot rally. *Members of the MIT Insect Lab—including Brooks, in the eyeglasses and white shirt—pose for a family portrait with a few of their creations.*

had a much larger battery of sensors and behavior circuits—and a gripper arm. In the end, Herbert was able to roll into offices and steal empty soda cans from people's desks.

And so it went, with robots like Seymour, Tom, and Jerry. The robotics community, which was still thinking in terms of robots that would need to carry around a Cray supercomputer to do much of anything, was impressed. However, as Brooks and his crew were well aware, Allen, Herbert, and their brethren all suffered the same drawback as any wheeled robot: they were confined to moving across smooth floors or, at best, a roadway. If robots were ever going to match the ability of living creatures to navigate across rugged, unpredictable terrain, they were going to have to have legs. And legged locomotion was something that conventional roboticists had so far found exceedingly difficult to achieve.

"They were worried about stability," says Brooks—making sure that the robot wouldn't fall over and break something. "But it occurred to me, why worry about that at all?" Falling down would certainly be a serious matter if your robot were the size of an elephant. But not if it were little. "Insects fall down all the time," he says, "but they still manage to get around okay."

Brooks therefore resolved to build a small, insect-like robot based on his subsumption architecture. "But when I got to thinking about it," he says, "I realized that the number of actuators you would need to control all those legs would be much, much higher than in the other robots. How would you coordinate them all? And then I realized that you didn't really *have* to coordinate them"—at least, not at the level of individual motors and sensors. With the subsumption architecture he could assign whole clusters of behavior circuits to each leg to work out the detailed motions of that limb independently. He could then get by with just a handful of centralized circuits that would broadcast commands for high-level behaviors such as "walk" or "turn" without worrying about the details.

This philosophy was embodied in Genghis, a tour de force of decentralized control. Built in 1988 by Angle, who was then an undergraduate, and programmed by Brooks, the six-legged robot was completed in just 8 weeks. As promised, most of Genghis' 57 behavior circuits are dedicated to local activities in the legs and whiskers. Only five of those circuits are required for overall coordination: two for walking, one for steering, and two that cause Genghis to track people with its infrared eyes.

The payoff is that Genghis in action seems eerily life-like. When it is first powered up, the robot's "standup" behavior pulls all six legs under the body and causes it to rise. Then the "walk" behavior sends it forward, with the legs alternating in much the same gait that real insects use. As Genghis starts to clamber over an obstacle such as a book, "force balancing" behaviors start pulling various legs in and out to keep its body from tilting too far to one side—even as the legs continue to move the robot forward. And if one leg should catch as Genghis tries to lift it over the edge of the book, a "leg lifting" behavior causes it to back down and try again—only this time lifting the leg higher.

Angle is currently hard at work on an heir to Genghis: Attila, which he hopes will be the prototype for a six-legged miniature Mars rover. As traditionally planned by the National Aeronatics and Space Administration, a mission to return samples from Mars would require flying one or two huge mobile robots weighing about a ton and costing about a billion dollars apiece, Angle says. But the Insect Lab's idea is to pepper the planet's surface with about 30 little robots like Attila, each one carrying only one or two instruments and each one weighing only about 3 kilograms. "You'd get much greater coverage of the surface," Angle says. "And you wouldn't be confined to the safest landing point. If one of your robots failed, you'd still have 29 left."

Mission planners at NASA's Jet Propulsion Laboratory have expressed interest in the idea—although they haven't yet been able to follow it up with financial support. But the Insect Lab group sees the idea not just as a way of exploring Mars, says Angle, but as a whole new paradigm for robot space exploration. Referring to their robots' ability to operate without constant input from the ground, the Insect Lab motto is "Fast, Cheap, and Out of Control."

In any case, says Angle, Attila will be much better adapted to rugged terrain than Genghis is by virtue of having knee joints on each leg. If Attila should fall off a rock and land on its back, moreover, it could rotate its legs 180° in their hip sockets and stand up again. For the same reason, Attila's body is designed to allow cameras and other instruments to rotate back to the top.

In addition, says Angle, Attila's legs will carry most of the robots 150 sensors for detecting position, force, surface hardness, surface color, and the proximity of objects. "Legs are really important for sensing the environment," he says. Attila will also carry A-Eye, a camera system that will track moving objects. Angle hopes that the robot will be able to stand up by June.

Outside of the Insect Lab, meanwhile, other roboticists and AI researchers have been watching all this with a combination of interest, excitement, and skepticism.

On the one hand, Brooks is greeted with packed lecture halls wherever he speaks. "Nobody thought you could do such seemingly intelligent things with such simple mechanisms," says Charles E. Thorpe, who works on more conventional mobile robots as head of Carnegie-Mellon's NAVLAB. "Rod has made a big splash among robotics and AI people. They don't all agree with him—but they all pay attention."

"It's very interesting stuff," agrees Kurt Konolige of SRI International. "Rod's major point about modular behavior and levels of behavior is very different from what has gone before, and people have really taken to that. It's always been taken for granted that a robot needs to model [compute a mental representation of] its environment. But now people ask, Why? For precisely what tasks do you need to model the environment?"

On the other hand, even the friendliest observers have to wonder whether Brooks' robot insects can ever be anything *more* than insects. "The drawback is that his robots are only probabilistically intelligent," says CMU's Thorpe. That is, you can never be quite sure what they are going to do. "You could build a whole fleet of robot mice to clean up the crumbs in your house, and if they get 99% of the crumbs, that's okay. But you wouldn't want to build a robot chauffeur that way, because a 99% chance of getting to work safely is not okay." Furthermore, he says, because a Brooks-style robot has very little memory and no internal data structures, "you can't even tell it something as simple as 'Go through the fourth door on the left,' because it doesn't have any concept of 'door,' or 'four'."

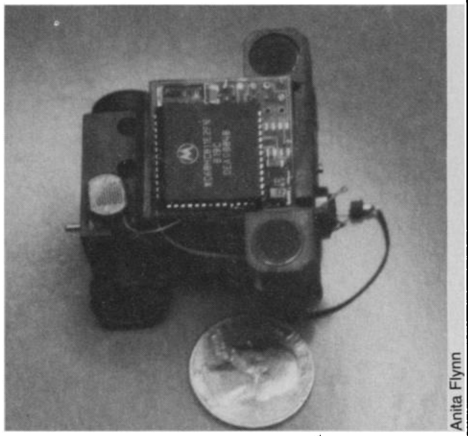

Anita Flynn

"World's largest 1-cubic-inch robot."
Squirt's brain and brawn fill 1.3 cubic inches.

Brooks, for his part, is a bit impatient with such criticism. Devising robots that can deal with such concepts is the standard approach, he says: "It's the way everybody else is working, and I don't see such great success that I should spend the rest of my life working that way, too." Yet he acknowledges that finding more sophisticated ways of controlling his robots has become a top research priority at the Insect Lab.

But Brooks and his crew have hardly given up on wild-eyed dreaming in the process. Consider their effort to push robotics along the path taken by microelectronics 20 years ago. "If you look at most mobile robots, says research scientist Anita Flynn, "you find that most of the bulk and the cost is in low-tech items—the motors, the chassis, the power supplies, and so forth." But with the subsumption architecture, she says, "you can just put some small integrated circuits on board, and the behaviors can be designed into the microcircuitry of the silicon." So why not try to make the motors and everything else as small as possible, too?

Enter Squirt, which Flynn refers to as "The world's largest 1-cubic-inch robot." (It actually measures 1.3 cubic inches.) Built in late 1988 as an exercise in reducing a robot to its bare essentials, Squirt is able to seek out dark corners and do its bug imitation with only 1300 bytes of computer code in its control system. However, most of Squirt's small bulk is still taken up by its motor and power supply. Is it possible to do even better?

Maybe, Flynn says. About the time the MIT group built Squirt, they heard about people trying to put motors on a chip by etching microscopic rotors right into the silicon surface. "So we said, Cool! We can build robots on a chip." Thus arose the concept of millimeter-sized "gnat robots," whose motors, brains, photovoltaic power supplies, and light sensors would all be on the same piece of silicon. These gnats are still a distant dream, admits Flynn, not least because the existing silicon micromotors are not nearly powerful enough. But the group is investigating several approaches to making them more powerful.

And what use would gnat robots be? Well, says Flynn, they could serve as miniature construction workers, aligning optical fibers or bonding wires to chips—or for that matter, acting as a surgeon's remotely operated hands for delicate eye surgery. They could serve as janitors and maintenance workers in such hard-to-reach places as space telescopes and planetary probes, keeping the optics clean and the instruments in tip-top condition. They could serve as tiny, autonomous sensors for military intelligence and space exploration.

Indeed, as she and Brooks point out in a playful research memo entitled "Twilight Zones and Cornerstones," the applications of gnat robots are limited only by the imagination. If you were a Navy captain, for example, why not just toss a few million gnat robots down the side of your ship to munch away at the barnacles one by one? For that matter, why not let swarms of gnat robots patrol your garden for pests, or trim your grass blade by blade?

Admittedly, they say, it's going to be quite a while before gnat robots are developed to that point. But they maintain that the possibilities are worth thinking about even so—because having swarms of tiny robots will fundamentally change the way we conceive of what robots can do.

■ M. MITCHELL WALDROP