

From Simple Rules, Complex Behavior

FOCUS ON INTELLIGENT CONTROLS

By layering simple behaviors, engineers are creating autonomous robots that can interact with the world around them.

By Alan S. Brown, Associate Editor

Robots have overrun Hanover Fair. Exhibitors at this, the world's largest industrial exposition, held each year in Germany, display hundreds of robotic arms, actuators, sensor systems, and feedback mechanisms. The fair even launched a "Robotation Academy" with Volkswagen Coaching, the automaker's consulting arm, advising attendees on robotic manufacturing.

More astonishing, though, are the autonomous robots. AirPenguins fly above the booth of German pneumatic automation supplier Festo. The RoboCup German Open pits competing teams of autonomous robots against one another in a soccer tournament. On the side of the playing field, mobile autonomous robots run in patterns and patrol the booths of a surprising number of companies.



Individual ants have simple behaviors but perform complex tasks by working together.

The growing presence of autonomous robots at Hanover echoes the interest seen worldwide. In the United States, for instance, autonomous robots range in sophistication from simple Roomba household vacuum cleaners to the complex vehicles that competed in Defense Advanced Research Projects Agency's Urban Challenge. In Japan, consumer electronic firms see their programmable toy robots as steppingstones to fully automated "personal assistants."

No one need worry about an independent-minded robot taking his job soon. In a preliminary RoboCup match, one team's goalies fell down as soon as they tried to move. Vent downdrafts occasionally push AirPenguins into Festo's booth. During the Urban Challenge, plenty of vehicles plowed through gates or crashed into walls or other vehicles.

Progress is palpable even as developers wrestle with the most vexing of control questions: How do you create a robot intelligent enough to operate without direct guidance? This is very different from programming an automated control system in a factory, where every piece of equipment and every

possible interaction is understood ahead of time. A truly autonomous robot must be able to interact with a constantly changing—and therefore unknown—environment.



Festo's AirPenguins fly and interact as a flock. These robots are featured in this month's online exclusive at www.memagazine.org.

Once, theorists thought autonomous robots needed artificial intelligence, the ability to assess the environment and make informed judgments before acting. Today, developers are putting less emphasis on this type of “thinking,” which is far too slow for real-time decision making.

Instead, engineers are looking for clues in the behavior of social insects. Ants have attracted the most interest: While an individual ant has only a limited range of simple behaviors, the way it interacts with other ants produces responses that make the nest appear intelligent.

“When simple behaviors work together, they can create what appears to be complex behaviors to a naive observer,” noted Bryan Adams. Adams is a principal investigator at iRobot, developer of the Roomba, the world’s most widely owned robot—and one inspired by insect behavior.

LEARNING FROM ANTS

To understand what Adams means by “complex behaviors,” consider how some species of ants find the shortest route to food. When a randomly foraging ant finds food, it grabs a piece and wanders back to the nest, leaving a trail of pheromones behind. Other ants set out from the nest, following that pheromone trail to seek provisions. At first, they follow it readily because the scent is recent and strong. Further away, though, pheromones have begun to evaporate, and the ants begin to wander from the trail. The ants that find the most direct route to the food and back leave the strongest scent, and their trail is the easiest to follow.

At the highest level, this looks like rational behavior. Yet it derives from very simple instincts: walk randomly until you find food, bring the food back to the nest, follow the strongest scent back to food, and repeat. Those rudimentary instincts, combined (or layered), produce complex, problem-solving behaviors.

Rodney Brooks, a researcher at the Massachusetts Institute of Technology and a co-founder of iRobot, based in Bedford, Mass., paid close attention to biological systems as potential models for robot behavior. “It was the launching point for his robots,” said Adams, who spent seven years in Brooks’ laboratory. “Up until then, researchers thought robotic intelligence meant building complex models of the world that the robot would act upon. But if you look at animals, they’re not doing that type of modeling.”



Roomba, a robot vacuum cleaner, solves complex problems by using a few simple rules.

BlueBotics' Gilbert maps its environment using wheel encoders and uses sensors to avoid moving obstacles.

In 1989, Brooks developed Genghis, a six-legged robot with compound eyes, to embody his layered approach. It was an immediate success and appeared on the cover of *Popular Science*. “The robot had a behavior that it engaged in by default, it walked,” Adams explained. “As it received inputs from the world, different behaviors took over from the default.”

Genghis could, in fact, engage in several behaviors, including chase, stand up, and avoid or walk over obstacles. Those behaviors were prompted by the environment, the same way ant behavior is cued by the scent of food or pheromones. The resulting interactions mimicked the behavior of insects. Indeed, Brooks has described Genghis as having a “wasp-like” personality.

The same approach now guides the Roomba, a disk on wheels that scurries around furniture and backs away from walls while randomly vacuuming rooms. Roomba’s instruction set looks something like: wander and vacuum, go left or right upon hitting an object, back up or spiral when caught in a corner, and find the docking station to recharge when low on power.

Is that all there is to a Roomba? “We obviously don’t give out all the details of our algorithms,” Adams said.

He notes, however, that many of the principles behind Genghis also govern Roomba. “You have a behavior that the robot engages in by default,” Adams said. “A default behavior might be to drive. If the input is that the front bumper is depressed, it might back up. If you start layering those behaviors, you get complex behavior.”

Brian Cusack, an adjunct professor of mechanical engineering at Cooper Union in New York City, makes a similar argument. Cusack’s introductory robotics course ends with a wheeled robot sumo wrestling competition. “There are only about 50 lines of code in a complicated sumo robot,” he explained. “Move around. Do I see the edge of the ring? If yes, turn away. Am I touching the other robot? If yes, push with all my might.

“We’ve found that artificial intelligence is hard,” Cusack continued. “Algorithms like those used to play chess are incredibly difficult to implement in a robot that needs to make decisions in real time.

But intelligent-looking behavior, behavior that mimics intelligence, is not hard. You don't need a supercomputer, so you can do it for \$200 rather than \$200,000."

SURPRISING BEHAVIOR

Watch a Roomba go about its chores and it becomes obvious that it makes many less than optimal (or even stupid) decisions. Which raises the question, can robots learn to behave, pardon the expression, more intelligently? And can they do it with simple commands that enable them to function in the real world?

Stefan Wrobel thinks so. Wrobel, director of the Fraunhofer Institute for Intelligent Analysis and Information Systems near Bonn, is active in RoboCup competitions.

"In robotics," Wrobel explained while scores of students put the finishing touches on their RoboCup contenders, "behaviors are the building blocks of a goal you want to achieve. Robots that learn and adapt to the environment go beyond manifesting behavior. They learn while performing tasks."

At its least complicated level, a robot could use simple behaviors to confirm its model of the environment. "When searching for power, it would learn whether its world model is or is not correct. If there is no power source in the northwest corner, it could remember that," Wrobel said.

On a more sophisticated level, a robot could assign priorities to its own behaviors, such as when to move or what to grab. According to Wrobel, designers give robots goals, and when the robot does something right, it receives a "reward." This is exactly how animal trainers use food rewards, though robots respond better to a number in an algorithm than to a piece of kibble.

"The robot doesn't know exactly what it did right," Wrobel said. "Most actions are complex interactions. But the actions that get you the most credit at the end of the day are the ones the robot is most likely to repeat."

Unfortunately, actions do not take place in a vacuum. They happen in specific states, or operating parameters, such as location, sensor input, and battery charge. A single 8-bit sensor can record 28 or 256 different states, and a robot may have a dozen or more sensors. It takes a lot of computing power for the robot to decode what it did right.



One approach, and Wrobel's latest line of research, is to attack the problem holistically. "When I think of a robot's representation of the world, I think of measuring and storing information. But it's obvious that's not how the human brain works. The brain isn't digital. A robot's sampling rate completely loses the dynamicism of the world. It is more power-intensive to recognize a wave from samples than it is to see the wave. So we're looking at ways use analog chips to create complexity in different systems."

For all the work in getting robots to discover their environment, providing a little human guidance is still important. That was exactly what Grégoire Terrien, an engineering manager at BlueBotics, headquartered in Lausanne, Switzerland, did with Gilbert.

Gilbert is an autonomous robot that patrolled the perimeter of the BlueBotics booth at Hanover Fair, along the side of the RoboCup playing fields. Before the fair opened, Terrien walked Gilbert around the booth to help it map the space, a process known as localization.

Nesbot is an autonomous robot that delivers coffee to the desks of Nestle workers who place orders over the Internet.

There are many ways to localize a robot. Terrien tilted his head towards three robots running in a tight circle across the hall. "Some of those robots are using sensors to follow a line, but it's not very flexible," he said. Then he pointed to three towers with spinning beacons on top. "The other robots are triangulating their location from the beacons, but in a factory, a warehouse, or a home, those beacons could be blocked. To be truly autonomous, a robot has to adapt to the environment, not the other way around."

To acclimate Gilbert to his environment (a booth, since Hanover Fair does not allow robots to roam free) Terrien walked the robot around, enabling Gilbert to map the booth's walls. "Gilbert also recorded data from encoders on its wheels that are accurate within 1 centimeter," Terrien said. Employing both the map and encoder data, Gilbert can calculate its position. It also uses a laser scanner to find objects it must avoid, such as chairs and people.

Does the robot really know where it is? Robots are sometimes switched off, or picked up and moved somewhere else (the so-called kidnapped robot problem.) Their sensors may make small errors that accumulate over time. A robot may think it knows its position, but it is sometimes mistaken.

BlueBotics, like other autonomous robot designers, takes a probabilistic approach to localization. Terrien pulls out a mark-covered map of where the robot thinks it is. Most of the marks cluster within a small ellipse, but there are scattered markings everywhere else. Gilbert constantly recalibrates that ellipse, which represents its location with a 95 percent probability. But if it becomes lost and can't find a landmark to reorient itself, it will stop.

The approach looks impressive when applied to Nesbot, a robot BlueBotics built for Nestlé Nespresso, a coffee service provider. Nesbot is essentially an autonomous coffee machine. It takes orders from a detachable handheld PC or from the Internet, then navigates Nespresso's offices to deliver the beverage to the employee's desk.

TEAMWORK

While Terrien was talking, RoboCup matches were under way. The long-term goal of the competition is to develop a team of robots capable of winning a match against human opponents by 2050. But if the play of the humanoid division is any indication, that timeline may be optimistic.

The University of Bonn's NimbRo team takes the field against a first-time RoboCup competitor. At the mid-line, two NimbRo robots march in place while the other team's robots stand still.

As they stomp from foot to foot, the NimbRo bots look intimidating, like veterans trying to scare a squad of rookies. Is that why they do it? "No," said Sven Behnke, head of Bonn's Autonomous Intelligent Systems Group and NimbRo's team advisor. "We do it because it's hard for humanoids to start and stop."

The NimbRo robots soon prove Behnke's point. As the referee blows the whistle, the NimbRo robots clomp off after the ball while two of their opponents immediately fall over.

Humanoid robots show themselves to be limited in other ways. They must stop before kicking the ball. They cannot pass effectively. Few can lift themselves off the ground. They are also painfully slow.

Wheeled robots have more game. They speed around the field, bumping the ball to one another and then finally bouncing it towards the goal. There is no hesitation. Everything happens in seconds.

The reason the robots can work together and still make rapid decisions is because they are organized hierarchically. "Robots have to react to different things as individual robots and as team players," Behnke said.



The RoboCup is a soccer competition for autonomous robots. To win, robots must play as both individuals and team members.

Each of Behnke's robots have four levels of control: the entire team, an individual robot, isolated body parts, and single joints. At the team level, the robots are focused on plans for the immediate future, while at the individual level, they react more to their environment. "The mechanisms are simple all around, but their interactions with the environment creates some complexity," Behnke said.

For example, each of Behnke's robots fixes its position by comparing a map of the green field with such visual cues as the white sidelines and goal. The robot moves toward the ball by default; as it moves, it calculates the ball's trajectory and plots a curved course that would place it behind the ball to position itself for the kick. Once the robot makes these calculations, the general instructions are sent down to the robot's arms and legs, which send even more detailed orders to the actuators in its joints.

"The robot would be helpless if it tried to control everything through one central controller," said Behnke. "Humans are the same way. The brain decides on a plan and the spinal cord communicates it through local reflexes to the rest of the body. It's simple and reactive."

Simple and reactive are good, but consistent is even better. That begins with physical robustness. For every minute of RoboCup play by teams such as NimbRo, which went on to win the championship, there are untold hours of preparation and recalibration. At tables around the playing field, scores of students are repairing, upgrading, or salvaging robots.

Ericson Mar, who teaches robotics at Cooper Union, ends his class with a robot tank battle. Students outfit their robots with all kinds of sensors, from electronic compasses and wheel encoders to sonar and infrared range finders. Even so, according to Mar, "Winning boils down to consistency." The robots that win are the ones that can do the task over and over again."

Yet the challenge of developing physically robust robots pales in comparison with the obstacles posed by achieving consistent behavior. Just like ants trying to follow a path of fading pheromones, autonomous robots learn from random behavior. As robots learn and adapt, they often show unpredictable behavior that no one expected.

behaviors to build complex interactions—that enables them to create robots that do useful and interesting things in real time. Only they can't do them consistently enough for most real-world applications.

Many challenges remain. Autonomous robots need better localization and mapping. They must learn faster and more efficiently. They should work together better as a team. They need to be more robust.

Above all, they need to get smarter. An ant works randomly because it has no choice. It stumbles around until it creates the shortest path to food. Humans can see the whole picture and pick out that path immediately. Now that engineers have developed robots to the level of ants, can they lift robotic sophistication to the next evolutionary step?