Visionary Arms

The less a robot knows at first, the better it may fare in uncertain settings

By IVAN AMATO

Moving a piano demands precise planning. Though Schwarzenegger types might enjoy hauling massive objects around unanticipated obstacles, the scrawnier masses strive to minimize improvisation during this sort of vein-busting activity. Before budging the leaden instrument even a smidgen, most would-be haulers try to map out the move in detail. Wielding tape measures, they carefully note the dimensions of the piano, doorways and other potential obstacles as they plot the shortest possible path between the piano's present location and its future resting place.

Most of today's industrial robots – which carry out such repetitious jobs as welding together the same car parts day in and day out – operate like piano movers. The computers that control them know everything ahead of time. They know the exact distance between the car parts and the welding tool attached to the robot's arm. They know precisely how much the arm needs to move to get that welder to the right spot. And they know just how much time a good weld takes.

This know-it-all approach works quite well for robots that perform utterly predictable operations, notes electrical engineer Vladimir Lumelsky of Yale University. But Lumelsky suspects that a learn-as-you-go approach might work better in other situations, and the robotic system he is now developing may prove him correct. "What we want our robot to do is operate in an insufficiently known environment," he says.

For example, a hospital administrator might want a food-delivering robot that can travel through corridors without bumping into boxes on the floor or clusters of people that weren't there the day before. A director of a day-care center might want a robot that can clean up every day after a tribe of preschoolers. Or an apple grower may long for a robot that can collect fruits of varying sizes from random locations without bruising the harvest.

The difference between robotic welders and hospital-worker robots rests within their computer-controlled approaches to the world. With a standard model called the Piano Mover, a preinformed computer swiftly guides the robot along the most efficient pathways. But if anything changes in this robot's limited environment, the computer becomes ignorant and the result is automated ineptitude. If an errant car chassis arrives just inches out of line, for instance, a robotic arm operating according to a rigid Piano Mover model might instead weld a spot of air.

Lumelsky envisions more flexible robots that operate according to a different mathematical view of the world, which he describes as the South Pole Search model. This model trades the foreknowledge of the Piano Mover – and the ultimate efficiency in movement that results from such knowledge – for the ability to adapt robotic movements to an unpredictable and changing environment.

Like a monarch sending an explorer off to find and claim the South Pole, the robot's computer-controller always "knows" the robot's purpose, or "global goal," but it doesn't know what specific geological obstacles the robot may encounter on its way. To provide this vital intelligence, on-robot sensors detect immediate obstacles and continuously inform the computer, which then uses motion-planning programs to guide the robot around those local obstacles.

Since the South Pole Search model assumes only minimal foreknowledge about the environment, says Lumelsky, a robotic Robert Scott should be able to start anywhere on the globe and get to the South Pole via a reasonable – though somewhat roundabout – path.

Building this kind of flexibility into robots might elevate the mechanical drones of today's assembly lines to the kinds of adaptable, labor-saving tools imagined by the early visionaries of robotic technology. To hasten that day, Lumelsky is studying how humans repeatedly succeed at the same "global goals" even though the information available to them may change each time.Truck drivers, for example, maneuver through countless city and highway settings, facing traffic patterns that change from one moment to the next. They avoid accidents by using sensory cues to guide their decisions to turn, brake or accelerate.

Driving would prove a frightfully hazardous business for robots that operate according to the Piano Mover model, which would assume exactly the same driving conditions in every setting. But a robot driver programmed with the South Pole Search model could adapt to new settings, even though it wouldn't discover what lurked around the corner until its sensors got close enough to "see" it.

Lumelsky has studded a robot arm with about 500 infrared sensors, which give it the look of a cactus mounted on a dentist's X-ray machine. The arm's surface serves as a "sensitive skin," a veritable compound eye that enables it to sense objects in its surroundings. Although some engineers are investigating sensitive skins that respond to tactile stimulation, Lumelsky argues that heavy robotic arms need to "see" objects from a distance before they "feel" them.

"NASA says that every time one of its robotic arms bangs into something, it costs them half a million dollars to fix things," Lumelsky notes. On a space station, a misguided tap from a repair robot might prove catastrophic to budgets, equipment and even lives.

So let the arm “see,” says Lumelsky. Each infrared sensor on his robotic arm consists of an emitter, which sends a beam of energy into space, and a detector, which picks up reflections from any object within about 6 inches of the arm. This multiple-sensor tactic, resembling radar and called “proximity sensing,” fills in the blind spots that would limit any practical arrangement of cameras, Lumelsky points out.

“You want to produce this aura around the arm with no holes or gaps, so that nothing could penetrate undetected,” he says. The currently restricted 6-inch “field of vision,” makes the arm slow. It takes 30 seconds or so to move about 5 feet and could not, for example, get out of the way of a lobbed basketball.

But so long as these robots perform tasks that don’t require the calculation of an optimal solution, which takes tremendous amounts of computer time, their speed will increase with further development, Lumelsky says. To that end, managers at the Japanese company Hitachi have already entered into a collaboration with Lumelsky. Hitachi aims to build the next generation of a more rugged proximity-sensing skin – studded with thousands of sensors – enabling the next generation of Lumelsky’s robots to see farther with higher resolution and precision.

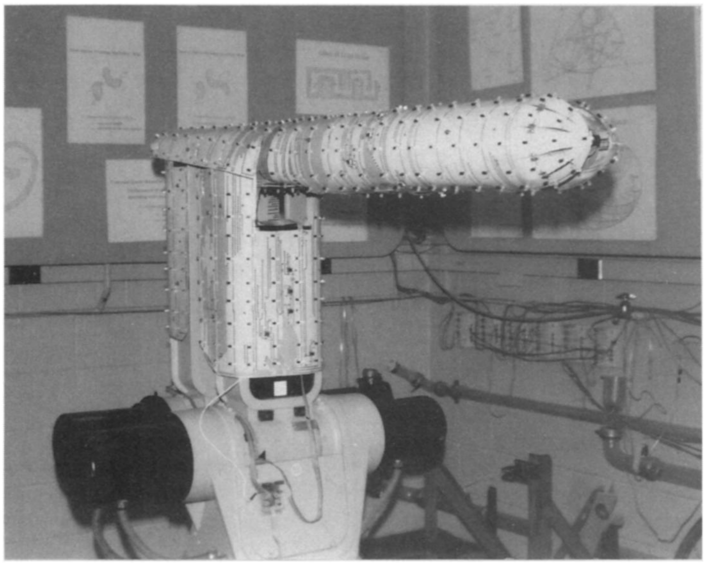

Studded with about 500 infrared sensors connected to several computers (not shown), this robotic arm can ‘see’ and maneuver around objects.

V. Lumelsky

But first steps first. In the present system, each of three computer processors controls one of three “skin” sections, each of which sports about 170 individual sensors. Raw signals from a sensor get digitized and preprocessed before they enter the motion planning programs based on the South Pole Search model. When sensors indicate that something is near, the associated processor injects this information into a mathematical algorithm that generates a small step toward the robot’s global goal. The arm slowly bobs and weaves as it tries to avoid the obstruction, each movement occurring in smooth succession within its half-foot “aura of visibility.”

Raw sensory signals may indicate obstacles, but they don’t automatically choose among the many pathways the arm could take to get by the object. That’s where a variety of decision algorithms come into service. Several years ago, Lumelsky and a graduate student discovered that the physical constraints inherent in any arm design help to simplify the mathematical description – called a configuration space – of the spatial relationship between the robot arm, its reach and objects within its reach.

To describe the mathematical problem, Lumelsky uses the analogy of a bug crawling on the surface of a doughnut or, in geometric terms, a torus. For a robotic arm with two joints, a torus can represent all possible locations of items within the arm’s reach. In order to negotiate around obstacles, “the arm imagines itself as the little bug on a torus,” Lumelsky explains. The motion planning program, which incorporates the robot’s global goal, traces a line from one point on the torus to another and excludes paths containing points corresponding to objects detected by the robot’s sensors. Since many paths can connect any pair of endpoints, other criteria – such as requiring that the route not exceed a stipulated length – enable the processor to settle on a specific sensible move without spending lots of extra time searching for an unnecessarily more efficient step.

By building a robot that can achieve certain goals through any of several sufficient moves rather than requiring a single optimal move, Lumelsky frees the computer-controller from carrying out unacceptably time-consuming calculations. At the same time, he says, such a robot may behave in more lifelike ways. Living creatures rarely move or behave optimally, yet their continued survival shows that their choices are sufficient to achieve their goals.

At times Lumelsky’s robot does show some of the less-touted features of intelligence, such as indecision. For instance, if he puts an object on each side of the arm, just barely within its “aura of sensitivity,” the arm darts back and forth in attempting to avoid both obstacles. “It behaves like an animal, like a cornered cat,” Lumelsky says. “It really gives you a strange feeling.”

Lumelsky hopes his efforts will lead to robots that make better companions for both their human and robotic workmates. In one important type of human-robot interaction, called teleoperation, a human operator handles a small “master arm,” which controls the movements of a bigger “slave” arm located some distance away.

“People have trouble in teleoperation,” Lumelsky says, because most human teleoperators find it difficult to think geometrically, imagining themselves as bugs on a torus and then translating that perspective into specific manipulations of the master arm in normal three-dimensional space. Such people perform poorly in computer simulation studies in which they are asked to plan the motion of a two-dimensional arm moving among simple planar obstacles. Even when they have complete information about the obstacles in the simulated environment, their performance speeds up only slightly. Robots with less information outperform humans in such tasks, Lumelsky says.

He interprets these results as a prescription for human-robot teleoperation teams. “In our system, we have an autopilot that takes over at the places where a human is not good at it,” Lumelsky says. When near an obstacle, the robot would shift to computer control, becoming an autonomous navigator. Its sensors and motion-planning algorithms enable it to negotiate around objects that a slower human teleoperator could not avoid. The human operator, however, controls the overall motion by setting intermediate and global goals. The autopilot mode is so smoothly integrated that the human operator can’t tell when it takes over, Lumelsky says.

Ultimately, he hopes his efforts will extend beyond human-robot teamwork, leading to autonomous robots that can work together without witlessly beating each other into piles of parts. Lacking a “boss” computer to orchestrate their motions, two or more independent robots working close together wouldn’t last long. But with sensitive skins and the flexible South Pole Search model, they might cooperate successfully in countless tasks ranging from car assembly to space station repair. And, Lumelsky quips, they might even enjoy a bout of fencing now and then. ☐