Cognitive Sciences Explored in Chicago

The beauty of being a cognitive scientist is that, while trying to learn how people think, or to build computers that mimic the human brain for use in industrial robots, you get to sample a rich broth of scientific pursuits: linguistics, anthropology, and philosophy as well as psychology and, of course, computer science. Experts in—and "tasters" of—all of these pursuits mingled at the University of Chicago from 7 to 10 August during the thirteenth annual Conference of the Cognitive Science Society. Here are some of the highlights that drew 500 researchers to the meeting.

Computer Vision Moving Closer to Reality

The buzzword is "active vision." That's what robots need if they are to identify and target enemy missiles, harvest fruits and vegetables, serve as mechanical housemaids, or explore space or the sea depths. In any such application, the challenge for computer vision systems is to pick out specific objects from a wide and changing field of vision. But until recently, such active vision systems have existed largely on paper because of the enormous number-crunching required and the costs of the systems' various components.

That's why participants in a workshop devoted to the topic at the cognitive science conference were quite excited by the great strides being made in two critical areas: the development of algorithms that allow computers to select desired images from a visual melange and the miniaturization of the cameras that serve as the robots' "eyes." Where progress is still needed, says Pete Bonasso of the MITRE Corp. in McLean, Virginia, is in the development of compact, lightweight power sources, especially for robots that would move out on their own—in unmanned vehicles for exploring planetary surfaces, for example. For that reason, computer vision is more likely to be used first in applications, such as spotting speeders or red-light runners, where the system as a whole doesn't have to move.

When researchers first tried to build computer vision systems, they ran into trouble, says computer scientist Michael Swain of the University of Chicago, who chaired the workshop, because they designed their machines to "look" at everything in their visual fields. The algorithms the researchers wrote were unable to process the large amounts of information in the vision fields fast enough to allow the robots to respond in real time. So much of the recent progress with the algorithms came with the realization that the machines' vision can be more selective. "The new rules for the design of computer vision are that you don't have to know everything, everywhere," Swain says. "You need only information that meets the goals of a particular behavior."

Eric Schwartz

In a nutshell. *The miniature camera, measuring 0.3 x 0.3 x 0.4 inches, may one day help catch red-light runners.*

Take the system designed by Eric Schwartz of New York University and his colleagues. It uses a program that enables it to survey a wide field of vision, but focuses clearly only on objects in the center of the field. That, of course, could have its downside, so the system was also designed with "attention algorithms" that detect action at the periphery, thereby telling it where to look next.

Schwartz's group has also attacked the problem of camera miniaturization. At the workshop, the team unveiled a piece of nifty hardware: a miniature video camera, small enough to slip inside a pistachio shell, that could serve as the system's eye. In fact, Schwartz envisions his system in actual use—catching red-light runners. It will first read the lawbreaker's license plate and then the computer will write a ticket on the spot. Other potential applications of the camera include surveillance at automated cash machines and as a built-in target finder for guns.

But while these algorithm and camera improvements have helped reduce the size of the hardware needed to run robots, further reductions are needed if the hardware is to be used in more mobile robots. Consider the problem encountered by MITRE's researchers. They've just built a robot with a state-of-the-art active vision system that may eventually be used to help an unmanned space vehicle avoid potentially lethal obstacles. But their prototype robot was a clumsy, rough-moving device, burdened by 10 pounds of computers and about 100 pounds of batteries needed to power it. So the engineers still have their work cut out for them as they try to build mobile and all-seeing automatons.

Computer Learning Gets Mixed Grades

For about 2 decades, computers have been touted as tomorrow's helpful assistant to the teacher. Students of all ages could learn much better—or so the theory goes—with the aid of computers that could simulate such real-life experiences as flying an airplane or conducting an ecology experiment. They could get immediate feedback on how they were doing, and the simulations would be cheaper (and in the case of flying, safer) than the real thing.

But while such computer-assisted learning works quite well in some situations, usually with adult professionals such as pilots, it hasn't been completely successful so far with the inhabitants of kindergarten to grade 12, according to a panel at the cognitive science conference. The panel members, all from Apple Computer's labs in Cupertino, California, conceded that, in trials of their own software as well as that of other manufacturers, "big problems" had been encountered in constructing simulations that work well for youngsters.

The main problem seems to be that schoolchildren, many of whom cut their computer teeth on "Super Mario Brothers" and other video games, get seduced by the whiz-bang nature of the technology, says Apple's Wayne Grant. He cites as examples experiences with two schoolroom simulations adapted for use on Apple computers. One is a program that explores the predator-prey relations between ants and beetles and the other probes the environmental impact of dam construction. Instead of patiently changing variables, such as ant-colony size or the health of the ants, one at a time to see how each affects the beetles' ability to prey on the ants, the students changed many variables at once, or jumped from variable to variable without fully exploring any. As a result, they came up with quick, but often incomplete conclusions. "The students took a video game approach to the simulations," says Rick Borovoy, also of Apple.

Now that the Apple researchers have a better understanding of the hurdles they face, they are at once seeking help and proposing some solutions. "We know students are amused and motivated by simulations, but we are just beginning to learn how to make simulations work in an instructional setting," Apple computer scientist James Spohrer said at the meeting. And ina nod to his audience, he added: "We are looking to cognitive scientists to get ideas about how this can be done." Meanwhile, though, one potential solution under investigation at the company is the development of the computer equivalent of lab books to help students plot a logical course of inquiry through simulated experiments, instead of using video game tactics. ■ ANNE SIMON MOFFAT

Atoms Do the Two-Step on Crystal Dance Floors

Physicists have long tended to picture an atom moving on a surface as something like a loose marble on a layer of close-packed marbles, hopping from hollow to hollow on the surface without disturbing the atoms underneath. But that's only part of the picture. Instead of hopping over a surface, it seems, an atom is often more likely to dive into it, displacing another atom, which then pops upon the surface nearby. That's the message now coming from a merger of an old microscopic technique with a theoretical analysis by physicist Peter J. Feibelman of Sandia National Laboratories.

The wanderings of a single atom might not seem to count for much. But an accurate knowledge of how atoms migrate across the topmost atomic layer of a metal is vital for understanding—and hence controlling—many basic surface phenomena. Corrosion, crystal growth, and catalysis, to name just a few, are all carried out by atoms dancing on surfaces. A refined understanding of how their dance is choreographed is also important, remarks Feibelman, "because we think it can help guide people in making materials." For example, he conjectures, a clearer picture of the dynamic behavior of atoms on surfaces could guide materials scientists who are attempting to create superstrong metallic alloys by building up atoms-thick layers of different metals.

Early hints that atoms on surfaces might not be as simple as loose marbles actually came decades ago from images made by field ion microscopes (FIMs), devices that were revealing the positions of individual metal atoms on surfaces long before the scanning tunneling microscopes now used to map atomic topography. In an FIM, a strong electric field between a sharp probe and a crystalline metal surface ionizes atoms of helium, deliberately leaked into the sample chamber, wherever the electric field is greatest—at the peaks created by atoms resting atop the crystal plane. The helium ions act as a tracer, generating an image of the peaks. And while scanning tunneling microscopes tend to drift, an FIM can repeatedly image a single patch of surface to trace the fate of atoms there.

As far back as the late 1960s, surface scientist Gert Ehrlich of the University of Illinois at Urbana-Champaign and others were using successive FIM images to show that when atoms from a vapor of one metal land in the valleys between ridges of atoms on the surface of another metal, they often burrow into a neighboring ridge, pushing atoms of the surface metal into the next valley. Feibelman's analysis shows that such behavior may be common at metal surfaces. His model, which simulates the movement of aluminum atoms across an atomically flat aluminum layer, suggests that it takes less energy—as measured by the balance of bond-making and bond-breaking—for newly arrived atoms to move by displacing an atom in the topmost layer of the crystal to a nearby site than by hopping across the surface without digging in. When an atom hops from valley to valley,

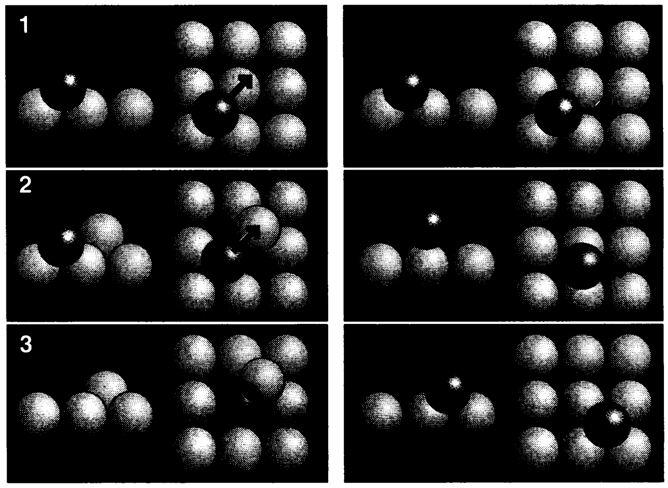

Trading places. *A metal atom often migrates across a crystal by displacing an underlying atom* (left) *rather than hopping.*

"you have to pay a toll for going over a bridge" that is higher than the cost of displacing another atom, Feibelman says.

In addition to the retrospective support from Ehrlich's work, Feibelman's model is gaining new support from FIM observations made by Feibelman's Sandia colleague, experimental physicist Gary L. Kellogg. Kellogg and others have observed platinum, iridium, and nickel atoms migrating across crystals of those metals by Feibelman's predicted exchange mode.

Feibelman's model doesn't rule out the hopping motion of the traditional picture; instead it predicts that the dominant migration pattern will vary depending on such factors as temperature, the kind of metal, and whether the wayfaring atoms move alone or in clusters. Indeed, Kellogg and his colleagues found—in keeping with Feibelman's predictions—that clusters of two platinum atoms actually migrate faster by the exchange mode than single platinum atoms do; clusters of three platinum atoms, on the other hand, migrate both by exchange and by hopping.

Ehrlich welcomes the new interest in his old observations. Together with new studies by him and others, it's leading to a much richer view of crystal growth on its atomic level, he says.

But so far the sharper picture of surface choreography applies only to metal atoms moving on metal surfaces. According to Steven George of Stanford University, who studies the surface diffusion of molecules such as carbon dioxide or hydrocarbons on metal surfaces, more complicated systems remain far out of current theoretical reach.

To be sure, Feibelman concedes, "We're just beginning to learn the rules." Subtle as these may seem, adds George, they're just what engineers will need to know as they try to design advanced new materials from the atomic level up. ■ IVAN AMATO