SEEING THE NEED FOR 'ART'

A NEW GENERATION OF 'SELF-TEACHING' COMPUTERS REPRESENTS ANOTHER STEP IN NEURAL NETWORKS

By KAREN HARTLEY

A robot, searching diligently for a particular kind of rock, crawls across the Martian soil. Suddenly, it comes across something totally unfamiliar, let's say a Martian. If the robot is equipped with today's typical computer, chances are that by the time it figured out that it didn't recognize what it saw, the Martian could have had dinner and hopped a transit to Jupiter for dessert.

But with a computer designed to process information like a human, that same robot might be capable of recognizing the Martian's unusual nature, getting out its camera and zooming in for home movies.

Such systems – patterned after how researchers think the brain's nerve cells connect, relate and relay information – are called neural networks (SN: 1/24/87, p.60). A discipline that in the past has drawn only a handful of professionals made its progress visible last week when about 1,700 researchers, designers and engineers met in San Diego at the First International Conference on Neural Networks to discuss aspects of these systems, including knowledge processing, cognitive science connections, vision, speech recognition and robotics.

One type of neural network, called adaptive resonance theory (ART) architectures, creates and organizes categories for objects and has the ability to respond immediately to its experiences. Networks with adaptive pattern recognition aren't new, but the ART architectures created by Gail A. Carpenter of Northeastern University in Boston and Stephen Grossberg of Boston University are designed to teach themselves new categories and continue storing information without trashing bits and pieces it might be a good idea to keep. Carpenter unveiled a new generation of the architecture, ART 2, at last week's meeting.

Upon receiving a preprocessed image, the ART system puts it through several steps. The system works in a series of levels. Over time, images that are learned are represented in categories on upper levels of the system. When a new image of a truck is fed into the first level, for example, it activates one of the memory categories and is sent up to be matched with that category. This "bottom-up" process is a common property among adaptive systems. What makes an ART system unusual is that at the same time the signal is going up, the receiving category is sending a signal "top-down" to the first level to make sure an adequate match exists. In essence, Carpenter says,

Carpenter et al.

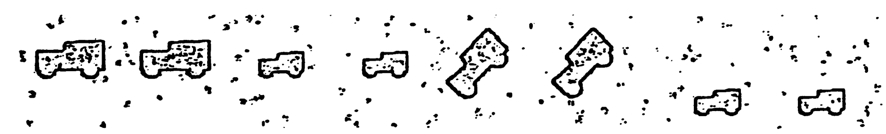

ART 2 is able to categorize identical pictures of trucks together, despite background "noise" and changes in size and position that could confuse identification.

"A category looks back down and says, 'What am I learning; what category should this be in?'" If the match isn't close enough, an orienting subsystem is activated, which automatically closes off the activated category and searches for a category that would give an adequate match. If no adequate match is found, the system creates a new category. The system "learning" process about what belongs where is strengthened when it matches an existing category.

Before reaching ART 2, the images go through a preprocessing system. In the example of the picture of the truck, the object first is detached from the image background. It then goes through a simple neural network processor, which extracts and fills in any insufficient boundaries, such as if there were a border missing from the truck. The network assigns each boundary a confidence level indicating how good a boundary it really is. Next, the image of the truck goes through a standard "Fourier-Mellin" transform, which computes the invariance for the picture if it is moved to a different spot, rotated or expanded. The resulting image is used as the input for the system. Eventually, ART 2 itself may be equipped with a preprocessor that performs the same functions, Carpenter says.

The first generation of ART architectures, ART 1, categorized images in binary, which limits the system to categorizing information represented in black or white. For example, if there were a topographic map in five different shades representing height, the binary input would make everything above a certain height black and everything below that point white.

ART 2 resolves that problem by categorizing images in continuous analog input, enabling it to make finer distinctions in analyzing and categorizing images. If the ART 2 were to input the topographic map, it would categorize all five plateaus, with the larger height corresponding to the larger value.

Because of its ability to make these fine distinctions, ART 2 is able to categorize a "noisy" or imperfect image, Carpenter says. For example, preliminary computer simulations of "noisy" images of four types of trucks showed that ART 2 correctly classified 40 pictures into four categories. ART 2 also successfully categorized trucks that had been rotated or sized differently. For the same reason it can categorize a picture embedded in "noise," ART 2 is able to make more subtle distinctions in classification. Take two preprocessed pictures of similar mountains, for example, except that one has a crater in the middle. While ART 1 would likely classify both in the same category, ART 2 has the ability to make out the finer distinction of the crater and classify it in another category, if necessary.

The system is also self-scaling, meaning that the more complex a picture, the more the system will tolerate a mismatch at a single location in the image. In addition to recognizing and categorizing images, ART 2 can react, if necessary, to what it has "seen." For example, if it came across the Martian on Mars, ART 2 would quickly realize that it didn't recognize the image, and then react by alerting an operator, Carpenter says. There are two reasons for this. One is that unlike conventional computers, it is massively parallel in nature, meaning that all the processing is going on at once.

The second reason for ART 2's speed is that once it self-stabilizes, or learns a base of information, the system automatically turns off its search system and directly accesses a category when it sees an object, bypassing what could be a lengthy search process as the system expands its base of information. The search system comes back on only when something arises that doesn't fit into any existing category, Carpenter says.

Although research is promising, rock-seeking robots on Mars and other applications of the ART architectures here on earth, such as use in underwater exploration, are a long way off. Before such applications can be considered, a massively parallel processor for ART 2 needs to be built, and that must first overcome its own problems before researchers can even begin thinking about the bugs trial simulations might bring. □