Teaching machines to recognize shapes

December 12 2011

Geoffrey Hinton gave the NSERC Herzberg Gold Medal lecture at Ryerson about how machines can be taught to recognize increasingly complex shapes and objects. His team has developed a program that can identify objects in photographs.

As any parent knows, teaching a toddler to recognize objects involves trial-and-error. A child, for example, may not initially recognize a cow in a picture-book after seeing the live animal on a farm and being told its label. In fact, a child may mistake a cow for a horse. After all, both animals have four legs.

Applying that principle of human learning to artificial neural networks, or machines, is the domain of Geoffrey Hinton, a professor of computer science at the University of Toronto and a fellow of the Canadian Institute for Advanced Research. A pioneer of artificial intelligence and

neural networks, Hinton is an expert on machine learning and has also made major contributions to the fields of cognitive psychology and neuroscience. In recognition of those achievements, he was awarded the 2011 Gerhard Herzberg Canada Gold Medal for Science and Engineering from the Natural Sciences and Engineering Research Council of Canada (NSERC). The country's highest prize for science and engineering, the honour celebrates Canada's top researchers.

Each year, the winner of the NSERC Herzberg Gold Medal delivers a lecture about his or her research. Sponsored by NSERC and the Royal Canadian Institute for the Advancement of Science, the public lecture was hosted by Ryerson earlier this month.

During Hinton's presentation, entitled "How does the brain recognize shapes?", he described how computers can learn in similar ways to the human brain and respond intelligently to the intricacies of the real world. To be certain, simulating the brain's computing abilities is no easy feat. Just consider what the human brain can do, from identifying patterns and making predictions to learning from examples and using big-picture thinking.

Teaching machines to automatically perform these high-level processes has many applications in our data-intensive world. Among them, facial recognition capabilities, quality control systems, making medical diagnoses and conducting financial forecasting. Hinton and his collaborators have developed algorithms used in applications such as creating better systems for voice recognition, automatically reading bank cheques and monitoring industrial plants for improved safety.

In his lecture at Ryerson, Hinton first showed how machines can be trained to recognize handwritten numbers that are very distorted. From there, he demonstrated how computers can predict the next character in a line of Wikipedia text or create an animated model of human

movement.

Hinton also explored how machines can be taught to recognize increasingly complex shapes, including those that may vary widely. Indeed, his team has developed a program that can identify a thousand different types of objects in photographs. The computer provides several guesses about the nature of an object, and the correct answer is usually within its top five guesses.

The computer's first guess is often incorrect. But, Hinton notes, even its wrong answers are still plausible. For example, a mound of cashews was determined by the computer to be lentils, chickpeas or beans. In addition, a quail was mistakenly identified as an otter – a reasonable error, says Hinton. The bird in the photo has a sleek coat that resembles wet fur.

"I'm an apologist for neural networks," he joked.

Hinton's research is supported by the Natural Sciences and Engineering Research Council of Canada, the Canadian Institute for Advanced Research, the Canadian Foundation for Innovation and gifts from Google and Microsoft.

Provided by Ryerson University