What "learning" machines will mean for medicine

Cite as: *CMAJ* 2017 April 24;189:E615-6. doi: 10.1503/cmaj.1095413

Artificial intelligence (AI) is having a major moment in Canada. In February, senators heard a call for a national plan to keep ahead of rapid advances in AI, and a month later the federal government pledged $125 million for precisely that strategy. Within a week, a new AI research facility called the Vector Institute opened in Toronto, catapulting Canada to the cutting edge of the field. The institute is backed by the Ontario and federal governments, and some 30 companies, including Google. The tech giant is also launching a second Canadian-based "supercluster" or global hub for AI research in Toronto; the first opened in Montreal in 2016.

The Vector Institute will specialize in "deep learning," AI that's modeled on the structure and function of the human brain. The goal is to develop computer programs that can learn, predict and adapt with the plasticity and intuitive judgment of a human mind, but better.

On the eve of Vector's Mar. 30 opening, the Senate Committee on Social Affairs, Science and Technology heard from experts how medicine will be among the first sectors transformed by deep learning AI.

"We'll see it apply to diagnostics initially, to surgery and perhaps prevention," explained Alan Bernstein, president and CEO of the Canadian Institute for Advanced Research, a major partner with Vector. Computers have long had the advantage of speed and larger memory banks. But until recently, humans have had the superior advantage of learning and prediction, that is, the wiring and rewiring of our brains in response to complex patterns in our experiences. Bernstein gives the example of seeing a ball roll onto the road: "You're going to look for a child running after the ball, not because you've seen the child, but because you know from experience that balls don't roll onto the road by themselves."

Doctors make diagnoses using the same processes — not by rote memory and checklists, but a kind of intuition based on deep layers of neural connections, hardwired by study and clinical experience.

Deep learning is now closing this gap between computing and knowing. Stanford researchers developed an algorithm that could diagnose skin cancer with similar accuracy to seasoned dermatologists. The program made its diagnoses based on a "learning set" or database of nearly 130 000 images of skin lesions, and applied the lessons from those images to unidentified specimens.

The same type of program could be used to identify disease using "any medical diagnostic one can think of," said Bernstein. Feeding such a program population health data or genetic information could also unlock new avenues for treatment and prevention. "That computer can sift through all the noise and identify the signals."

Fully autonomous robot surgeons are still the stuff of science fiction, "but certainly that's the goal we're all looking for," said Dr. Christopher Schlachta, medical director for Canadian Surgical Technologies & Advanced Robotics. Certain robotic surgical devices already give surgeons "pushback" if they go outside safe boundaries. Paired with AI, "the future of surgical and medical care lies in the interposition of a computer between the patient and the health care provider," Schlachta said. "We know from robots involved in arthroplasty surgeries even the best surgeons can achieve contact

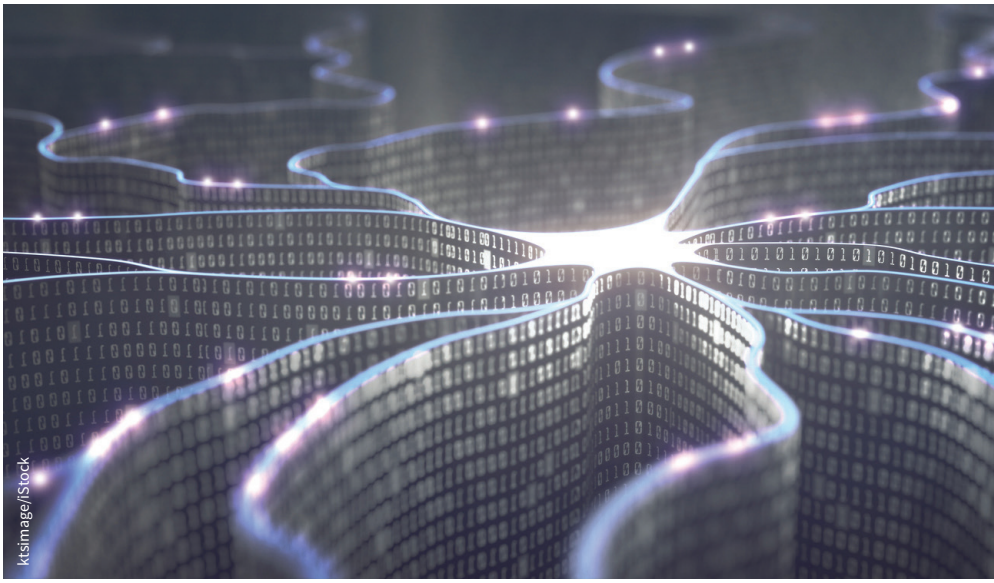

Canada is investing in "deep learning" technology that will enable machines to triage, diagnose, refer, comfort and operate on patients, redefining both the art and science of medicine.

© 2017 Joule Inc. or its licensors

rates with their implants that will be in the 70%-80% range. You use a robot and you're well over 90%."

The biggest barrier to this brave new world is simply that people don't trust robots. Even when using semiautonomous systems, "we still have surgeons preferring to actually guide the instrument themselves," said Schlachta.

Bernstein argued that once AI is proven to be more reliable than humans, "any responsible physician would only trust the computer."

That of course raises the question, why have a doctor at all? Bernstein told senators that the Vector Institute's scientific advisor Geoffrey Hinton has warned "if you're in medical school training to be a radiologist, you should change fields."

"I think that's an overstatement," said Bernstein. For the foreseeable future, "we're going to need the combination of the machine and the clinician working together to give a good diagnosis and also to give that sense of wellbeing that comes with human contact."

Schlachta pointed out that the invention of the calculator didn't put accountants out of work. When it comes to surgery, "the surgeon would assist with the preoperative plan and then the robot would carry it out."

The real challenge will come in keeping physicians and surgeons up to speed with technology that can learn and adapt very quickly. "We have to consider retraining the entire workforce at regular intervals," said Schlachta. "We're not currently equipped to handle that."

However, "one advantage of computer-assisted surgery is that the technology is digital so you don't really need a patient for training," he said. That protects against the kind of Wild West training that occurred, for example, when keyhole gall bladder surgery was introduced. "Surgeons got together in groups and travelled around the country and taught weekend courses on an animal model and then on Monday the surgeon would assault their first patient."

Senators questioned whether greater dependence on simulation in training and AI in practice would blunt physicians' skills over time. Schlachta agreed it would, but countered that few people suffer from lack of horse-riding skills now that we have cars. "As long as we're confident that those technologies will be readily available when we need them, I think we're going to be fine."

More concerning are the ethical problems, such as who sets the rules and chooses the data that guide AI decisions. Deep-learning programs will work from vast repositories of information, so the rationale for their decisions will be unknowable, much like our own intuition.

To some degree programmers can build rules, said Schlachta. Like the Hippocratic Oath, Asimov's first law of robotics is to "do no harm."

"The other thing that gets overlaid on that is compassion," Schlachta said. "The greatest fear everyone has is an unfeeling machine."

Lauren Vogel, *CMAJ*