The future of health care could be elementary with Watson

As we draw closer to a future where artificial intelligence may play an active role in health care, should we be putting our trust in a supercomputer with a seemingly infinite amount of medical knowledge or a clinician with decades of experience in the field?

IBM's AI supercomputer Watson seems to be making that decision for us.

Since beating two top trivia champions in an episode of *Jeopardy!* in 2011, including all-time winner Ken Jennings who walked away with US$3 172 700 in total earnings, Watson has quietly been learning the ins and outs of the health care industry.

From being trained by oncologists at the Memorial Sloan-Kettering Cancer Center (MSKCC) in New York City in optimizing treatments for lung cancer to streamlining the approval of medical tests with the largest for-profit health benefits company WellPoint, Watson has definitely been keeping busy.

The real benefit in using Watson throughout the health care system is to "be able to train Watson so that oncologists in the community, who may only see a handful of cancer patients a year versus our thousands, can have that same level of expertise," says Chris Hickey, a spokesperson for MSKCC.

In addition, new advances in oncology don't reach everyone at the same time, but having a supercomputer that is automatically updated with the latest research at the click of a button will provide oncologists with a great advantage, she says.

Using natural language processing will also allow physicians and patients to express symptoms and preferences to Watson, which will let the supercomputer see patients as individuals and not just as numbers on a page.

"Watson works in natural language, you don't have to break it down into key words — it's not a search engine that uses keywords. It actually completely understands the question you have raised and you ask that question in natural language," says Dr. Martin Kohn, chief medical scientist for care delivery systems in IBM Research.

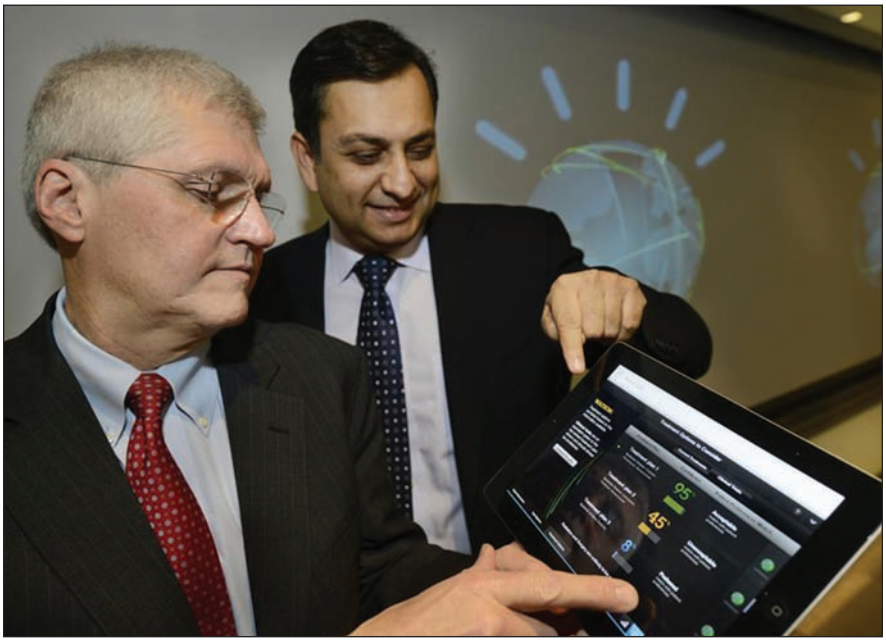

IBM

Dr. Mark Kris (left), chief of thoracic oncology at Memorial Sloan-Kettering Cancer Center in New York City, and Manoj Saxena (right), IBM general manager of Watson Solutions, work with the first Watson-based cognitive computing solution for oncology.

"So you could say 'Watson, I have a headache and my right eye hurts, and my left knee is swollen and I had a fever of 102.5 yesterday, and I'm sneezing.' You could ask it exactly that way … and Watson will understand it."

From that statement, Watson would look through large volumes of literature, in addition to the patient's electronic medical record, and then return with ideas and suggestions for the patient that would be individualized and supported by evidence.

"This file of information would then be shared with the patient's primary care team," says Kohn. "So not only does the patient have access to personalized evidence-supported ideas, but also has somebody to talk to about it with … and bring the patient into the decision process."

Watson can also assess the resources available to clinicians in a given hospital and tailor its suggestion to them based on their equipment and location, while oncologists can reverse-analyze Watson's suggestions to determine from which specific medical literature it drew its conclusion.

Currently, Watson is focusing on hypothetical lung cancer diagnoses, but IBM hopes to reproduce the knowledge it gains from working with oncologists to other cancer types and diseases down the road.

"We'll spend a lot of time getting the first training right, and then all of those lessons will apply to the next disease and the next disease and the next disease," says Thompson.

"Going forward it's going to be hard to imagine a decade from now that most physicians aren't going to be using some form of this technology, just to make sure they haven't missed some alternative."

Watson isn't physically at MSKCC or WellPoint, but resides in IBM sites and is accessed remotely on computers and tablets with appropriate controls for privacy and security. The end goal is

that Watson will be accessible through the cloud for users in the future.

"We're working with oncologists from [MSKCC] and WellPoint in this effort to teach Watson enough about patients with cancer so that Watson can help bring forth suggestions for therapy that allow the patient and the physician to make a good choice for that patient," says Kohn.

Ultimately, though it's not being worked on yet, the goal is that patients could also use Watson for their own diagnoses and treatment information.

Kohn says this is part of the concept of the "empowered knowledgeable patient," the idea that the more a patient knows about his or her health status, the better manager of their own health they will be and the more prepared they will be to participate in the decision-making process.

"Watson, in preparation for this project, read every medical textbook available in the world. It then read the vast majority of medical journals that they could get a licence for," says Dr. Craig Thompson, president and CEO of MSKCC.

"When we first encountered it in our training sets, it was as good as a third-year medical student at answering questions. Every day it's a little bit smarter — right now it's as good as our fellows."

Thompson adds that the training of Watson through hypothetical cases with the help of oncologists will be completed this year, at which point it will be beta-tested to start working with actual patient files.

"It'll take the rest of the year but we expect it to start to gain the same 20 years of experience [as an experienced physician], in a much shorter period of time because it's being constantly fed information," he says.

"Then it will be available like a senior physician just the way we call our colleagues up every day and say 'you've dealt with a lot of these cases. What would you do next?' Except it will be stored information from the physician that trained Watson in that set of disease diagnosis and treatment."

Despite the advances to date, oncologists have faced some challenges in the training phase that could prove troubling if not addressed.

"The challenges in training it have been the nuances in medical language," says Thompson. "Watson is still learning about this, but that's this issue of natural language processing. When it was taught *Jeopardy!* it read dictionaries. … In a medical radiology report when it says 'suspicious of lung cancer' that doesn't mean that you should decide that the patient has lung cancer."

"The computer needs to know how to weight the word 'suspicious.' So we've learned some things about how doctors communicate probabilities and teach the computer how to think about those kinds of words."

Another problem that has arisen in the training of Watson is the concept of family history and the genetic understanding of cancer.

"The computer doesn't completely understand family trees, so if the family history says the patient's father had lung cancer that doesn't mean the patient has lung cancer. And that sounds really obvious to people, but it's not always so easy in writing the algorithms to analyze medical information," says Thompson.

"So there's been lots of training sessions and then corrections — it takes a whole team of physicians in each of the diseases we're training Watson in [along] with the Watson research team to really make sure that the computer isn't learning incorrect facts."

These problems are just the tip of the iceberg for AI experts like Dr. Jesse Hoey, assistant professor of computer science and a researcher in artificial intelligence and health analytics.

"One of the major limitations of this type of system is not what it can do for the 99.9% of people, but what it can fail to do for that 0.1% of the people," he says.

"When you think about the system, it can play *Jeopardy!* Well, that's an incredibly precise game. There's one and only one answer and you have to get it exactly right. But what if the *Jeopardy!* question was more vague? Then would it be able to still answer the question?"

The real risk with Watson is its limitations in bringing the intangible human component to diagnosis, Hoey says.

"The things that [Watson] can't do are the funny little intangible things that humans can do. And a lot of the things that humans can do are based on gut feelings — so AI does not have gut feelings."

"Doctors are not just data-processing machines, they're humans that talk to other humans and they operate on gut feeling some of the time. … We all know that AI isn't able to handle that. There's people working on that but we're so far from being able to do that."

Yet Hoey isn't convinced that the hype surrounding Watson won't lead the partnership between AI and health care down the wrong path.

"If you start building medicine… so that it's all only based on data that's processed by machines I think it's probably a dangerous road to go down at this stage, just because artificial intelligence has not developed that intangible aspect of human interaction," he says.

"The way that people want to be cared for in hospital, their sort of emotional needs, is incredibly important and I don't think Watson's able to do that. So I think it would have to be approached with a lot of care and I'm not sure that they are approaching it with a lot of care. I think people are wowed by technology and they're willing to jump in with both feet without thinking too carefully about what they're getting into." — Adam Miller, *CMAJ*

***CMAJ* 2013. DOI:10.1503/cmaj.109-4442**