

Where is AI headed in 2018? Your phone will know you better than your friends do, researcher predicts

January 5 2018

From self-driving cars to finding disease cures, artificial intelligence, or AI, has rapidly emerged as a potentially revolutionary technology – and the pace of innovation is only set to speed up.

To get a sense of where the field is headed in 2018, U of T News sat down with the University of Toronto's Richard Zemel, a professor of computer science and the research director at the Vector Institute for artificial intelligence research.

He was just back from the annual Neural Information and Processing Systems (NIPS) conference in Long Beach, Calif. – a once staid academic gathering that's more than tripled in size over the past five years, drawing dozens of giant corporations.

Zemel's take? Get ready for a world where businesses enjoy unprecedented insight into their products and services, Toronto continues its ascent as a major AI research hub, and digital assistants like Apple's Siri and Amazon's Alexa become ultra-personalized.

"It will be like being friends with someone for many years," predicts Zemel, who spoke on the sidelines of an event at the Creative Destruction Lab (CDL), one of U of T's numerous entrepreneurship hubs. "The computer or phone may know more about you, potentially, than anyone else."

What can we expect to happen in AI over the next 12 months?

One of the things that will affect [people](#) in their daily lives will be personalization. People know about personalized assistants like Alexa and Siri, but those are just the first generation. They're going to get a lot better in the next six to 12 months. They will be able to really understand what you're asking and be able to formulate answers and get to know you better – not just look things up in your calendar or on the web.

Is that a function of improved speech recognition capabilities?

Speech will be part of it. But there's all this other information about you that's available. It's your daily habits, what you do and where you go. So, if people allow it – if you give it access to your emails and photos, what you look at online, watch on TV and the books you read – it's going to be a much bigger package. It will be like being friends with someone for many years. The computer or phone may know more about you, potentially, than anyone else. So it's a question of combining all that information and getting a real profile of your tastes.

Beyond personal assistants, there's a lot of other industries that are potentially going to be impacted by this technology, if they aren't already. What other sectors do you think we may be hearing about?

Education is one example. There will be more systems that learn how you learn best. These could be online learning tools that are custom tailored to you. There will also be lots of manufacturing applications. Here at CDL, there's a lot of companies who are using sensing

technologies to find out what's happening in the environment – ranging from smart cities down to a company I was just chatting with that's putting sensors in cows' milk to determine how healthy it is. All of these things that were typically very expensive to assess can now be done with a few sensors and a lot of training data.

Medicine is another area where you have a ton of data, although there's a lot of privacy and data-sharing issues in that. The big companies like Microsoft and Google have talked for years about getting into it, but they've always stopped because of privacy issues. But I think there's enough momentum now that there will be progress in health, ranging from health records to medical imaging diagnoses and robotics in surgery.

On the research side, which areas do you find exciting right now?

One of the most exciting areas – and it's reflected in the research I'm doing – is called transfer learning. That's the idea of performing a new task without a lot of training data. This has a lot of applications in business. Let's say a robot has to climb hills and take out the garbage, and it has a lot of data it's trained on to do that. But now you give it a new task – moving a bin from one place to another, and it's never done that before. So now it has to transfer its knowledge to this new task. The novel thing here is you're training it with a huge amount of data, but you're testing it on something else.

I'm guessing that's more difficult to accomplish than it sounds.

Exactly. That's the interesting piece in all of this. The things that seem easy to us, because people do them naturally, are typically the biggest

challenge for these systems. That's true for perception and speech. I mean, we speak pretty easily, but it's taken computers a long time to figure out speech and produce speech. All these things we take for granted are big challenges.

Any other areas of research interest?

Another one goes back to what we were saying about personalization. If it takes off, you will need to allow the [AI] system to see all your personal details. So it becomes a question of whether you're going to be hesitant to release your personal details because of privacy and fairness issues. So there's a lot of research now – a really growing field – called fairness in [machine learning](#). I do a lot of research in this. In past years, there were just two or three papers at NIPS. This year there were 20.

Is this a technical issue or an ethical one?

It's both. It's an ethical and societal issue to define what it means to be fair, but the technical issue is how do you build a machine learning system that embodies those principles? It's a very interesting area. The challenge is defining it in a good way, and then you take this definition and formalize it into a mathematical statement we can use to train up the machine learning.

There's also a field called FAT ML, which is fairness, accountability and transparency in machine learning. That's going beyond fairness and privacy to ask whether you can get explanations from the system, so it can be useful for doctors and lawyers. In those high-risk situations, you need systems that are more interpretable. That's an increasingly important direction, too.

[U of T University Professor Emeritus] Geoffrey Hinton made a

statement last year about how the current paradigm for deep learning needs to be thrown out the window so we can start over. What do you make of his comments and what should the rest of us make of them?

What he's talking about actually follows a trend in machine learning. And that's to build more structure into the system. Now, a lot of people don't like that point of view. They think of it as a plain vanilla system and you should allow it to learn everything. But he's saying you actually need to build in some structure – capsule networks – where you learn about parts of objects and some parameters associated with them and how they're related. There's a lot of work in that area. There's always been this debate between a sort of tabula rasa view of learning to something where you start with some structure and learn on top of it. So throwing out current [deep learning](#) might mean, in my view, that you want to incorporate some sort of structure and the key question is: What's the right structure?

AI has been a big story for U of T and Toronto, particularly with the creation of the Vector Institute last year. Can you give me an update on what will be happening at Vector in 2018?

It's exciting. We finally have the space to move into with desks, chairs and everything. People were worried we wouldn't have enough critical mass to get going. But, actually, when we put it all together, there will be 90 people moving in, including students, post-docs and faculty. And that's just the full-time people. There will be a lot affiliates coming in part-time from other parts of Toronto and the province, including the Universities of Guelph, Waterloo and McMaster. It's going to be a real hub of AI activity. We'll do some growing as well, hiring some additional research scientists and bringing in a new batch of grad students in September. We're hiring some post-docs, as well as software

and research engineers.

How will the relationship between Vector and U of T work?

The idea is people will be working at U of T in their faculties, but are also cross-appointed to Vector, so they will move back and forth. Many will also have students whose main desks will be at Vector, but they will also be teaching and holding talks on campus. The way I think about it is as an additional facility with a lot of good researchers who will facilitate collaboration.

In a broader sense, how do you see the concept of an AI hub developing in Toronto?

There's a lot of AI around. A lot of the hospitals are doing AI and health is an application that we're very interested in at Vector. We're going to try and co-ordinate things to get the hospitals working and talking with each other and sharing data. We could also play an important role working with businesses when it comes to finding talent. One of our main aims is graduating master's and Ph.D. students. We're not going to co-ordinate all AI, but we can be an important resource and hub for research.

What do you think are the misconceptions rattling around out there about AI?

One thing people don't realize is that machine learning systems, the way they are right now, require a lot of data and a lot of labelled data – thousands of examples with annotations. If you don't have that then you're in the research field, not the applications field. People need to

know that. You need a lot of data.

Is it difficult to get access to sufficient data in a less populous country like Canada?

Not really. There's data everywhere. It's just a question of harnessing it and figuring out how to get labels. We're big players in this. People are trying to emulate what's going on here. I had meetings with top people all over Europe, asking "How did you do Vector? We want to copy you." I think we're sitting on a model for the rest of the world.

Provided by University of Toronto

Citation: Where is AI headed in 2018? Your phone will know you better than your friends do, researcher predicts (2018, January 5) retrieved 7 July 2024 from <https://phys.org/news/2018-01-ai-friends.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.