Christina Griffith
May 2019

Artificial Intelligence For The Real World: A discussion with Professor Dan Roth

Professor Dan Roth does not appear the way you would expect: he is an intimidating figure dressed all in black, more reminiscent of a bouncer than a pioneer in computer science. Roth is in high demand: his schedule is tight, and there are several students waiting to meet with him on this afternoon. The computer and information science wing his office occupies is populated by ultra-modern, colorful chairs surrounding small tables, inviting collaboration in open spaces. The walls are almost completely lined with whiteboards bearing the faded evidence of intense mathematical calculations with dry-erase markers.

Mathematics, it so happens, is where Roth began his journey as an innovator in machine learning and artificial intelligence. Born in Haifa, Israel, Roth served in the Israeli Defense Forces after graduating Summa cum laude from Technion with a BA in math. Israeli culture values education, elevating teachers and scientists to hero status. This emphasis on the importance of education had a powerful impact on the direction he would take in his career.

In the army, his skills were put to work in research labs. Once he realized all the interesting things that could be done with computer software, he was converted. Upon completing his service, he attained a PhD from Harvard in computer science.

Roth is the Eduardo D. Glandt Distinguished Professor at the Department of Computer and Information Science at the University of Pennsylvania. Before coming to Philadelphia, he was a Founder Professor of Engineering at the University of Illinois at Urbana-Champaign's computer science department. In 2017, the International Joint Conferences on Artificial Intelligence Organization named Roth winner of the John McCarthy Award, which recognizes mid-career AI researchers for advances in the field.

"Artificial Intelligence gets a very bad rap," Roth laments, "mostly from people that have no clue what I'm talking about." He speaks with an accent, and with his hands, smiling and maintaining eye contact, completely relaxed. His command of his field allows him to translate the most complicated concepts in computer science to someone with a background in the humanities. It isn't just popular culture, he insists, but also what is written in the media that produces this disconnect between what the public think about AI and what he actually does.

Developing computers that can carry out tasks that have not been programmed, and to process natural language data is Roth's pursuit. The fact that today we can tell our Siris and Google Assistants to put items on a shopping list, call your sister, or set a timer does not mean we have technology that can learn or have conversations with us. Natural language processing (NLP) is exceedingly difficult, something most people don't realize because it's something we do literally all the time. Roth identifies the two principles that make language understanding so difficult: ambiguity and variability.

"Everything we say in language is ambiguous," he explains. "Everything depends on the context in which it's being said, who it's being said to, who is saying it, you know…When you say the word table, what is table? Table has multiple meanings." The realization that an object as simple as "table" is actually very ambiguous is where we begin to understand the problem of getting computers to process language.

The variability in language is the way in which we convey the same meaning in different ways. Putting together variability and ambiguity, Roth sums it up as, "Everything you say in language could be many, many things, every meaning that you want to convey can be said in many, many ways." While much of his published work has made great progress in the ability of computers to process language, invariably the media has headlined press releases and award-winning papers by stating the problem of AI or language processing has been "solved." Roth shakes his head. "Everyone that actually has played with Siri, Alexa, your favorite things– knows that they have no clue what you're saying."

Humans, he points out, "put together many things in order to understand." We do this without thinking, it's how our brains function. What Roth and other scientists are doing is trying to create machines that can learn to do this. Programming a computer to respond to a factual question requiring an answer that can be calculated or searched based on known data is easily done today. Programming a computer to respond to the question of whether there are more black or purple objects sitting on the table is not yet possible. "That is something that the three-year-old can answer; but think about it, this is a complex question. It requires understanding the question. I need to look for black things, I need to look for purple things. I need to count them, I need to compare them. We don't know how to do this."

There is one particularly irksome problem for which Roth believes machines that can process natural language are the solution: information pollution.

Roth views this problem far more broadly than just fake news. As he sees it, there is an immense amount of information available to people from endless sources across the internet.

There are search engines, social media sites and the groups you can join within them, platforms like YouTube, message boards, blogs, online encyclopedias–and all of these are providing information without any context or verification process.

Searching for information on almost any topic gives the user literally millions of results. Roth throws up his hands: “Who’s saying what? What’s really truth? What’s myth? What can be substantiated? You know, there’s so much stuff out there. So to me, this is *inherently* a natural language understanding problem that we can make progress on.”

Roth and his students are at work on creating a way for artificial intelligence to do what a human behind a keyboard cannot. Instead of search results being presented as a list of 10 million documents organized by relevance to the keywords or by date, AI can help us “navigate the information polluted world.”

“We think about things as a claim. A claim could be, vaccination causes autism. Or it could be animals should have human rights. And then, given a claim, we want to find perspectives.” Roth posits that most interesting claims have multiple perspectives. Not every claim is true or false, perspectives may lie along a spectrum or are conditional. Maybe you believe all animals should have human rights, maybe you think the opposite, or maybe you think only cute animals deserve human rights. In each of these cases, the next step is to identify evidence that supports the perspective.

All of this organization and evaluation requires that the application understand what it is looking at and reading: it must be able to understand natural language. “If we understand language better, rather than presenting information as a bunch of documents that google retrieved for me, I’m going to organize this information in perspectives with evidence.” Roth predicts that eventually, algorithms will be able to assign levels of expertise to the people credited with the evidence in support of or against a claim, and that at the final stage, it may be sensitive enough to determine “trustworthiness.”

The eventual result of Roth’s work is not a machine that learns, thinks, or performs tasks the same way we do. The future of artificial intelligence is not to be like humans, as Roth sees it, but to “augment” us. With the aid of AI, humans will be able to do more, and do it better.