MIT MANAGEMENT EXECUTIVE EDUCATION

CSAIL

MIT SLOAN SCHOOL OF MANAGEMENT

MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY

ARTIFICIAL INTELLIGENCE:

IMPLICATIONS FOR BUSINESS STRATEGY

ONLINE SHORT COURSE

MODULE 1 UNIT 1
Patrick Winston Video 2 Transcript

© 2017 MIT Sloan School of Management / GetSmarter All Rights Reserved (not authorized for commercial use)

DELIVERED IN COLLABORATION WITH getsmarter

Module 1 Unit 1 Patrick Winston Video 2 Transcript

PATRICK WINSTON: In about the mid-1980s came what has come to be called the AI winter; confident projections did not work out. Very few programs replaced human experts. Many startups aimed at replacing human experts failed. Venture capitalists started to use the term another AI as a label for poor investments. But not all the startups failed. I helped start up a company we decided to call Ascent Technology and we called it that because our products were targeted at airlines and airports.

Importantly, our goal was not to replace experts but to do things that couldn't be done by either people or computers working alone. We focused on the more efficient use of expensive assets for airlines and airports, including people. Our first produce enabled better gate management. Everyone loved our gate management system, but nobody bought it, they wanted to know who we already had under management and of course we were a startup – we didn't have anybody under management yet but we got a big break during the Gulf War – where we helped build a system called DART which was for planning force deployments to the Middle East.

We reused much of the technology that we developed for managing airline and airport resources. After that, when airline and airport customers asked who we had under management, we said well, the entire US Department of Defence, how is that? And, of course, that worked, and now if you fly to a destination in the United States, there's about an 80 per cent chance that your aircraft will pull up at a gate determined by an Ascent system.

Today we have a whole suite of products aimed at working with people to manage resources that neither a computer system nor a computer can manage as well alone. One of those products ensures that the work force is used efficiently, but it doesn't just increase efficiency, it enables workers to trade shifts with other co-workers and to choose the schedule they prefer, making for a happier work force. And that's something that really appeals to me, using AI to help people have a better life.

But now back to the AI winter. There was slow but steady advance of knowledge and companies that made new things possible like Ascent succeeded. Also companies whose products were aimed at routine work succeeded. But nobody paid attention to AI, for the most part before, the third wave began to appear. Then in 2010, we got Siri and we could talk to our telephones. In 2011, the Watson program, the one that played jeopardy, beat the human champion. So, now AI was beginning to show signs of a kind of resurgence, a third wave.

IBM, encouraged by all this, committed itself to cognitive computing in a big way. Google, Facebook, Amazon, Microsoft are all in. In general, the third wave has been enabled by lots of computing and lots of data, both of which enable a new kind of statistics. Some call this new kind of statistics machine learning. Others call it computational statistics, but whatever you call it, one important part of it is deep neural nets.

And one signature program based on deep neural nets is the Google captioning program. For this picture, it impressively suggested a pretty good caption. A group of young people playing a game of Frisbee. Of course, it's important to realize that programs like the Google Picture Captioner appear to have more understanding than they actually have. The

© 2017 MIT Sloan / GetSmarter All Rights Reserved (not authorized for commercial use)

captioner doesn't know what play is all about, it doesn't know what it means to be young. It doesn't even have any understanding of motion because it's never seen anything other than still pictures. Nevertheless, it's enabled to do a pretty good job of captioning, but we shouldn't think that it actually is thinking in the same way that we think, just because it's behaving in some respects like we would behave.

It's interesting that the biggest idea in deep neural nets was worked out by Paul Werbos in 1974 in his Harvard PhD Thesis. He showed how neural nets can learn using what is called back propagation which in math speak is what we call gradient ascent. For 30 years or so, few people paid attention until a program by Geoff Hinton and his co-workers showed impressive results in a picture classification contest. Because neural nets are called neural nets, it's easy to think that they work like the neural nets inside of our heads, but evidently not. There's a whole community of people who like to show how to mislead these deep neural nets. Both of these pictures look like school busses to us, but not to one of the picture classification programs. One thinks the one on the left is a great school bus, and the one on the right isn't a school bus at all. You could even make images that fool neural nets into thinking that images are school busses even though they clearly are not. This one looks like noise, another one is black and orange stripes – not school busses for sure, but deep neural nets can sometimes be fooled into thinking that pictures like these are something that they aren't.

So, where are we? Where are we with machine learning and especially with deep neural nets? No doubt about it, deep neural nets do amazing work, but they don't see like us, they don't think like us and they don't even think at all. What they focus on is perceiving and recognizing. They're an important engineering triumph, an important part of the story, but not the whole story and that's the bottom line. We need more than machine learning than deep neural nets have been shown to offer. Nevertheless, machine learning and deep neural nets have excited and worried some well-known people.

In 2014, Elon Musk said, "With artificial intelligence we are summoning the demon. AI is our biggest existential threat." A pretty worried kind of remark and I must say AI is not on my top list of five. He's not the first one of course to have such a grim forecast. Here's a puzzle. Who said, "Once the computers get control, we might never get it back. We should survive at their sufferance. If we're lucky they might keep us as pets."? Sounds contemporary but that was Marvin Minsky in 1970.

That was about when Alan Newell and Herbert Simons said, "Symbol systems modelling human problem solving are the answer. We'd better alert the people." Turned out we didn't have to alert the people. A few years later in 1982, Edward Feigenbaum had another answer. His was that rule-based systems are the answer. In his book, he argued that Japan was about to use rule based systems to dominate the world economically. It didn't happen. Then around 1990, Rodney Brooks had yet another answer. He believes that his subsumption architecture is the answer. He was inspired in part by evolution and how evolution advances by building layer on top of layer. In his subsumption based robots, the lowest layer is about obstacle avoidance, then comes wandering, next exploring, and so on up. Brooks said that once we figure out how to make robots at an insect level, the rest will be easy.

Well, using subsumption to build robots is an important idea but the rest isn't easy. We have the Roomba vacuum cleaner, but we don't have human intelligence. Now we're in the deep neural net era and it's been brought on by people such as the pioneering and persistent Geoffrey Hinton and Yann LeCun and others. A program out of Hinton's research group impressively won a picture labelling contest after which interest in deep neural net technology exploded.

© 2017 MIT Sloan / GetSmarter All Rights Reserved (not authorized for commercial use)

But what should you make of all of this? I think what you should understand is that there are many ideas and each has been promoted by enthusiasts as the answer. Each brought excitement and worry and each will continue to have scientific and commercial value. Each idea is surely part of the story but not the whole story. Marvin Minsky made this a big point in his books. He said it basically, “We need it all. We need all kinds of representations, all kinds of methods, and ways to learn all working together.” AlphaGo which is the program that beat the World Go champion, Lee Sedol, is sometimes viewed as a demonstration that deep nets all you need.

Actually, it’s a pretty good example of needing it all. AlphaGo works because it’s a blend of deep neural nets together with Monte Carlo search reinforcement learning and carefully designed board features all of which have been around for a long time – putting them altogether is what makes AlphaGo the program that it is.

© 2017 MIT Sloan / GetSmarter All Rights Reserved (not authorized for commercial use)