

The Cloud Computing News India


News And Info On Cloud Computing And Virtualization

<http://cloudcomputingnews.in/artificial-intelligence-out-of-the-futurists-lab/>

Artificial Intelligence: Out of the futurists' lab!

26/09/16

September 26, 2016



Kathryn Hume, President, Fast Forward Labs, talked about the future of artificial intelligence, the ways in which it's impacting enterprise operations like security.

Let's talk about the future of artificial intelligence, the ways in which it's impacting enterprise operations like security say we've seen IBM Jeopardy – IBM put out the Watson – the Watson tool that beat Jeopardy in 2013. Shift to 2016.

We've seen Google DeepMind build AlphaGo which is a tool using a technique called reinforcement learning, a set of artificial intelligence algorithms that put in position a system of rewards to train systems to excel at a particular task – beat Lee Sedol who is the leading Go champion at this game.

Executives say, interesting that DeepMind can do this but what does this mean for my Churn analysis? What does this mean for my ability to optimise marketing programs? What does this mean for my ability

to protect my networks? The answer is, it's a really hard question to solve. It's not trivial. It's very hard to go from that particular original fable use case of the gold to find the actual yield in the crops that will come from these applications. So the question that I'd like to pose is, where is this land? How do we find it?

Since we're in Silicon Valley one of my mentors is a man names Geoffrey Moore who in the 1990s wrote a book called Crossing the Chasm – bible of marketing theory and how technologies move across their adoption life cycle. Right now in the AI space we're in that very early part of the distribution curve.

We're working with the early adopters. We're building out vertically specifically applications that are the early instances of what this might become but don't yet foretell large wide scale adoptions in the pragmatic larger marketplace that's actually going to have the larger impact of these technologies across the industry.

So early apps aren't always the killer apps. DeepMind winning Go isn't necessarily the application of reinforcement learning that's going to really transform inventory and supply management. But it's possible for us not to predict exactly what those killer apps are but to hone our skills to recognise and pivot our strategy when they occur.

So for the rest of the talk I'd like to just give two examples of where we're already seeing this occur in the artificial intelligence space, to give some insight into where we think the space might evolve in the future. Fast Forward Labs, at my company we study technologies that are on the threshold of shifting from the academic sphere into wide commercial adoption and applicability.

We educate our customers as to how they can apply them in their own business processes and business environments by building out prototypes with reports that explain what these tools are and what one can practically do with them and then advising them on where they can be applied in their individual environments. Our first work was on a technique called natural language generation.

I'm sure everybody has heard of natural language processing which is starting with lots of unstructured text that humans write. It's messy. It's not amenable to the numbers and digits that computers work with. So these techniques are focused on putting structure into unstructured text. Natural language generation is the opposite.

We start with structured text – not text, excuse me – structured fields like an Excel spreadsheet that has rows and columns with data entries inside of it, and then automatically write articles that use qualitative language to communicate those quantitative insights. When this application – when this set of technologies was first released to the market we thought that the killer app was going to be in automated journalism.

I'm sure you've seen in Forbes and the Associated Press articles that are written by computers that describe sports performance, weather reports, company earnings reports, lots of descriptive, non-

interpretive, relatively repetitive type reporting. That was the focus of this technology when it first came to the market.

Since that time Automated Insights and Narrative Science who are two of the key players in this space have realised that the real value of this tool is not in journalism per se but rather in narrative applications of business intelligence. Taking the mess of numbers that exist in data warehouses there was a first generation of tools like Click and Tableau that would provide nice visual interfaces to provide charts and graphs to executives.

They said that's not clean enough. We want it even simpler. We want it in human language where we can very quickly have insights into how our business is performing. We started off with gold – automated journalism – and shifted into something relatively more prosaic, right, narrative business intelligence but that has really had a yield for business. The second technique we've seen evolve is in the space of deep learning.

It's very hot right now. Most of the time when we talk about artificial intelligence we are referring to these artificial neural networks that are powering a lot of new applications and capabilities. We did a report focused on using these techniques to automatically discern the objects that are in images. So the application you see there, hook up to your Instagram feed and it will classify and reorder your pictures according to the objects that are in them.

Just as a funny side note these systems – we're going to talk a little bit later about supervised versus unsupervised learning. So they're supervised. They have to start with a training data set in order to get a vision – what their vision of the world might look like so that they can perform. My colleague, Hilary Mason, likes to take pictures of the New York City subway system on her way to work.

The training data set that we used for this tool had no images of subways in it. So it used to classify all of these gates as correctional institutions, as prisons, which both tells us about the limitations of the tools and also give some insight into why New Yorkers have their temperament when they go to work in the morning.

We start off with fun applications to classify images. Who cares? Where is the crop? A couple of cool things I've seen on the market come out since we worked on this project, one is applying – this is coming out of the artistic world – there's a technique called style transfers where artists start with famous paintings like the Starry Night from Van Gogh, abstract out the style and then enable people to go through their Facebook page and turn their selfies into works of art imitating the grand masters. This is actually using the same set of techniques that enable us to see objects and images also enable us to abstract out the style and graft it onto other images. So this is an app called [Picasso] from a start-up out of St Louis.

With a little bit more gravitas and enterprise importance a set of start-ups, one based in Silicon Valley called Orbital Insight that's taking satellite data available from a whole host of new small satellite

providers and enabling us to use these convolutional neural net deep learning techniques to gain insights into macroeconomic activity where the data was not formerly available.

The image on the right is a picture of shadows from buildings in Shanghai that hedge funds are using as a proxy to try to get an estimate of macro-economic activity going on in areas where there is not traditional market and pricing data. The third application, my personal favourite, comes from a San Francisco start-up called Enlitic that's using deep learning image technology to automate radiology to examine chest x-rays.

Stunningly they are getting the speed of their system able to complete the work that a typical radiologist would complete in one month, in 104 seconds. So it really could have massive impact on work flows, the way in which hospitals are managing their work force in the future. Final example – and its transition into the presentation from Cylance – comes from the world of text. We did a report using deep learning techniques, a slightly different style than those used to process images, to automatically summarise text.

The application here we start with an input data set, say a relatively long article in the New Yorker or the Atlantic magazines here in the US. We build a model of the meaning of the article and then use that to select out the sentences we think best represent the meaning of the article as a whole. So to shift around the goal here is not to absolutely revolutionise journalism but to shift around the reading experience so that users can start with the skimmed main points of the article and then read the entire article at will if it suits their purposes.

What's interesting here is we shift – once again gold to yield and crop – is that we were focusing our models on aspects of text that are relevant to summarise, relevant to the human reader. Stuart and his team at Cylance are adding in additional features from text to discern whether or not incoming data is malicious or not.

So same sort of technique, just shifted to a slightly different use case. The moral of the story I think going back to The Tortoise and the Hare one of my favourite Aesop fables I'm sure everybody knows of, is that we're at the beginning. We really are at the beginning of this.

We are not really sure where things are going to pan out and how they're going to pan out. But there is massive opportunity if we stop paying attention to the hype and really focus on the discrete applications and use cases that are going to be available on the early market.

(Courtesy: NetEvents)