

NETEVENTS

GLOBAL PRESS & ANALYST SUMMIT

DRAFT

Conference Debate Session V— Where Will Artificial Intelligence Make the Biggest Impact? And Where Will It Fizzle?

Chair: Jason Bloomberg

Contributor, Forbes; President of Intellyx

Panellists:

Dr Vinod Peris

Senior Vice President Central Software Group

Sam Liang

CEO of AISense

Greg Fitzgerald

Chief Marketing Officer of JASK

Jason Bloomberg

All right well, good morning everybody, I'm Jason Bloomberg, President of Intellyx and they've asked me to talk about artificial intelligence, how it's making a big impact - or how it's fizzling, I guess we'll find out.

I have - they gave me five minutes to talk about AI, it's like how do you take a subject like artificial intelligence and boil it down to five minutes?

A bit about me, Intellyx is an industry analyst firm, we focus on agile digital transformation, AI is only one of many disruptive [trans] and enterprise IT we focus on.

Here is the question, how do I introduce all the controversies around AI in the time allotted? Well the answer is cartoons. So I'll give you a series of cartoons, five minutes of cartoons and each one has a point. You've probably seen some of these, block-chain, I'll give you a moment to read it.

Obviously one of the key challenges with the AI is the vendor hype and what we call AI washing, this washing suffix. We did cloud washing where vendors started putting cloud in their marketing before they figured out cloud and now a lot of vendors are putting AI machine learning in their marketing before they really have much of anything, so that's one of the trends, it makes it hard to tell the real players from the pretenders.

Next cartoon. There we are. Machine learning is part of the AI story but it's one of the simpler parts and in terms of overall adoption, machine learning is probably more widely adopted than other forms of AI, cognitive computing, deep learning. So a lot of the times you hear about AI, what they really mean is machine learning, it's a powerful set of technologies, machine learning specifically, but it's not the whole AI story.

Okay so next cartoon, you've probably seen this one too. So of course this brings up questions of control and responsibility, right? The Uber autonomous vehicle accident continues to be in the news, just recently they determined that the software in the car told it not to stop for an emergency and that's part of why a fatality resulted.

Who is at fault? Well if it's the software that's the problem, is it the developers at the company? These are questions, you know the ethics of AI that we have yet to really figure out and there's a lot of very complex subtleties to who is responsible for what, when AI does something on its own.

So, next cartoon. I'll give you a moment. Of course this is a big data problem, if the answers look wrong just keep stirring until they look right. This is one of the challenges as well, the challenges of bias, accuracy and veracity.

Now the word veracity, whether or not something is true or not, that's one of the Vs that we're familiar with from the big data world, right? Three Vs or seven Vs, volume, variety - well veracity is on the list. But in the big data world veracity has always been a challenge because we haven't really had technology that told us whether some bit of information is actually true or not. Is it accurately representing the real world or is it something inaccurate and fictitious? 3:46

Well this becomes a big problem that AI can potentially solve. The fake news problem is one of these big hairy challenges, can AI solve the - can AI recognise fake news? Can it protect us from fake news? Can it actually give us text-based information that is unbiased?

Also do we really want unbiased information? Where people who lean left prefer left-leaning biased information, people who lean right prefer right-leaning biased information. You give a crowd of people truly unbiased information and nobody will like it, right?

Next cartoon. So of course we'll talk about this in the panel, the killer robots program, Skynet from the terminator movies or Hal 2000 or whatever. A very common

Hollywood trope is the AI as evil, so technology run amok, technology causing problems outside of the control of humans.

This, unfortunately is one of the primary themes in the public discourse about AI. AI comes up, you see it on TV and it's like the first thing they ask is, what about Skynet, and this is sort of distracting from the real problems and the real promise of AI. So it's a big distraction more than anything else. But it's something that we have to deal with, being professionals in the industry, is how do we help our customers avoid the distractions.

Okay, next cartoon. [Unclear]. So this one illustrates invasion of privacy and actually it's interesting that Alexa was just in the news, I don't know if you saw this, that Alexa - somebody's Alexa, somebody's own personal Alexa in their home recorded a private conversation and sent it to an acquaintance of the couple having the conversation, unbeknownst to them.

What had happened was, something in the conversation sounded like the word Alexa, which is the word that indicates that Alexa should start listening and then something in the conversation sounded like, record a message, and then something in the conversation sounded like, send a message. Alexa then responded, they didn't hear it because they weren't paying attention. Who should I send the message to, and then something in the conversation sounded like somebody on their contact list.

It's a very unlikely sequence of events but it still violated their privacy. So generalise that, how do we know that AI isn't going to do this? How do we know that AI isn't going to do this in a corporate environment where a sequence of unlikely events leads to a severe compliance breach or transferring money without a human being touching anything? It could be any number of different problems where it's just a sequence of unlikely events but you have enough events and something will happen, right?

Okay so you knew Dilbert was coming, I saved Dilbert for later. I'll give you a moment.

Of course the big boss, the dome headed guy, is using machine learning to track customer preferences and use that knowledge to manipulate them. So this is the whole Cambridge Analytica problem. AI not just doing evil things on its own but AI in the hands of evil people doing things that are evil, and now it becomes a more powerful tool for manipulating people and for causing all manner of issues. This is also a big concern and a very real world concern given some of the news today.

Okay well that's my five minutes of cartoons, so let's go ahead and have the panel. I guess we'll start with some brief introductions, and my questions here on my phone. Why don't you have the panellists each introduce yourself, say a bit about what you're doing with AI and what do you think is the strongest use case for AI today. So pick one and what - in terms of the real world, not the hype, what is really going on in AI today.

Let's start here and work down.

Vinod Peris

I'm Vinod Peris, I'm the Senior Vice President of the Central Software Group at CA Technologies and regarding your question, Jason, the enterprise really has always been

about data and at CA we've been working many years with our enterprise customers, we have a lot of data and we use data to make decisions to measure ourselves, our metrics.

But if you think about it, data has typically been something like hindsight, it's something that you look - after the quarter you look at your results, or if you are doing your program management you look at milestones that you've missed. Or you have static thresholds that you're measuring your operations against.

Really what we are doing with AI is changing this to be more predictive. We're looking not just at what you've missed as red flags, but alerting you that you're likely to miss this milestone in your program.

Or if you're monitoring your data centre, if it's systems that you're monitoring, what we're looking for are anomalies in there, not static thresholds. Before you would look at your CPU or garbage collection or things like that, but as now, it's more what's the behaviour of the users and the pattern of the day and what do you expect, and are you seeing some anomaly from what you expect?

Another area where we're using AI is really in payment security and I'm sure many of you use credit cards and are really pissed off when your credit card company denies your transaction. But really it's the bank that's trying to balance the friction in your payment and their risk. By doing behavioural analytics, you can remove that friction just because you know what to expect without increasing the risk for the banks. So these are all the aspects in which AI is really changing the game for the enterprise.

Jason Bloomberg

Very good. Sam?

Sam Liang

My name is Sam, I'm the CEO of AISense, we are a speech recognition technology and product start up in Los Altos. I totally agree with Vinod in terms of data. We focus on the voice data. So if you look at how people communicate, while you're spending two days here in this conference, most of the time you are either speaking or listening.

When we started the company a few years ago I was thinking, I can search my email 10 years ago within seconds, but I can't search anything I heard this morning, I keep forgetting things. So we see this as a huge opportunity here.

In enterprises there are all kind of statistics. Some said that 30 per cent of all the employee's time is spent in meetings. But if you are a higher level manager you probably spend more than 50 per cent of your time in meetings, phone calls, video conferences. Where is this data? People have to take notes diligently, you have to type, you have to write on paper, notebooks. You can't search anything on your paper notebook. When you type, yeah it's good, but it takes effort.

We are trying to automate this process. We say, hey why don't we just use speech recognition to help with this? In addition, because [you] spend so much time in meetings, every hour is worth thousands of dollars, all this money is wasted if the

meeting and the information is now saved. So we see this actually as tremendous value to save all the voice communication. Whether in person, over the phone, or over video conferences.

One example we saw was Bridgewater Associates, probably the biggest hedge fund in the world, managing \$150 billion. The founder, CEO Ray Dalio, said, we have to use his radical transparency principle, we record everything. I guess that's one of the reasons they are so successful.

They actually came to us, they said we have recorded 15 to 20 years' worth of meetings, but we can't make sense out of this. So they love our technology, they invest in us and they want to use this. They analysed a huge amount of voice data there and extract the insights out of it. Moving forward they want to use this for every new meeting as well.

We're building technologies to understand human to human conversations which is different from Alexa, which is human to machine communication. As I mentioned, people spend much more time speaking to other people than talking to Alexa, so we want to make sense out of that.

Jason Bloomberg

All right, Greg?

Greg Fitzgerald

I'm Greg Fitzgerald with JASK and we're looking at artificial intelligence being applied for better situational analysis. Basically the overabundance of data that a human physically cannot interpret on their own and applying - like I teach my daughter, good decision, bad decision, simple as that.

We're applying it in such a way that it just is enhancing and augmenting the intelligence that a human has to then make this subjective decision on whether they need to act on that information and it's starting to change the world to enhance what we do as people.

Jason Bloomberg

In terms of some of the challenges that AI faces, what are some of the key roadblocks? They could be, you know, on the human side the sort of scepticism that we talked about in the cartoons, or it could be technical roadblocks as well. What do you see are the limitations to the progress? This really could be anybody.

Greg Fitzgerald

If I could take a quick one, in my belief AI is like a child, its learning. There's supervised learning and unsupervised. If you look at, again my 11 year old, I taught her what a bus looks like so she can get on the school bus. It's big and long and yellow and has big wheels, that's supervised. But when she goes out to the street and there's a selection of buses, could be a city bus or other, she has to make that decision, hence the unsupervised.

AI is still developing in terms of its intelligence. I think the challenge is still feeding it the information and trying to help it to learn the right way. Because in the cybersecurity

world, AI is, to your point, being used by bad people to get ahead of the good people and in my world I want the good people to get ahead of the bad people. So it's who can train the best, and again in my world a lot of it has to come down to a moral compass of decision making that is applied through that technology, so there's a lot of learning that still has to go.

Vinod Peris

I'd like to add a couple of things. One is regulation, especially if you look at things like self-driving cars. We all know that it's going to be a better driver than us, at least better than me. But it's going to take a long while before you convince policy and governments to allow that, because every little fatality that you see out there is going to put a big roadblock and slow down the innovation.

The other thing is, it's all about data. The more data we can collect the better it is. Another - one of the aspects that will slow things down in collecting data is the concern around privacy. Now it's not like I'm not a privacy advocate, I don't want to go there. But the point is that many simple things like we can do like usernames and IP addresses and stuff like that now will be much, much more restricted because of GDPR and PII and this will also slow down our ability to advance in this area.

Sam Liang

You should look at AI relatively. People say what is AI? AI is anything that hasn't been done yet. We go back 50 years ago, everything we have today could be considered as AI 50 years ago. Once it's done it's not attractive anymore. You can do OCR pretty reliably today, but 10-20 years ago it's still a fantasy. Same as speech recognition, we can recognise most of the words today, but still not perfectly.

Regarding privacy, for our product privacy is one of the major challenges. People have the psychological resistance to being recorded. Are you going to use this against me some time later? Are you going to sue me some time later? Also reliability in the self-driving cars and - statistically I believe even the current self-driving car is already much more reliable than human drivers. It doesn't do DUI, it doesn't fall asleep. There are still bugs but I believe it's already killed fewer people than human drivers.

Regarding the Alexa incident, it's a bad privacy incident for Amazon but if you think about it, if you have a human secretary, she or he could make a mistake and misunderstand your instruction as well. So which one is more frequent? Of course whenever the machine makes a mistake people question AI. It's a psychological or philosophical question there.

Jason Bloomberg

Almost 100 years ago, Alan Turing sort of set AI onto a particular course with the Turing test, a test which, ironically was designed to mimic humans, but that wasn't the point of the test. If you go back and see what Turing was actually trying to do, he realised that there was no real way to test for human intelligence or human-like

intelligence in a computer. So he came up with the Turing test, something testable, you could test whether or not a computer could fool people.

Ever since then, this question of whether or not we want AI to mimic human behaviour sufficiently well to fool humans into thinking that AI is human, has now been a bar that is set but it's not really clear whether or not we want to achieve that or not.

But you see that out in the - especially in the consumer oriented world. Alexa is trying to be more and more human-like. Even though we realise when we're interacting with Alexa or one of the similar consumer oriented products, that they are - we know we're interacting with a machine, it's still more and more human-like and we now have this human nature response of reacting to something that is similar to a human in a human way.

The same is true of just robots, even if they are not particularly intelligent, if they are soft and cuddly then we treat them like animals or children even though we realise they are machines. So this is one of these challenges.

Then you see products like Sophia, I don't know if you've seen Sophia, it's this human-like robot that has some rudimentary AI and has got a lot of press recently.

Or even more recently is the Google - what was that called? The Google...

Vinod Peris

Duplex.

Jason Bloomberg

Duplex, Google Duplex. Where the AI was able to interact with a human and set an appointment with a human being unaware. The question is, is this the right way for AI? Should we be trying to mimic human behaviour? Or as the products - you folks are talking about predictive analytics and voice transcription and cybersecurity, pattern detection. Those aren't necessarily things that humans could do well, those are things that AI can do better than humans.

So which is it? Do we want to have AI be more like humans or do we want to have AI solve problems that humans aren't particularly good at solving?

Vinod Peris

Yep so I can jump in here. I think there's nothing wrong with mimicking humans. If you think about it, if you want the machine to interface with a human it's only natural that they do it with our interfaces which are speech, listening, natural language and so forth.

Where I think we are crossing the line from an ethics perspective is, you don't want to use it for deception. You as a human being, want to know when you're talking to a human and when you're talking to a machine. I'm sure many of you have been pissed off like me when you get these robo-calls, because they're getting more and more sophisticated these days, they try to act like they drop the phone in the middle or something like this to just hook you in, right?

So I think that's where I would draw the line. You don't want to deceive. Same thing with Duplex, I think a couple of criticisms that I've seen is, one they didn't announce up front that this is Google Assistant or something like that, and two is they were mimicking some of the pauses like mm and things like that, which human beings do, which seemed deceptive.

I think you want to draw the line at being deceptive. Otherwise just using it to interface with humans is perfectly fine.

Sam Liang

The key question is actually - Elon Musk asks the question, is AI going to destroy human beings? I think my view is it's still a question too early to worry about. The goal is actually not to replace humans any time soon, it's actually to augment human intelligence.

With machines - well with cars we run faster, with speech recognition and other technologies you take notes easier. There is research done that shows that, with image recognition for medical radiology, if you look at the accuracy of the human doctor and the accuracy of the robots, a human doctor is still slightly higher than the robots. However when you combine them, the combined accuracy is higher than either human, or the robot. That's a good example of having AI augment human intelligence.

I'm very optimistic. I think with robots humans will become more efficient and more productive, life will be more enjoyable. I was chatting with [Jerry] Jerry was saying that [David Cheriton]'s view is that AI is eventually going to save human beings from messing up. All the machine learning will try to learn the best way to train, to build new robots. So it's not the other way around, it will actually help human beings be better.

Greg Fitzgerald

I agree with Sam.

Vinod Peris

I would like to just counter a little bit of what Sam was saying. I don't agree with Elon Musk either, I don't think you have to worry about killer robots.

Jason Bloomberg

That was the next question.

Vinod Peris

I think it's - robots are - AI is making us dumber, and that's what I worry about. Let me give you an analogy. If you look at GPS, and I have to admit, I'm woefully directionally challenged. But in the last five years I've gotten much worse thanks to GPS. With AI it's not inconceivable that you will become much worse.

Even if you take the case of the Tesla accident on 101. What happens is, when you're on autopilot it gives you back control at a moment when you're least equipped to deal

with it. This is the challenge of AI, I think it will kill us in this sense. Not so much maliciously, but by making us so dumb that we won't be able to take charge of a difficult situation.

Greg Fitzgerald

Actually to that exact point, driving with my boss, Greg Matin, and he has a Tesla and we were driving down to this event and he just, in the front seat, turns around and talks to me. Not even paying attention to the road. I jumped, reactively, like what the hell are you doing? Then he gets on and starts texting. He'd put on the autopilot, the car is slowing down, it's shifting lanes and I was like, this is crazy, because I had not experienced that before, and he said yeah the problem is if I get in a real car, I forget to turn the blinker on.

[Laughter]

Greg Fitzgerald

I forget to stop. So to your point, I agree with you Doctor, that's exactly what I am kind of concerned about, the natural human behaviour is to look for shortcuts, our brain is always looking for shortcuts. AI is now a shortcut, which then could make us, in a generation or two, literally not understand the situation we're in. Then if I need to drive a manual car - great example in Europe, I've got to rent a car, it's a manual stick shift. But everything in America is automatic, I've got to go back to remember how I drove a car in the first time, to be able to drive that car in Europe.

Jason Bloomberg

There was an Air France crash five or 10 years ago. The airbuses have been so automated that the pilots have very little to do, it's just basically push a button and the thing flies itself. They got disoriented and it led to a crash because the pilots weren't flying the airplane, it's the same problem. I guess there was AI in that story, although it's five or 10 years ago.

I think we're about out of time. Oh questions from the audience, let's have questions from the audience. Yeah we have a few.

Audience Q&A

Marcelo Lozano, IT Connect & CNN

My name is Marcelo Lozano, IT connect and CNN. What can they tell us about the ethical of algorithms?

Jason Bloomberg

The ethics of algorithms?

Marcella Lozano

Yeah.

Jason Bloomberg

Okay, interesting.

Vinod Peris

I could take a crack at that. More in terms of - think about bias, when we think about machines we think about algorithms, we naturally think that they're perfect, they're unbiased. This is something that is a cautionary tale for us in this business.

If you look - I think it was May 2016, ProPublica had an article where they were looking at the US criminal system and they have something called a risk predictor. Which says how likely is this person who is coming up for bail to commit another crime? They've been using this to decide how much bail to set or what to do. They went back and looked at the cases and looked at the statistics afterwards and found that this risk predictor was racially biased.

It's something that we have to be concerned about with algorithms as well. Ultimately they are made by humans in some cases, and two they are trained by data that humans are collecting. Both of these can result in ethical challenges in biases.

Sam Liang

Just to quickly add to that, everybody's - your belief, my belief, it's just a function of the data we got before, the things we heard, the things we've seen, so machines are no different. Based on the data it has seen, it forms its belief. So to make it more ethical you just need to give it better data.

Jason Bloomberg

Let's take another question because we're almost out of time.

Hector Pizarro, DiarioTi

Hector Pizarro from DiarioTi - I have a question for Sam. Would the output of your speech recognition solution be understandable for IBM Watson for instance? Because Watson's strength seems to be its ability to make sense of unstructured data. I think Watson can make sense from a meeting recorded and transcribed by your solution.

Sam Liang

Actually, I didn't totally get the question, how do we compare to IBM Watson?

Hector Pizarro

The question is, whether your solution would be understandable, the output of your solution, would be understandable for IBM Watson and to make sense out of a recorded meeting.

Jason Bloomberg

If you feed the transcription into Watson.

Sam Liang

Yeah we haven't done direct integration with IBM Watson but I believe that's doable. It's actually a big challenge for us, or a major direction we are going into, to make speech recognition more personalised, it's more context aware.

Right now, whether or not you use IBM Watson or Google Voice API, they just have a generic speech recognition model, it doesn't know who you are, it doesn't know your history. However the language, when you're communicating with somebody, if you know some background of that person you can understand him much better, even if you don't hear every single word, you can infer a lot of things based on historical conversations.

I see that as a huge challenge and a huge innovation there, to make speech recognition much more context aware based on where you are, when I'm talking here versus what I'm talking about at home, I'm talking about totally different topics. But if my cell phone understands the location context, it can actually change model dynamically. Both the acoustic model and language model.

Also if I recognise your voice I can look up - the robot can look up your LinkedIn profile, can look up your Facebook, your Twitter. Even before you start talking for five minutes, the robot can already infer a lot of topics that you are going to talk about. That can help the speech recognition system to be more accurate and to understand you better.

Jason Bloomberg

Do we have another? Are we done? I guess we're done. We do have more questions but I guess there's plenty of networking time with our panellists so I guess we'll have to leave it at that, right?

Manek Dubash

Yeah sorry, we are out of time.

Jason Bloomberg

Okay will do. Thanks to everybody on the panel.

[Applause]