Raising the Bar Sydney 2018

Lamiae Azizi – Artificial unintelligent

Welcome to the podcast series of Raising the Bar Sydney. Raising the Bar in 2018 saw 20 University of Sydney academics take their research out of the lecture theatre and into bars across Sydney, all one night.

In this podcast you'll hear Lamiae Azizi's talk Artificial unintelligent. Enjoy the talk.

[Applause]

So, thank you everyone. I'm going to stand up here because I can't see you all. So, I usually start my talk saying that I'm very honoured to give a talk somewhere.

I'm very honoured that you've given your Wednesday to come and listen to me, but I'm finding it a little bit challenge because this is the first time I'm going give to give a talk without my formulas, without my computer, so I feel like thrown in the sea in a sense, right.

So my talk tonight is going to be about artificial intelligence. And I always envision my talk in a way, but just before coming here, my mum called me. She's a very supportive and loving mother who always calls me before talks, but she always gives herself the right to give opinions about my talks. Well, she's not a scientist right? So maybe we all have that kind of mothers in a sense.

So she had one statement and one question. Her statement was, well, you're going to, I can understand, you're a mathematician and you're going to talk to people about your research in a pub. How did we end up here?

Anyway, so, she did overcome the shock of it. The question was an interesting question. She said, you're going to talk to people about artificial intelligence or artificial and intelligence, but what does that have to do with your research? You're doing research in statistical machine learning.

So obviously her statement was, why did you give yourself the right to talk about something that you maybe don't understand? So that gave me insight about how to reshape my talk in a sense.

Alex, as well, who is here, who works we me at the Centre of Translational Data Science at Sydney Uni told me the same two days ago about how people perceive artificial intelligence or what they think about artificial intelligence. So I thought that maybe I should start with a little bit of history around artificial intelligence before moving on, because it seems like there are different views around artificial intelligence.

But before that, I'm going to tell you I think that we're going to spend like 25, 30 minutes or a little bit over in this intimate place, nice. I'm going to be honest with you. I'm going to talk about very cool things maybe, I think are cool, maybe you will find them cool, but I have one intimate goal, during all my talk is, [inaudible] your brain around maths.

My head of school is here. Thank you, Jackie, for giving up your evening, but this is putting more pressure on me. Because that means that probably I will lose my job if I'm not doing a good job. [laughter] So you've had a beer, that's good.

So what I'm trying to make you understand are the core of even the most like fancy or cool technology there is maths basically. All these algorithms that you hear about, mathematics is at the core of it.

So I think regardless if you like math or you don't like it, it's important that you understand the minimum of it if you're interested by this technology or if you want to be part of this AI evolution in a sense.

All right, I [inaudible] to give this example. If you just use the technology as it exists, that makes you a carpenter of the school. If you understand the mathematics behind it, at least the minimum, and you can manipulate it, then you become the, you have the skills of the master carpenter. So choose where you want to me, all right.

Now, let's go back to this question by my mum about statistical machine learning and artificial intelligence. So if I asked everyone here, what was artificial intelligence, would you all agree with me that this is computer system designs to use that data to teach themselves?

How many people agree with me? Everyone? No? Not there? Nobody? Yeah, okay.

So I think of artificial intelligence as the broad concept of systems designed to carry out tasks in a way that we can see they're smart. Right. Okay. No, okay. So, no? Nick? Okay.

So let's go back now. Is artificial intelligence this new hype? Is artificial intelligence a new thing or it existed for many years? What do you think?

Many years.

Many years, right. So if you like mythology like me, you will find that even in the Greek myths, there were mechanical men designed, do you know that? That they were designing mechanical men that can mimic our own behaviour.

And if you consider the history of the early European computers, you will understand that they were perceived in a way that they can mimic arithmetic's basic and memories, so in a way that they can mimic our behaviours.

Artificial intelligence has been there for a long time, but what happened? Why did it give back to us?

It goes, so our understanding of how our minds work has progressed over years, all right. So we change it.

So rather than considering AI as just complex calculations, the work in the AI field moved towards more like how we can teach machines to learn themselves and not teach them everything we need, they need to know about the world we live in and how we understand it. Okay. So that makes it a bit more complicated.

Let's just go back to AI. AI has two, I classified it in two classes, broad classes, the generalised AI and the applied AI.

So the applied AI is the systems that are designed to trade intelligently in a sense the stocks or shares or the systems that are, we use them to manoeuvre autonomously, for example. The generalised AI is what I'm going to talk to you tonight about.

So the generalised AI is more the systems that are, so the generalised is more a broad concept where the systems can use the data to teach themselves.

And to get there, we've seen the new thing coming out from the generalised AI, which is machine learning. All right.

So now we understand, machine learning is a part of the generalised AI, and this is the new, so do you all understand what is machine learning in a sense? No? Jackie says no. Okay.

She says no.

Okay. So who knows what is machine learning? Yeah? Nobody there? Okay.

So machine learning is just people tend to think of it as a subset of artificial intelligence that use data in a sense to, a subset of artificial intelligence that thinks that we can just use data to let the machines learn by themselves.

So it's a bunch of algorithms that we've designed to find the patterns of the data, right, in the data, and try to make the machines teach themselves how to understand things around the world. Okay.

So this is what I can talk about artificial intelligence, because machines, the most coolest things are happening there, in the machine learning basically and not in the applied AI where we talk about robotics and all of these things. I'm not going to talk about robotics because that's not my field. It's part, it's part of it. The techniques and algorithms that we design for machine learning are used to teach the robots how to learn and interact with the world in a sense.

Okay, so how did we come to this machine learning. So this was a realisation. The [inaudible] in machine learning started with a realisation by Arthur Samuel.

Do you know who is Arthur Samuel? Arthur Samuel was a pioneer in computing gaming. So his realisation was that instead of teaching the machines how to, how to understand everything around the world, we can just teach them to use that data and to think as a human being.

So the machine learning is about mimicking the human decision making processes in a sense. So we use that data, the data is information, any kind of information, everything around us is data. So we use this information to teach the machine to mimic decision making of humans, how we tend to make decisions if we have facts around us.

Does that make sense now? The difference?

So people think of machine learning as a subset of AI, but in fact, I think of it as the con state of the art, because this is the force driving the developments that we have in AI.

So if AI is progressing now it's because of the machine learning. So the machine learning is based on algorithms.

The first developments in machine learning was the neural networks. Everyone hears about deep learning, right? Yes, yes? No? Yes.

So, deep learning, the neural networks are just the systems that are driving or supporting the deep learning systems basically. And they are algorithms, or they're metas if you want to think about them like that, they are a computer system designed that works based on a system of probability. Maths, right.

So the system of probability, when fed with the data, it tends to make decisions, make predictions, make statements and by feeding, by adding to it a loop of feedback, it can learn about things and decide what are the approaches because they are said if they are wrong or right, they, and they decide which kind of approaches they need to take in the future. So this is how we can make predictions in a sense. They improve.

Okay. So that's the machine learning. Now we understand, artificial intelligence now is driven by machine learning developments, and machine learning applications are everywhere.

They can read texts and tell you, decide if you're happy or not happy, classify you as happy or not happy. They can read pieces of music and decide to give you advices about what kind of music you should listen next to change your mood, if they detect that you are not happy or they classify you as unhappy person.

So they are everywhere. They are driving the natural language processing, where we process texts and we try to understand how people interact with people through the texts in a sense.

So they are everywhere, and in my, in the advertisement for this talk, it was put that I'm going to tell you that the machines are not smart yet because they are not equipped with the mathematical reasoning.

But I just told you that the neural nets that supports one of the big or the major impressment is deep learning are based on a system of probabilities, or basically there is maths there, right. So everything is maths in a sense.

But first let me tell you, what made this deep learning, or maybe you know, well made this deep learning coming to the surface again. Neural networks have existed for a long time, but they were buried in a drawer for a little bit, and then they came back.

So do you all remember when Google deep mind a few years ago, like two years or--sorry, two years ago. Two years ago. What did they, what did they tell us?

That the algorithm have bet, 18 time world champion in the ancient game, Go, the ancient Chinese game Go. Do you know what is the game Go, right? Right.

So why was that impressive coming back? Why did we, why did AI at this stage amazes people? It's because this game was considered for a long time as the most ever complex game that humans have invented, because it has more possible move configuration than the number of atoms we have in the universe, and that's combined with the complex strategies and the complex mind boggling around it, so it has always been considered and for a long time that it's going to be just something that humans can achieve. Machines will never beat them, right?

So from there, it started to be, this machine learning algorithm started to be amazing people. AI came back to the surface with that, and people started inventing and ramping up their games.

So China last week announced that they are going to invest in 2.41 trillion Australian dollars in AI because they want to be leading AI basically. Because AI is going to change our lives, is going to make, is going to start making decisions for us in every single domain of our lives, so soon we're not going to be having a doctor because a machine will decide for us if we need to get operated or not.

So AI is going to change the world, right. But despite the success that we observed in Alpha Go, we had also some failure, a major failure it came to the real life.

So if we consider like real problems outside the gaming imaging where we can just classify the images based on the elements they contain where machine learning algorithms really excel, we have a number of failures that showed us that we're not maybe yet there. The machines that we thought they were intelligence, they probably are not that intelligent.

So, for example, Google flu trends, have you heard about that? This is where Google tried to build a kind of, an algorithm to estimate the number of influenza-related cases based on queries of people and on Google Search around influenza and their visits to the physicians.

So despite it's initially success, and it had an initial success, a huge success, but despite that, we realise that they over estimate by one factor to two in subsequent years when we had a look at the number of visits related influenza visits.

So clearly there is a problem. What we told to have, what we told to be a successful revolution that would solve all our problems, it appears to be having, still having it's major problems.

You're listening to Raising the Bar Sydney 2018.

>> So what went wrong then? If we look at the information that this machine learning systems use, it's most of the time based just on statistical information that we have in the data.

So what does that mean? In other words, this algorithm tries or this methods tries to optimise over a space of parameters that they receive just, or they try to optimise parameters over just inputs that they receive from the environment.

So why is that a problem? Okay. So this is where I want to tell you a story about the long process learning.

Do you know who is Judea Pearl? He is one of the biggest computer scientists, a professor the UCLA, who got the Turing Prize. The Turing Prize is the high distinction in computer science. People usually refer to it, sorry Jackie, it's like the Fields Medal, the Nobel Prize of computing, because everyone who doesn't have a Nobel Prize needs to say that this is equivalent to the Nobel Prize of the Field.

So Judea Pearl tends to just take an analogous, like explanation of this machine learning systems and how they learn, and he tends to think of it like a long process, which is very similar to how the natural selection process drives the Darwinian evolution, right.

So it can explain how over millions of years species like snakes and eagles have developed supervision systems, but it cannot explain how us as humans have developed glasses, telescopes over just a little bit, this is a super evolutionary process, just over thousands of years. Right.

So the difference between humans and species is the fact that humans have this [inaudible] of their environment, right, where they can distort [inaudible] their environment by mentally imaginations, acts of imaginations, where they can imagine hypothetical environments where they can learn and plan.

And this is the difference. So our machines that we are just building now cannot reason about experiments or actions. They reason just from the observations we have.

While humans, if we think about how we make decisions, we just don't make decisions based on the facts that we have in our hands. We make decisions based on experiences, based on asking what if, and this is the question, what if I act, this the questions about actions. The other questions is what if I had acted differently? It's retrospective reasoning. Okay.

So when we start imagining these kind of questions, then we start having hypothetical scenarios and hypothetical outcomes that could happen in a situation. This is what our machines are not able to do at this stage. They are just observing things.

So what does it mean? This is probably an observation that we need to take into account and equip our machinery with that kind of reasoning.

And it turns out that if we go back to philosophy and even mathematics, we will find something called causality reasoning, or causal reasoning. Do we know what is causality? Causation or not correlation, right.

To mention [inaudible] has a very nice explanation of causation on his human logic, have you seen that one, human logic? Where he talks about causation, and he says that as humans we tend to assume that when we observe an event be happening, [inaudible] event A, that event B, that event A causes event B, so this is causation in a sense. More than just correlation.

So now what is the thing about causation? Causation can answer questions around what if I act, even if I have not acted right, and it can even answer questions around what if I have acted differently? Which we cannot do in our observations.

Imagine a patient, right, going, we give him or her a treatment, and then we observe the outcome of this treatment. I'm not going to tell them to come back again and ask them to act because we didn't understand the outcome or it wasn't what we wanted to do, to come back and have a different drug. It's not going to happen, right?

So we need to have an environment where we can observe, we can learn from the data that allow it as well to have this hypothetical scenarios and be able to have hypothetical outcomes and things about how our decisions could change if we had run the experiment differently or we had collected the data differently.

So it turns out that what is missing that we can overcome the limitations of our machines not able to act intelligently but equipping them with the causality reasoning.

So here is what theory and the maths tell us. The theory and the maths tell us that if we follow the causal reasoning, the theory tell us that there are three classifications about the logic of how we can use causal information.

So there is a short distinction between these three classes, and it is a hierarchy with three levels where it's a one-directional hierarchy, it goes from the top to the bottom and not in other sense, and the top is the most powerful, is the most powerful layer, where it can answer such questions like what if I had acted differently.

It turns out that when we look at these three layers from the theory, and we try to map the machine learning systems that we, the current state of the order we have developed until now, we find that most of these techniques or most of these algorithms and metas lay in the association level.

So despite the very impressive achievement of deep learning systems in some applications, we are still reasoning in the association level, which means that we're still doing exactly what other techniques that just operate at this level.

So we're going to just answer questions that are like, what at the same time tell me about a disease, what a survey tell me about election. It's not going to answer us questions about interventions, which are at the second levels, interventions like what if I take an, if I take an aspirin would that cure my headache. Or these kind of interventional questions, right.

And the third level is what we call the counterfactual level, and the counterfactual level is have you ever heard about counterfactual? No.

So counterfactual is a philosophical approach to causation that was introduced by Lewis in 1973, who say that instead of having the regularity of the analysis, it's what he calls the promise in alternative analysis.

Okay, so the counterfactuals appears to be as the key success of understanding causal facts. What does he mean? See, the cause happens, if the effect happens, right. But if C had not happened, E would not have happened.

So this is the counterfactual. And this seems to be a great tool to answer questions such as what if I had acted differently? Okay.

And all of this seems to be philosophical, but this isn't really maths, it's all probability statements that we write and then we try to derive and try to understand. So what is missing for our machines to go to the next level is this causal reasoning tools that need to be implemented in addition to just the usual way of doing our machine learning stuff. Okay.

So this is what is missing for now.

Now, imagine we've had this causal reasoning tools imbedded in our systems. That means that we can answer questions that are, we can have hypothetical scenarios. We can answer hard questions. We can make machines think, or these algorithms that are the driving forces of the artificial intelligence systems, we will make them think as humans think in a sense.

In the last few years, machine learning has become a bit obsessed with something called ethical, the ethics of artificial intelligence, right, or the ethical algorithms because machine learning tends to call everything algorithms even if they are really models. Just because it turns, because it's driven by the computer science community, so we need to call them algorithms, because it makes more sense than models.

So ethics, ethics. So why did we become very obsessed with that.

In 2016, I was in a conference in Barcelona, and this conference is the neural processing information systems. This is one of the biggest conferences in machine learning and artificial intelligence. It covers around 6000 people from around the world [inaudible] and academics and--and Boston Dynamics group. Do you know what is the Boston Dynamics Group, it's a robotics group. So they tend to create the new generation of robots.

But these robots are not yet the robots that learn and interact with the world, but they are robots that are at this stage just, they teach them how to stand. They teach them how they can really walk. They teach them how to grab things. So it's control based at this stage.

It was very amazing to see that they brought the robot there on the stage, and they were showing us the capabilities of these new robots. We were all impressed with that, really impressed, because that's really a major advancement in the field of robotics. And we thought that it would never happen, to be able to see that kind of robots like really standing there with these capabilities.

But, at the same time, where everyone was excited, standing on the chairs, trying to have a look and see what's happening with the robots on the stage, I think everyone got scared with what we were saying. We were seeing something that was really unbelievable, scary. Like this is the people we're going to be living with maybe in a few years. Or maybe not in our generation. I tend to think we're not going to see it. Maybe other people will see it, our children or, but we never know.

So this is the kind of, and this, so now what's going to happen that we will buy these robots, and people who work on statistical machine learning or machine learning will try to implement these algorithms that we implement in other ways in these robots to make them interact and behave like humans.

So on Monday, I was talking to someone who said, oh my God, I think that robots will take over the world, and while I can think of that they will be taking care of me when I am an old person. I'm like, yeah, that's probably not going to happen that quickly, but anyway.

So this is something that we are all scared of, and this is where we started realising that these algorithms that we build from the machine learning, I mean we started realising it a little bit before, but that scene in itself made a lot of people who work in the field who are driven by the science fiction ideas, because the people who work on machine learning or artificial intelligence are really driven by science fiction ideas. That's what we think of. It's like, how can I make something really cool, and you know, things that will maybe never happen we could implement them and make them work.

But even these people realised at this stage that there was something that we need to do around this machine learning algorithms. So let me step back to this machine learning algorithms and explain briefly how they work.

We call them black boxes. Forget now about the causality. This is what we should be doing for the next generation of machine learning algorithms, to be able to make them think and reason in a more human way than they are now.

But think about them as now how they work. They are just black boxes. You give them an input, any kind of information you are interested in taking something out of it or understanding something about it, something happens inside, and then you have an output.

It's not completely 100 percent true, but this is the, like the, how to say this, sarcastic way of presenting the machine learning technique. So we call them black boxes. Because all we're interested in is the prediction. We don't understand why, how. We don't understand what happens inside, in particular the deep learning, the architecture is very complicated, but we don't really understand how we got to the prediction.

So people started worrying about, ethics is just a big word that they use, but it doesn't mean the ethics maybe from the social sciences we think of, but it's more about how can we make these models or these algorithms more [inaudible]. We understand what happens there, because at this stage we don't understand. We have an input, and we get an output, and then we see the prediction, and we're happy with it.

So when we started hearing about things called explainable AI, interpretable AI. Have you heard about that before? These terms? Okay.

So this is exactly why people started thinking. Now we get this prediction. We don't know what happened inside. What I'm going to do now is opening the black box and trying to understand what is happening inside.

So instead of just understanding how we got to the--instead of understanding the prediction, we're going to be trying to understand how we got there. What is the path that we follow to get to that prediction?

I tend to think that with the causal reasoning that will be solved, but the community has moved towards trying to build new ideas of how to justify, explain why when we give these machines an input it gives us an output.

We started thinking about that because these algorithms or these systems will be used to decide about our lives now, and many places around the world, governments or business, want to use these algorithms to make decisions around of, you know, deligating the task of decision making to the machines.

So if we don't really understand how we got to the prediction or how we got to the result, the output, it's a problem, right? Because we just make decisions based on opaque things.

So, we started thinking about interpretability, justifiability, and there are other problems that came up.

So we are in the era of big data. Okay. So we all know that. We're all excited the big data, maybe not everyone, but--I don't like the big data name. But anyway.

So, we're going to be collecting information from everyone. We already does it, about everyone. So we're all excited that what's going to happen, we're going to move from average, you know, decision for, decisions based on the average individual to more solution tailored to personalised, to personal solutions.

So what does that mean? For example, in health, one of the big challenges is what we call the personalised medicine or personalised healthcare. That means that we're not going to be all taking aspirins for if we get a headache, but we're going to be having aspirin A for you and aspirin B for me, because we're completely different based on a number of characteristics that are maybe your genomics, your genetic coding, my environment, and other things.

So we're going to be targeting solutions, and we're going to be using these algorithms. There are problems. It's not just the justifiability and interpretability that seems to be something that we could solve in a sense or another.

But now we get onto other problems, privacy. Privacy is a big problem. Are you worried about the privacy of your data? Yes, absolutely.

Particularly in health. Just like I hear people say, you know, I'm happy for you to take my data and do research and try to understand diseases and cure cancer, but be careful. I don't want you to be using my data and give it to insurance company to give me personalised suggestions about how to ensure me. All right.

So this is the kind of things that with personalization it might happen. So privacy is a big problem. And they are not going to tell you that privacy is, privacy has not been solved before or that's a new issue.

Sorry--it's clear that I'm not used to being on the stage.

So I'm not going to tell you that that's new, but what's happened, before, the privacy, from the theoretical computer science community, people were talking about differential privacy. Have you heard that before?

They were saying just [inaudible] your data a little bit. Don't give me your data, and I'm going to be able to answer any queries related to the data because we have theory that, that will, that ensures that even if I don't see the raw data, I can just give you the right answer.

The problem that it turns out that with this big data and high dimensional spaces from the math side, it's very hard. This differential privacy doesn't, it breaks basically. It doesn't work anymore.

So privacy is an issue. Interpretability is an issue. There are other issues around also how to trust these algorithms for the deployment all of these problems.

So the community is very active in that. I'm not going to tell you that we're going to solve all the problems. They will need a lot of involvement from--I think the problem is not about stopping the artificial intelligence or machine learning developments, but we need a little bit more regulation as from the governments and more than for us in a sense.

I tend to think that we're going to be anyway going towards AI, human-like machines because it's going to, it's becoming faster, and people or the progress in the machine learning is going very fast, and people are really very excited about how we can be making the machines carry out the tasks that humans are not going to be able to carry out at this stage to make decisions just because of these big data problems where we can't be able to process all this amount of information and understand what's happening.

So, it's going to be happening. I'm excited because, because, because it gives me a job, right. I can, I can be useful for something, but I think what we should remember is just that people are working hard to make this evolution and using it for the better. It could go wrong, but if you do math, it's not going to go wrong, because we will be able to work out and at least understand when things break when they break. Thank you very much.

[Applause]

Thank you for listening to the podcast series of Raising the Bar Sydney. If you want to hear more podcasts from Raising the Bar, head to raisingthebarsydney.com.au.