



Artificial Intelligence and Augmented Intelligence for Automated Investigations for Scientific Discovery

AI3SD Interview with Professor Andy Stanford-Clark
18/11/2019
AI3SD Conference 2019

Michelle Pauli
Michelle Pauli Ltd

27/04/2019

AI3SD Interview with Professor Andy Stanford-Clark

AI3SD-Interview-Series:Report-1

27/04/2019

DOI: 10.5258/SOTON/P0022

Published by University of Southampton

Network: Artificial Intelligence and Augmented Intelligence for Automated Investigations for Scientific Discovery

This Network+ is EPSRC Funded under Grant No: EP/S000356/1

Principal Investigator: *Professor Jeremy Frey*

Co-Investigator: *Professor Mahesan Niranjan*

Network+ Coordinator: *Dr Samantha Kanza*

Contents

1	Interview Details	1
2	Biography	1
3	Interview	1

Andy Stanford-Clark, chief technology officer of IBM UK and Advisory Board member for the AI3SD Network takes the AI3SD Q+A.

1 Interview Details

Title	AI3SD Interview with Professor Andy Stanford-Clark
Interviewer	MP: Michelle Pauli - MichellePauli Ltd
Interviewee	ASC: Professor Andy-Stanford Clark - IBM
Interview Location	AI3SD Conference 2019
Dates	18/11/2019

2 Biography



Figure 1: Professor Andy Stanford-Clark

Andy is IBM's Chief Technology Officer and Distinguished Engineer. He has over 40 patents to his name and works extensively with Internet of Things technologies. Alongside his work at IBM he is affiliated with three different Universities: Visiting Professor at the University of Newcastle, Honorary Professor at the University of East Anglia and Adjunct Professor at the University of Southampton. He is also a Fellow of the British Computer Society. Prior to becoming Chief Technology Officer, he spent 20 years working in the area of Internet of Things technologies.

3 Interview

MP: How is AI, or machine learning, changing how we do science?

ASC: As a tool, AI is allowing us to explore spaces that we couldn't go into before. In particular, trying out an ensemble of ideas and looking through them all, using AI to either validate results or create the infill between real lab results and the predicted ones in between. It means we're essentially magnifying the capability of researchers. It's also evident that we're very much in the infancy of this science. Although AI itself has been around for ages, we've only recently had the huge computing power and the more sophisticated algorithms which allow us to do things like good speech recognition, for example, that we just couldn't do 20 or 30 years ago when the algorithms were there but we didn't have the compute resource. It means we can do things that we couldn't do before, which is hugely exciting.

At the moment it's 'let many flowers bloom and see what happens' rather than a methodical approach. I'm keen that we define 'science of AI' in the same way we have done with computer science and web science, through, for example, Wendy Hall and Nigel Shadbolt's work in Southampton. We need to think about how we formalise the science of AI.

MP: Do you anticipate that happening?

ASC: I'm sure it will. I think as people start to do meta studies on how AI is being used to solve certain classes of problems, those best practices will start to be recorded. And then someone will say, "we should formalise this and start teaching it as a science in university". I think we're probably a few years away from that now, even though I comfortably predict that it will happen.

MP: And when that happens, will ethics be a central part of it?

ASC: Yes, very much so. We're already becoming hyper sensitised to the bias problems with AI and the lack of explainability, the fundamental lack of transparency in the output of a system which is based purely on its input data. For example, adding a little black dot to a no entry sign fools a Tesla vision system into thinking it's a 30 mile an hour zone!

Now, we laugh at that, but if you think about optical illusions, for example the image that looks like either a vase or two faces, we use that as an entertainment, but, actually, if we had to act on our perception of that image and make a decision, then some of us will get it wrong. Some will choose one, some choose the other because of which way we saw it at first. That shows it's really hard for us and that it's going to be even harder for the AI systems. So accountability, ethics and techniques for avoiding bias, being aware of biases and recognising them will be crucial. Some of the biases might be intentional, but recognising unintentional biases is absolutely going to have to be part of that process.

MP: We talked a little bit about how AI is changing how we do science and you talked about solving problems, what about the impact on scientific discovery?

ASC: The 'robot scientist' approach is, in effect, trying to proceduralise the scientific method as described hundreds of years ago. That idea is to create a hypothesis, test the hypothesis, interpret the results, rinse, repeat. Keep going around until you've got the answer. At face value that seems like something you should be able to automate. However, there's a 'magic happens here' step in either the creation of the hypothesis or the interpretation of results. I

can imagine a system with a human in the loop where you're presented with set of options and asked to guide the automated system towards the direction it should head based on your intuition or your eureka moments, but having discovered things myself, there isn't a procedural method for getting from 'here's a problem' to 'here's the answer'. It's normally 'go and stand in the shower for 20 minutes and then, aha, got it'. And how would you ever teach a computer that?

I think it's more like if you talk to police officers who are solving crimes they tend to say, "it's not a eureka moment, it's huge amount of background legwork, knocking on every door in the street, checking all the registration numbers. I went down the road interviewing witnesses." It is that methodical slog, going through all the huge amounts of data that usually leads to the answer. And I think that's really where computers can help because if you say, "well, I'm going to have to try these 500 experiments to see which one gives the right answer," I'd much rather give that to a robot, which means you may actually solve things through meticulous exploration of a space rather than waiting until you have the eureka moment to go straight to the answer. I think it'd be a different way of doing science, but no less valid a one.

MP: But there will always be a place for human creativity?

ASC: Absolutely. I'm a firm believer in the eureka moment as a way of cracking really thorny problems for breakthroughs in science.

MP: And can a robot have a eureka moment?

ASC: I don't think it can at the moment, no. Not yet. Ask me again in 15 years' time.

MP: Well, I do want to ask you now to think ahead a bit and, although it's early days, what promise does AI hold?

ASC: We talk about narrow AI through to generalised AI. So narrow AI would be something like IBM Deep Blue that beat Garry Kasparov at chess. That computer could only play chess. It couldn't even play draughts, which is on the same board. Different pieces, different rules. It couldn't do that. If you think about where we are at the moment with things like the Amazon Echo, which seemingly can answer lots of questions, but through having 'skills', so you say, "Alexa, ask the barman how to make a bloody mary" and it'll invoke the barman skill to tell you how to make a bloody mary. But if you say, "Alexa, ask the barman how to get to Edinburgh" it'll have no idea. Were you to walk into a pub and ask the barman how to get to Edinburgh, every barman on the planet will have an opinion, right or wrong, about how to get to Edinburgh. So the barman skill on Alexa hasn't got that generalised knowledge. If you say, "ask the route finder how to get to Edinburgh," it'll know exactly that, but it'll have no idea how to make cocktails.

There's also no conversation aspect to it yet. If you think about scifi computers like the Star Trek computer, or Eddie the shipboard computer on Hitchhikers Guide to the Galaxy, or Holly on Red Dwarf, you can walk up to it and start talking about any subject and it'll give you an answer back. So it knows about you, it knows about your situation, it knows about the conversation it's had with you so far, it knows about the current state of the world, it has situation awareness, and it has a world knowledge equivalent to ours – what we've acquired over the last 40 plus years in terms of what we would probably call wisdom or just background knowledge, context. It merges all those different knowledges together. If you say, "what should we do next?" it knows how to answer that question. But at the moment, although we've made some progress into knowledge graphs and knowledge representation, we

don't really have the richness or representation to be able to create those databases. When we do it'll be great because it means that we can avoid having to put an AI through 30 years of evolution like we go through, we can just say, "there you go, there's the file. You've got it. There's your 30 years of education." We're way off that yet. We're in the very, very early days, and the AI we're playing with at the moment is nowhere near that kind of rich, generalised AI. We're probably 30 plus years away from that. I'm sure 30 years ago people were saying, "Oh, we'll have full AI in 30 years." So, I'll say 30 years.

MP: How about the intellectual debt issue - this idea that we can reach a solution or an answer and have no idea how we got there. Does that matter?

ASC: It does. It doesn't mean we can't do AI or you shouldn't do AI, but there is a class of problems where explainability is essential. However, with something like novel synthesis, if an AI comes up with a novel synthesis route for a new molecule, it doesn't actually matter how it came up with it. What's important is the novel synthesis route because now you can synthesise it and make a ton of money, so from that point of view it doesn't really matter, and we're probably in that situation now. We probably don't know exactly how some formulae do what they do or how electromicrographs work or how mass spectrometers really work, fundamentally. We just use the output they come out with, using it as a black box tool which does a job, and all that matters is understanding the result you get. There's a ton of stuff we can do without explainability but there will also be problems where we do need and will demand explainability.

MP: Finally, augmented intelligence. Tell me about that.

ASC: IBM tends to prefer the name augmented intelligence to artificial intelligence, mainly because we don't see AI as putting people out of work. Augmented intelligence is really the symbiosis of what humans are good at and what computers are good at. If you need to go through a ton of data, millions and millions of rows, you wouldn't give that to a human because it can be really boring and they'd fall asleep after five minutes. Whereas people can make use of the synthesised outputs, the information and knowledge and insights from that sea of data, to be able to make operational decisions. In other words, if that's the state of the world from our sensors and sensory data, then I'm going to make this decision. That's a really good harmony of computers and humans both doing what they're good at. And so the idea of augmenting intelligence is of the AI doing what it's good at, feeding into the bit that the human's good at – that's very much the way we see AI, augmented intelligence, developing in the future.