


A discussion about AI's conflicts and challenges



Source: https://techcrunch.com/2017/06/17/a-discussion-about-ais-conflicts-and-challenges/?curator=MediaREDEF&lipi=urn%3Ali%3Apage%3Ad_flagship3_detail_base%3BrGAg%2FVo8SgeWT3Ny1aGQPg%3D%3D

Date: June 17, 2017

Author: Natasha Lomas

Thirty five years ago having a PhD in computer vision was considered the height of unfashion, as artificial intelligence languished at the bottom of the trough of disillusionment.

Back then it could take a day for a computer vision algorithm to process a single image. [How times change.](#)

“The competition for talent at the moment is absolutely ferocious,” agrees Professor Andrew Blake, whose computer vision PhD was obtained in 1983, but who is now, among other things, a scientific advisor to UK-based autonomous vehicle software startup, [FiveAI](#), which is aiming to trial driverless cars on London’s roads in 2019.

Blake founded Microsoft’s computer vision group, and was managing director of Microsoft Research, Cambridge, where he was involved in the development of the Kinect sensor — which was something of an augur for computer vision’s rising star (even if Kinect itself did not achieve the kind of consumer success Microsoft might have hoped).

He's now research director at the [Alan Turing Institute](#) in the UK, which aims to support data science research, which of course means machine learning and AI, and includes probing the ethics and societal implications of AI and big data.

So how can a startup like FiveAI hope to compete with tech giants like Uber and Google, which are also of course working on autonomous vehicle projects, in this fierce fight for AI expertise?

And, thinking of society as a whole, is it a risk or an opportunity that such powerful tech giants are throwing everything they've got at trying to make AI breakthroughs? Might the AI agenda not be hijacked, and progress in the field monopolized, by a set of very specific commercial agendas?

"I feel the ecosystem is actually quite vibrant," argues Blake, though his opinion is of course tempered by the fact he was himself a pioneering researcher working under the umbrella of a tech giant for many years. "You've got a lot of talented people in universities and working in an open kind of a way — because academics are quite a principled, if not even a cussed bunch."

Blake says he considered doing a startup himself, back in 1999, but decided that working for Microsoft, where he could focus on invention and not have to worry about the business side of things, was a better fit. Prior to joining Microsoft his research work included building robots with vision systems that could react in real time — a novelty in the mid-90s.

"People want to do it all sorts of different ways. Some people want to go to a big company. Some people want to do a startup. Some people want to stay in the university because they love the productivity of having a group of students and postdocs," he says. "It's very exciting. And the freedom of working in universities is still a very big draw for people. So I don't think that part of the ecosystem is going away."

Yet he concedes the competition for AI talent is now at fever pitch — pointing, for example, to startup [Geometric Intelligence](#), founded by a group of academics and [acquired by Uber](#) at the end of 2016 after operating for only about a year.

"I think it was quite a big undisclosed sum," he says of the acquisition price for the startup. "It just goes to show how hot this area of invention is."

"People get together, they have some great ideas. In that case instead of writing a research paper about it, they decided to turn it into intellectual property — I guess they must have filed patents and so on — and then Uber looks at that and thinks oh yes, we really need a bit of that, and Geometric Intelligence has now become the AI department of Uber."



Blake will not volunteer a view on whether he thinks it's a good thing for society that AI academic excellence is being so rapidly tractor-beamed into vast, commercial motherships. But he does have an anecdote that illustrates how conflicted the field has become as a result of a handful of tech giants competing so fiercely to dominate developments.

“I was recently trying to find someone to come and consult for a big company — the big company wants to know about AI, and it wants to find a consultant,” he tells TechCrunch. “They wanted somebody quite senior... and I wanted to find somebody who didn't have too much of a competing company allegiance. And, you know what, there really wasn't anybody — I just could not find anybody who didn't have some involvement.

“They might still be a professor in a university but they're consulting for this company or they're part time at that company. Everybody is involved. It is very exciting but the competition is ferocious.”

“The government at the moment is talking a lot about AI and the context of the industrial strategy and understanding that it's a key technology for productivity of the nation — so a very important part of that is education and training. How are we going to create more excellence?” he adds.

The idea for the Turing Institute, which was set up in 2015 by five UK universities, is to play a role here, says Blake, by training PhD students, and via its clutch of research fellows who, the hope is, will help form the next generation of academics powering new AI breakthroughs.

The big breakthrough over the last ten years has been deep learning but I think we've done that now,” he argues. “People are of course writing more papers than ever about it. But it's entering a more mature phase where at least in terms of using deep learning. We can absolutely do it. But in

terms of understanding deep learning — the fundamental mathematics of it — that’s another matter.”

“But the hunger, the appetite of companies and universities for trained talent is absolutely prodigious at the moment — and I am sure we are going to need to do more,” he adds, on education and expertise.

Returning to the question of tech giants dominating AI research he points out that many of these companies are making public toolkits available, such as Google, Amazon and Microsoft have done, to help drive activity across a wider AI ecosystem.

Meanwhile academic open source efforts are also making important contributions to the ecosystem, such as Berkley’s deep learning framework, Caffe. Blake’s view therefore is that a few talented individuals can still make waves — despite not wielding the vast resources of a Google, an Uber or a Facebook.

“Often it’s just one or two people — when you get just a couple of people doing the right thing it’s very agile,” he says. “Some of the biggest advances in computer science have come that way. Not necessarily the work of a group of a hundred people. But just a couple of people doing the right thing. We’ve seen plenty of that.”

“Running a big team is complex,” he adds. “Sometimes, when you really want to cut through and make a breakthrough it comes from a smaller group of people.”

That said, he agrees that access to data — or, more specifically “the data that relates to your problem”, as he qualifies it — is vital for building AI algorithms. “It’s certainly true that the big advance over the last ten years has depended on the availability of data — often at Internet-scale,” he says. “So we’ve learnt, or we’ve understood, how to build algorithms that learn with big data.”

And tech giants are naturally positioned to feed off of their own user-generated data engines, giving them a built-in reservoir for training and honing AI models — arguably locking in an advantage over smaller players that don’t have, for example in Facebook’s case, billions of users generating data-sets on a daily basis.

Although even Google, via its AI division DeepMind, has felt the need to acquire certain high value data-sets by forging [partnerships](#) with third party institutions — such as the UK’s National Health Service, where DeepMind Health has, since late 2015, been accessing [millions](#) of people’s medical data, which the publicly funded NHS is custodian of, in an attempt to build AIs that have diagnostic healthcare benefits.

Even then, though, the vast resources and high public profile of Google appears to have given the company a leg up. A smaller entity approaching the NHS with a request for access to valuable (and highly sensitive) public sector healthcare data might well have been rebuffed. And would certainly have been less likely to have been actively invited in, as DeepMind says it was. So when it's Google-DeepMind offering 'free' help to co-design a healthcare app or their processing resources and expertise in exchange for access to data, well, it's demonstrably a different story.

Blake declines to answer when asked whether he thinks DeepMind should have released the names of the people on its [AI ethics board](#). ("Next question!") Nor will he confirm (nor deny) if he is one of the people sitting on this anonymous board. (For more on his thoughts on AI and ethics see the additional portions from the interview at the end of this post.)

But he does not immediately subscribe to the view that AI innovations must necessarily come at the cost of individual privacy — as some have suggested by, for example, arguing that Apple is fatally disadvantaged in the AI race because it will not data-mine and profile its users in the no-holes-barred fashion that a Google or a Facebook does (Apple has rather opted to perform local data processing and apply obfuscation techniques, such as [differential privacy](#), to offer its users AI smarts that don't require they hand over all their information).

Nor does Blake believe AI's blackboxes are fundamentally unauditible — a key point given that [algorithmic accountability](#) will surely be necessary to ensure this very powerful technology's societal impacts can be properly understood and regulated, where necessary, to avoid bias being baked in. Rather he says research in the area of AI ethics is still in a relatively early phase.

"There's been an absolute surge of algorithms — experimental algorithms, and papers about algorithms — just in the last year or two about understanding how you build ethical principles like transparency and fairness and respect for privacy into machine learning algorithms and the jury is not yet out. I think people have been thinking about it for a relatively short period of time because it's arisen in the general consciousness that this is going to be a key thing. And so the work is ongoing. But there's a great sense of urgency about it because people realize that it's absolutely critical. So we'll have to see how that evolves."

On the Apple point specifically he responds with a "no I don't think so" to the idea that AI innovation and privacy might be mutually exclusive.

"There will be good technological solutions," he continues. "We've just got to work hard on it and think hard about it — and I'm confident that the discipline of AI, looked at broadly so that's machine learning plus other areas of computer science like differential privacy... you can see it's hot and people are really working hard on this. We don't have all the answers yet but I'm pretty confident we're going to get good answers."

Of course not all data inputs are equal in another way when it comes to AI. And Blake says his academic interest is especially piqued by the notion of building machine learning systems that don't need lots of help during the learning process in order to be able to extract useful understandings from data, but rather learn unsupervised.

“One of the things that fascinates me is that humans learn without big data. At least the story's not so simple,” he says, pointing out that toddlers learn what's going on in the world around them without constantly being supplied with the names of the things they are seeing.

A child might be told a cup is a “cup” a few times, but not that every cup they ever encounter is a “cup”, he notes. And if machines could learn from raw data in a similarly lean way it would clearly be transformative for the field of AI. Blake sees cracking unsupervised learning as the next big challenge for AI researchers to grapple with.

“We now have to distinguish between two kinds of data — there's raw data and labelled data. [Labelled] data comes at a high price. Whereas the unlabelled data which is just your experience streaming in through your eyes as you run through the world... and somehow you still benefit from that, so there's this very interesting kind of partnership between the labelled data — which is not in great supply, and it's very expensive to get — and the unlabelled data which is copious and streaming in all the time.

How do we make the best use of a very limited supply of expensively labelled data? Understanding that labelled data is in very short supply — and privileging the labelled data. How are we going to get the algorithms that flourish in that environment?

“And so this is something which I think is going to be the big challenge for AI and machine learning in the next decade — how do we make the best use of a very limited supply of expensively labelled data?”

“I think what is going to be one of the major sources of excitement over the next five to ten years, is what are the most powerful methods for accessing unlabelled data and benefiting from that, and understanding that labelled data is in very short supply — and privileging the labelled data. How are we going to do that? How are we going to get the algorithms that flourish in that environment?”

Autonomous cars would be one promising AI-powered technology that obviously stands to benefit from a breakthrough on this front — given that human-driven cars are already being equipped with cameras, and the resulting data streams from cars being driven could be used to train vehicles to self drive if only the machines could learn from the unlabelled data.

[FiveAI](#)'s website suggests this goal is also on its mind — with the startup saying it's using “stronger AI” to solve the challenge of autonomous vehicles safely navigating complex urban

environments, without needing to have “highly-accurate dense 3D prior maps and localization”. A challenge billed as being “defined as the top level in autonomy – 5”.



“I’m personally fascinated with how different it is humans learn from the way, at the moment, our machines are learning,” adds Blake. “Humans are not learning all the time from big data. They’re able to learn from amazingly small amounts of data.”

He cites [research](#) by MIT’s Josh Tenenbaum showing how humans are able to learn new objects after just one or two exposures. “What are we doing?” he wonders. “This is a fascinating challenge. And we really, at the moment, don’t know the answer — I think there’s going to be a big race on, from various research groups around the world, to see and to understand how this is being done.”

He speculates that the answer to pushing forward might lie in looking back into the history of AI — at methods such as reasoning with probabilities or logic, previously applied unsuccessfully, given they did not result in the breakthrough represented by deep learning, but which are perhaps worth revisiting to try to write the next chapter.

“The earlier pioneers tried to do AI using logic and it absolutely didn’t work for a whole lot of reasons,” he says. “But one property that logic seems to have, and perhaps we can somehow learn from this, is this idea of being incredibly efficient — incredibly respectful if you like — of how costly the data is to acquire. And so making the very most of even one piece of data.

“One of the properties of learning with logic is that the learning can happen very, very quickly, in the sense of only needing one or two examples.”

It’s a nice idea that the hyper fashionable research field of AI, as it now is, where so many futuristic bets are being placed, might need to look backwards, to earlier apparent dead-ends, to achieve its next big breakthrough.

Though, given Blake describes the success of deep networks as “a surprise to pretty much the whole field” (i.e. that the technology “has worked as well as it has”) it’s clear that making predictions about the forward march of AI is a tricky, possibly counterintuitive business.

As our interview winds up I hazard one final thought — asking whether, after more than three decades of research in artificial intelligence, Blake has come up with his own definition of human intelligence?

“Oh! That’s much too hard a question for the final question of the interview,” he says, punctuating this abrupt conclusion with a laugh.

On why deep learning is such a black box

“I suppose it’s sort of like an empirical finding. If you think about physics — the way experimental physics goes and theoretical physics, very often, some discovery will be made in experimental physics and that sort of sets off the theoretical physics for years trying to understand what was actually happening. But the way you first got there was with this experimental observation. Or maybe something surprising. And I think of deep networks as something like that — it’s a surprise to pretty much the whole field that it has worked as well as it has. So that’s the experimental finding. And the actual object itself, if you like, is quite complex. Because you’ve got all of these layers... [processing the input] and that happens maybe ten times... And by the time you’ve put the data through all of those transformations it’s quite hard to say what the composite effect is. And getting a mathematical handle on all of that sequence of operations. A bit like cooking, I suppose.”

On designing dedicated hardware for processing AI

“Intel build the whole processor and also they build the equipment you need for an entire data center so that’s the individual processors and the electronic boards that they sit on and all the wiring that connects these processors up inside the data center. The wiring actually is more than just a bit of wire — they call them an interconnect. And it’s a bit of smart electronics itself. So Intel has got its hands on the whole system... At the Turing Institute with have a collaboration with Intel... and with them we are asking exactly that question: if you really have got freedom to design the entire contents of the data center how can you build the data center which is best for

data science?... That really means, to a large extent, best for machine learning... The supporting hardware for machine learning is definitely going to be a key thing.”

On the challenges ahead for autonomous vehicles

“One of the big challenges in autonomous vehicles is it’s built on machine learning technologies which are — shall we say — “quite” reliable. If you read machine learning papers, an individual technology will often be right 99% of the time... That’s pretty spectacular for most machine learning technologies... But 99% reliability is not going to be nearly enough for a safety critical technology like autonomous cars. So I think one of the very interesting things is how you combine... technologies to get something which, in the aggregate, at the level of assist, rather than the level of an individual algorithm, is delivering the kind of very high reliability that of course we’re going to demand from our autonomous transport. Safety of course is a key consideration. All of the engineering we do and the research we do is going to be building around the principle of safety — rather than safety as an afterthought or a bolt-on, it’s got to be in there right at the beginning.”

On the need to bake ethics into AI engineering

“This is something the whole field has become very well tuned to in the last couple of years, and there are numerous studies going on... In the Turing Institute we’ve got a substantial ethics program where on the one hand we’ve got people from disciplines like philosophy and the law, thinking about how ethics of algorithms would work in practice, then we’ve also got scientists who are reading those messages and asking themselves how do we have to design the algorithms differently if we want them to embody ethical principles. So I think for autonomous driving one of the key ethical principles is likely to be transparency — so when something goes wrong you want to know why it went wrong. And that’s not only for accountability purposes. Even for practical engineering purposes, if you’re designing an engineering system and it doesn’t perform up to scratch you need to understand which of the many components is not pulling its weight, where do we need to focus the attention. So it’s good from the engineering point of view, and it’s good from the public accountability and understanding point of view. And of course we want the public to feel — as far as possible — comfortable with these technologies. Public trust is going to be a key element. We’ve had examples in the past of technologies that scientists have thought about that didn’t get public acceptability immediately — GM crops was one — the communication with the public wasn’t sufficient in the early days to get their confidence, and so we want to learn from those kinds of things. I think a lot of people are paying attention to ethics. It’s going to be important.