

2016 Conference Transcription

Date	Thursday 31 March, 2016
Session Title	Intelligence - Panel Discussions 1& 2
Session Time	10:30-12:45
Moderator	Vikas Shah
Speakers	Lydia Nicholas and Darius Kazemi
Notes	n/a

Intro	<p>Hello and welcome to FutureEverything 2016 Festival Podcast Series. Over two days, in Manchester's iconic Town Hall, we task designers, artists, scientists, and many more, to rethink our resources from life, earth and intelligence, to community and uncertainty, our speakers asked what we might need less, and more of, in our near future.</p> <p>In this panel discussion we hear Darius Kazemi and Lydia Nicholas further discuss Intelligence as a resource.</p>
-------	--

Panel Discussion

Lydia Nicholas & Darius Kazemi

Vikas Shah	<p>In the morning Q&A, I let you guys off because I thought you might be a little bit shy, a little bit tired, not enough coffee. But this time I am fully expecting more questions from all of you, which will be lovely. I'll just kick things off.</p> <p>So in some ways, we've seen a lot of earth scientists talk about this anthropocene, about how humanity is changing the world in unexpected ways. It's going to have this existential impact. A question for both of you. Do you think we're entering into a technocene where technology is going to have a similar existential impact on our world?</p>
Lydia Nicholas	<p>I think they're the same thing. Humans, before they created agriculture, weren't having a huge impact on the world, at least not in the same way. You could talk about there's a corporate AI, I wonder if you could go back to think about the warring city state AI and the rules that needed to be enacted in order to survive</p>

	<p>there, and the kinds of things that began to develop. You needed to have more efficient agriculture. You needed to take over more land, you needed to slash and burn to create ways to feed your growing population. It's those sets of rules that do that. So I think it would be hard to define a difference.</p>
Darius Kazemi	<p>I agree generally. I like to talk about anthropogenic intelligence instead of artificial intelligence. So you could have asked interesting questions, like what second order affects would our AIs and technology have separate from us, in an independent way? But I think it's also important to know that AI is never going to be independent from us. It's always going to be a symbiotic relationship. In a weird way, the Matrix actually had it right. There's always going to be this relationship between the two.</p>
Vikas Shah	<p>You mentioned the Matrix, and with science fiction, we always have this slightly petrified or warped view of the world. Do you think that actually the public perception of AI is that, in a way, something which we're craving for that overlord of technology above us? Why do we take something and simply put it into catastrophe land?</p>
Lydia Nicholas	<p>I think one of the things about catastrophe and disaster stories is that they make us feel less responsible for fixing things in the slow trudging way that happens in the real world. So alongside the Longitude Prize, which is a ten pound prize for solving the problem of antibiotic resistance, we commissioned a series of stories around that, and then had some workshops around how you do responsible communication around disasters and existential threats to humanity. We kept coming back to the fact that it's about abdicating responsibility. If there's a zombie apocalypse, I don't need to go to work tomorrow. I don't need to do the effortful thing of creating a union, of joining a union, of trying to push back slowly and effortfully against the corporate machine. I just have to shoot people and run away and steal a helicopter, which feels much easier a lot of the time. [Laughter].</p>
Darius Kazemi	<p>Speaking of fiction, I would actually recommend that everybody who's interested in this topic to read the book called 'Speak' by an author named Louisa Hall. It's one of the best science fiction books I've read ever. It came out last year and it's about bots and post-apocalypse. The apocalypse that happens is actually a climate change apocalypse, and the bot apocalypse is this weird quiet secondary thing that just affects human lives, but it's not like the sea level is rising or anything. Definitely worth checking out Speak by Louisa Hall. I think mostly, I agree.</p>
Vikas Shah	<p>Darius, in your work, you're often using humour. It might seem perhaps like a glib question, but why humour? What is it about humour that you think is helping us learn something?</p>
Darius Kazemi	<p>Most of what I make is secretly educational. I like to think that what I'm doing is increasing people's procedural literacy when they engage with my bots over a long period of time, and of course it's more fun to engage with a thing like that when it makes you laugh every time you look at it, and you think about it sometimes in your off hours and all that. I think humour also slots in very nicely... when we think of humour, we think of spontaneity, but actually a lot of humour is very rote. Most of what I do is, I look at rote repetitive jokes that</p>

	humans tell in the wild, and then I think of a very simple algorithm that can replicate that rote joke.
Vikas Shah	Perhaps for you Lydia, as well, in your work, how are you finding AI and technologies around that are impacting how we're running our societies, and how we're running our governments, for want of a better phrase?
Lydia Nicholas	<p>I actually think that some of it is providing a spectacular opportunity, in that we've always been making government decisions in obscure and opaque ways. But if you start introducing a neural net, you have the chance to radically redefine what responsibility might be and what a citizen might be. So in some of the workshops we ran, there were all these wonderful and strange conceptions of how citizenship could be a statistical function, a relationship between your connectedness with a particular place, and there might be ways to think about how you balance opacity and accuracy in AI. So if you're trying to make a judgement about a situation, do you choose something that has only five factors, rather than five thousand, so that if you need to dig into it and check for bias and problems, you can go into it? Even if that gives you a slightly less accurate answer?</p> <p>All of those are things that have been true for a long time when you've been dealing with incredibly complex systems of government. That it's always been very difficult to work out how they work, or what factors are influencing what, but now you're introducing something that has a potential big scare factor for the public, and which you have this lovely line about this neural net cannot be understood even by experts, so you have to redefine your terms. I really hope that it's an opportunity that can be seized to think about how we might change our systems of law to adapt to the realities of modern life.</p>
Vikas Shah	That's wonderful. Thank you. There's a question just here from the front on this side.
Question 1	Thank you very much for your time, I love the talks. I had the privilege of going to see Paul Mason at St Paul's cathedral to release his book on post-capitalism. One thing that he talked about was using AI and big data for modelling the earth as a whole to enable a resource-based economy. So considering we don't know about ninety five percent of what these AIs are doing, is there any way of building AIs with rules, so that they're self-organising and mirror nature in a way that can enable that modelling in a resource based economy but is also beneficial to humans as well?
Darius Kazemi	My natural inclination is to say yes, of course there is going to be some way to do that. I think we have to ask ourselves whether it's a good idea to do that or not. What I'm mostly worried about in that situation is because you're mentioning central planning by computer. That's essentially what this idea is, right?
Question 1	I would say, decentralised.
Darius Kazemi	It could be decentralised too, but it's still an entity, even if it's a distributed entity that does the planning. Of course, you could say, what if we just do it in tandem with humans? Like the entity has a place at the table. But I think it's important to

	<p>consider the rhetorical affordances that software has when it does have a place in a decision making group of people. What we are trained to think is that software is impartial. So of course, the software is always going to be the tiebreaker if we have a tie. But software always has biases and in the case of these sorts of high level AIs, we don't even know what universe these biases come from. We might not even be able to understand what the shape of these biases are, and they might not even make sense to us as biases but they might shake out in the end to have some kind of second, third, fourth, fifth order effect on things. I like having computers at the table actually, when decision making has to happen, but I think we really need to retrain people on how to interact with computers, and how to work together with computers.</p>
Lydia Nicholas	<p>There are incidences of thinking about ethical service design of the interaction with AI. The best way I've heard it described is, as always, giving plural advice and multiple levers to press, to play with the system and work out its affordances as much as you can. So the idea of a computer managing the system from the top, I find quite unnerving. But something that can provide you with multiple answers and multiple ways of doing things, that you could then try and test out and run, is something that is a little less terrifying.</p> <p>The other thing is worrying about the biases within the model of the kind of information that it's easy for a computer to measure, and the stuff that it isn't. So it's much easier to measure something a particular kind of sensor can feed in, but more complex connections, community, cohesion and things like that are harder for it to measure, so it is filtered down its list of priorities. So you end up biased towards things which are much easier to measure and can feed into the model cheaply and quickly.</p>
Vikas Shah	There's a question from the front here as well.
Question 2	I'm interested in... there's always this massive rush to beat the Torian Test. Do you feel that there should be another kind of test? What do you feel of the Torian test generally?
Darius Kazemi	<p>I always think we're moving goal posts for artificial intelligence. For a long time it was 'well, if you can beat a chess Grand Master, then we've achieved true intelligence' and then it does, and it's like, wow. How about Go? Then it does. I'm waiting to figure out what the next goal post is that we're going to set. The Torian Test itself is weird and fraught because there are humans who fail the Torian Test. I think the Torian Test is a great thought experiment, but when the rubber hits the road, I don't think it's a terribly useful test.</p>
Lydia Nicholas	<p>I will not probably manage to sum it up as the original Tweet, but in the same fortnight that we managed to have an artificial intelligence beat Go, we proved we weren't able to keep a chat bot from expressing Nazi views for twenty four hours. No chat bot has managed to maintain an ethical position, or think deeply about class, privilege and race. So I think that those are interesting things for AI to have to...</p>
Darius Kazemi	There's our new goal post.
Lydia Nicholas	Yes, there's our new goal post. I'd like one of them to solve privilege.

Vikas Shah	There's one other thing that occurred to me when you were talking. I'm sure lots of people in the room today, especially after the party last night, would have had a paracetamol, an aspirin, or a neurofen or something. You take it and it works and that's great. You very rarely question, how do the molecules in this actually work, and what do they do? I think with AI, and with the algorithms that are in many ways running our lives, to what extent do we need to know that intrinsic function of how it works. Is it not enough to know that it does within safeguards? I just wanted to get both of your views on that.
Darius Kazemi	I think when I was talking about the rhetoric of bots earlier, and the rhetoric of artificial intelligence, I think that's the level that everyone needs to be focused on. Experts and practitioners, of course, need to understand the how as much as possible, but I think it's important to just black box the artificial intelligence and to say, alright it's an entity and it's communicating with me, what do I need to know about communicating with this? What are the protocols? Again, I talk about aliens all the time, because it is very much like being alien. Imagine if you were the diplomatic officer on a star ship that was going out to meet an alien race for the first time. You would need very special training to communicate with a race that you don't even know what kind of language capacity they have and all that stuff. So I honestly think that people should be prepared for dealing with AI, like you would be prepared for dealing with a theoretical alien race.
Lydia Nicholas	I think a fundamental difference between the black box of a human skull and the black box of an AI, is that, as a human, I can still take responsibility for my choices and my decisions. So I can say 'put your resources in this bucket, rather than that bucket', and if something goes wrong, I can be punished for that. Whereas you can't punish an AI for making a cruel decision. It doesn't even really understand what that means. So there's a need if something is a simple tool that cannot feel guilt or empathy, then that needs to be understood, because someone at some point needs to be able to make a decision about whether this is fundamentally ethical or right to use.
Vikas Shah	We have a question here and then we'll come to you sir.
Question 3	I'm interested in this idea how we can learn to tech with this artificial intelligence, alien intelligence. If I give you the job to be a director of the Zoo of artificial intelligence, what would be your choice of intelligence to put on show? What would you feed them, what kind of audience do you expect? What does this Zoo look like?
Vikas Shah	I'm sensing an art project in this already. [Laughter].
Lydia Nicholas	Someone is filling in their Arts Council bid right now.
Darius Kazemi	The Zoo is an interesting metaphor because I think Zoos are fairly cruel in a lot of ways. So I would feel very uncomfortable putting AI in a Zoo to begin with. I think that's about as much as I have in response to that, sorry.
Lydia Nicholas	At the moment I think Twitter is a decent Zoo for a lot of artificial intelligences, but it's one of those nice ones where they get to roam around free in quite a large environment and get fed followers. [Laughter].

Darius Kazemi	I do think there's interesting things to be shown with AIs talking to other AIs, and observing those interactions.
Vikas Shah	A question from the front here.
Question 4	<p>Thanks for your talks and analysis. They have all been totally fantastic. I'm very curious to know what both of you think, or have to say, about the relationship between programmer creator and the AI. I feel like there's an incredible void of analysis as to what that relationship is and an incredible desire and momentum in toward severing that link. There's almost no discussion of this continued relationship between creator of the AI and the AI, and when do those things actually become separate; in so far as if I make a machine that mows my lawn, or if I make a machine, there's a sense in which I've [inaudible 27:56] my own functions into this machine. Sorry, I'll disregard my example. You know what I'm asking. So what do you have to say about that?</p>
Lydia Nicholas	<p>I think there are two ways to approach it. I think there's one way where you can think more about an idea of distributed liability and responsibility. So we're talking about AI as if we're talking just about the function, but in fact, it's the dataset that makes a lot of difference in terms of the decision that comes out the end. So you're not just talking about the responsibility of the programmer, but the people that are running the dataset, the people that made the decisions about what to collect and how. Even specific decisions about [inaudible 28:37] are designed, how the data is cleaned at different points. What is disregarded as noise? What is accepted as something to delve into deeper? So there's a huge system of humans and very fallible humans, that contribute to the outcome of an AI.</p> <p>The self-driving car example is quite interesting, where you start to have to think about so how is it in a crash, who is liable? Is this something that we can distribute across a company? There are already laws in place, but they're beginning to be tested, and it's something that I can't actually quite follow the conversations that regulators are having at the moment about how all the different ways that it could work. But thinking about quite sophisticated ideas of liability.</p> <p>In terms of what's important in the authorship though, I keep coming back to the fact that I don't think the connection can ever fully be severed. I've had this conversation with Scot who's sitting next to you on Twitter a couple of days ago. The first example I had was there were a series of bot creators that were in my [inaudible 29:56] that were changing the way that their bots worked, because their queer followers were struggling with the Easter bank holiday weekend, going back home to families who are religious, who they might struggle to talk about their identities with. So they upped the rate of their comedy bots as a motion of support. So even though those bots in what they were doing had supposedly nothing to do with gender, sexuality, it mattered who the author was. So that's the kind of thing where, even if it seems that the bot is doing nothing to do with gender, race, class, power, it matters who wrote it and that we work towards diversity in authorship at all levels, even if it's quite an abstract and distant connection.</p>

Darius Kazemi	<p>I think Lydia did a great job of explaining the difficulty of pinning accountability in any organisation. Even when a space shuttle explodes because of an expanding O ring, it's who is responsible for that? Is it an individual at NASA? Is it a defence contractor? Is it the whole system? Do you fine, essentially, the AI that runs NASA? Do you fine it? Do you give us X percent of your resources as penance for this? Honestly, I don't think we have a great model of accountability period in corporations, governments, and things like that, so we're not starting from a very good baseline anyway.</p> <p>In clearer cut cases, like with the bots that I make, that's much more of a single creator and something that's been created. I have some fairly strong parental feelings about my bots. I see my responsibility as similar to that of a parent. It's got to stay separate from me, but I have to shepherd it and give it the right rules to follow, and guide it in its growth. I certainly think from a more pragmatic standpoint, that if I am going to take any kind of credit for the thing that I made, then I also have to take responsibility for when it messes up. So maybe if I were anonymously releasing these things and just letting them wreak havoc on the internet, I would be [inaudible 32:43]. But I get to come out to places like this and talk about this stuff, so of course, I have to take responsibility for what the bots do ultimately. So I do think there's a very strong connection between, at least in the case of individuals, or very small teams building these sorts of things. But of course, once the teams get bigger, you get to more general accountability questions.</p> <p>Of course, as Lydia mentioned, there are tons of different pieces of technology that go into these things too. So if I build a bot that used R&Ns do I blame Andre [Capathie? 33:27]? Do I blame the person who built the open source [inaudible 33:30] that I'm using? Is it my fault? It's probably a combination of all the above, but what do you do with that? It's a really interesting question.</p>
Vikas Shah	There's a question just from the front here as well.
Question 5	Thanks very much. I think my brain is firing on all pistons from both your talks and the earlier talks. I think this raised a lot of really important questions. Probably for me, it's raised more questions than answers about AI.
Lydia Nicholas	Sorry.
Question 5	No, it's fascinating. I suppose my question is really where this sits in the context of how this morning opened with the very urgent question and provocation about climate change? It would be interesting to hear your opinions on where you think AI sits within solutions around climate change. Do you think it can play a role, is it already playing a role with respect to finding and responding to more urgent challenges that humans are grappling with, given that there are some really interesting but still very many unknowns about the behaviour and decision making processes of AI?
Darius Kazemi	I think it would be very interesting to explore those avenues. I'm sure there are some think-tank's somewhere that are doing this right now, feeding in global climate models into recurring neural networks etc. Climate change is such an obviously important and potentially catastrophic thing that could make or break

	<p>the planet. I think AI is going to play a role in this sort of thing. Any situation where we have more data than a person can reasonably interpret is going to need to be interpreted by these things. Of course, the question is, again, what criteria these decisions are being made by. Are we telling it to just optimise by global temperature and what it thinks our new model for economics should be. Of course, there's the structure of the data that's sent in there too, the data that you see one of these AIs with can completely change the outcomes, and I can imagine duelling think-tanks. One think-tank is, we put all this climate data into power AI, and it said the answer is this. Well, they put this slightly different interpretation on the data in and we got this other this other solution. So I think AI is going to end up being a voice at the table in this sort of thing, but I don't think it's going to spit out the answer on a little slip of paper.</p>
Lydia Nicholas	<p>The crisis is a political one. You can use AI do whatever you want, to tell whatever kind of story that you want. If you want to optimise energy use across a system and say these are the best places to put solar panels, this is the best place for wind farms, then you could do that. But if you're the kind of government that then allows nimby laws to say, 'well, I don't want to have wind farms on my country estate, they have to be moved down wind', then that's going to happen. It's something that is a very human challenge to follow what might be the world saving answers, and it's one that will take courage, which I don't think is something that AIs have quite overcome that hurdle yet. Although, their neutrality and lack of fear as well might be the kind of bold statements they can make about, actually this is the way you should be organising your cities, or your [inaudible 37:47], are ways to reframe the debate in potentially a positive way.</p>
Darius Kazemi	<p>I do want to just piggy back off of that. I harp on this point a lot. Software, algorithms, artificial intelligence, can get away with saying things that a person can't. So if I, as a person, said 'we have to cull the population ninety nine percent', I would not feel comfortable making that recommendation, and if I did, I would probably have my peers at the table look at me like I was an evil person. But software can say that, and people will say 'well, that's interesting, how did it come up with that idea? Let's look into this'. So software comes from a very different position.</p>
Vikas Shah	<p>What about the flipside as well? I know, Lydia, that you've done an awful lot of work in terms of having a look at how technology is even making us smarter and our collective intelligence. Again, a question perhaps for you, and Darius if you want to jump on this one too, are these technologies actually going to help improve humanity? Are we going to end up smarter, or more capable, or more creative as a result?</p>
Lydia Nicholas	<p>I think both yes and no. Collective platforms used well have tremendous capacities to make us better at a lot of things we do. I was recently in Cambodia, and I interviewed a man that runs a sustainable tourism system, and he says that Trip Advisor, and without me saying that I worked in collective intelligence, the collective knowledge of that means that people can make deeper and more informed decisions. They can check it out. He'd been working in the field for fifty years and it was the coming of these kinds of platforms where someone might decide that they want to have an ethical holiday, though they've taken a plane,</p>

	<p>so their climate change contributions might not really make that much difference in the end, but anyway, they decide they want a more ethical holiday and now they can actually stick to those because lots of different people have gone and contributed comments, contributed ideas, about how that can work, recommendations, and then they can map out how to do that, which is something impossible just a few years before.</p> <p>On the other hand, it also works as a means of policing and of silencing people. If we're very aware that people can see everything we're doing, it's very hard to experiment radically and to try out new ways of being, to spend your Thursday experimenting with a different mode of dressing and thinking and speaking, so that you can work out, by Friday, whether that's something you want to do all the time. It's much harder to have those plural identities. So yes and no.</p>
Vikas Shah	<p>I guess on that one as well, FutureEverything has a huge arts component. Are we going to see, and I know obviously you had one of your bots involved in writing, many of them in fact. Are we going to see new forms of art emerging from AI and from bots? Are we going to have new ways of exploring what beauty is, what aesthetic is, in our world?</p>
Darius Kazemi	<p>Absolutely, we see that with all new technologies. Photography changed what painting was. You can pick any technology you want and it has some effect on arts. So yes, I do think that AI aesthetics is going to be a really interesting phase. Obviously there's a lot of people already interested in it. I'm interested in it, but also it's how new is this stuff really? I run something called 'NaNoGenMo', National Novel Generation Month, every year, and people spend November writing code that writes a novel. It's very similar to early surrealist and data experiments and Burroughs type cut ups, and all of that. So it's really just generating experimental writing, and it's a different kind of experimental writing, but it's a clear lineage.</p>
Vikas Shah	<p>We've only got a few minutes left so maybe time for one more. There's a questioner hiding right at the back there.</p>
Question 6	<p>Darius, you brought up the issue of definitions of artificiality and intelligence, and it got me thinking. Do you think there's actually an inherent problem, even in using the word artificiality? It seems to have become the regular linguistic way of talking about this alternative form of intelligence. I'm personally even a bit dubious about if there is such a thing as artificiality. I think there is such a thing as intelligence, but perpetuating the use of the word artificiality could be limiting, in terms of ways of thinking about this whole thing. Is there another word we could use? Or could we just say, look there's intelligence that manifests, or emerges in different ways. As we know about humans, we don't just have cognitive intelligence, there's the emotional intelligence and all sorts of types of intelligences. I'm just thinking about would you challenge the way that we just lock into that AI language really, because it's become a bit of a fixed notion, I think.</p>
Darius Kazemi	<p>I'm with you one hundred percent. I don't believe that there is a divide between natural and unnatural. I don't think that that exists. I'm in the same camp as Bruno Latour when he says we've never been modern, when he says there's not</p>

	<p>actually, natural and unnatural, it's all this flat ontology between everything. Whilst things are different from one another, they're not substantively different. So yes, I'm with you one hundred percent. That's why I'm playing around with new phrases like [anthropogenic? 44.10] intelligence and that sort of thing, because I do think that artificial is not great, it's not useful.</p>
Lydia Nicholas	<p>I agree with that. It's often useful to be more specific to talk about so we're using a neural net, and that has very different implications than something that's quite a simple set of rules that anyone can understand. That means that you might be talking about something where like, if this thing, that if it goes over this level then enforce this condition, that's something that I can understand. It's an intelligent system to an extent, it's an artificial intelligence system it doesn't matter, but as soon as it's something that can't be investigated, maybe that's something that I need to apply different rules, morals ideas to, because I need to work out how that fits into another system.</p> <p>So being specific about exactly what's going on inside of any particular kind of intelligence is sometimes useful, but the artificial divide is not particularly. Material culture anthropologists. I think that walking is a kind of technology. We walk in different ways to communicate different kinds of identities. Everything is technology and everything is natural.</p>
Vikas Shah	<p>I'm afraid we're going to have to wrap up the questions there, but could I just invite our founder and CEO Drew on stage. He'll do a wrap up in a moment. Before we do that, if you could all join me in thanking our speakers and thanking our first speaker in her absence as well. Thank you very much, absolutely fascinating.</p>
	<p>We hope you enjoyed this panel discussion and thanks for listening. You can hear the rest of the talks and discussions from 2016 at futureeverything.org/2016podcast.</p>

[Transcription ends]