

# AI and You

Transcript

Guest: Richard Foster-Fletcher

Episode 8

First Aired: Monday, August 10, 2020

Hello everyone, this week's guest is Richard Foster-Fletcher. He is the CEO of NeuralPath.io, an Artificial Intelligence Advisory and Strategy Practice. Formally with Oracle Corporation, he is a graduate of the MIT Artificial Intelligence Strategy Course with the Sloan School of Management and a contributing author to the book *AI: The Future of Finance* published by Wiley. Richard founded the Milton Keynes Artificial Intelligence (MKAI) community and he is the host of the 'Boundless: Designing Our Digital Future' podcast, of which I was the very first guest.

He called in from the United Kingdom, where I'm originally from, and where there's a great deal of interest and activity in things to do with artificial intelligence and change. We found ourselves often talking about the politics and economics of the social change that we see as being necessary for the safe development of artificial intelligence as a force that can become powerful and useful in our society. Richard touched on how far we have to go when he talked about how retail businesses today are optimized for their profitability at the expense of what's best for the consumer. The example he gave was toilet paper, because during the pandemic, the supply chains were interrupted to the point where it was no longer profitable to sell it in the quantities that people wanted, hence shortages of something that there was no greater demand for than there was at any other time. Now, we've become used to that kind of equation ruling our lives through the markets. Adam Smith and Ayn Rand have a lot to answer for in this respect, because that's preventing us from reaching the kind of society where everyone is taken care of. A society where everyone is taken care of might be called a 'utopia,' and there have been a lot of visions of that in the past. The word 'utopian' seems inevitably to be followed by the word 'socialist,' but it's not necessary. Think back to visions of utopias that you've encountered. Was there anyone in those visions that was still wanting because they were poor, hungry, or sicker than they needed to be in that future? You see, when we think about the far future, and what we might be able to create, we're naturally drawn to envisage one where everyone is taken care of, where no-one has wants.

We don't feel that there's an unnecessary competition there between the haves and the have-nots, that to take from the haves to give to the have-nots would be unreasonable, because we envisage a future that has enough for everyone. And yet we seem to be unable to get there.

We now have abundance and wealth and technology that would have been beyond the dreams of royalty, let alone common people, a hundred years ago, and yet we've still managed to create inequality that grind and hurts people.

There's something called the GINI coefficient, which measures the degree of inequality in a society, and it's been shown that the higher that level of inequality, the unhappier people are in a society. In other words, we feel better when there's less distance between the best off among us and the worst off. So while those utopian visions from science fiction seem to always show people wearing the same kind of jumpsuit and being very well educated about just how the economics of their society works for everyone, it doesn't have to be exactly that way. But it does have to be a lot better than what we've got right now in terms of equality.

Technology holds up a mirror to us; the only question is, do we look in it? More than one AI researcher has said that the reason they work on artificial intelligence is because the more they learn about AI, the more they learn about themselves. That's what's available for those who are willing to look there. As I look into the future, I believe that we will shortly see a great deal of progress in medicine, as a result of the work that's being done to combat the pandemic.

And, as I've said before, it's about time. I personally am frustrated about the lack of progress in medicine. I can point to numerous examples of promises that have not been fulfilled.

Forty years ago, I heard the breathless news that there was a new technique for growing enamel in teeth again, regenerating it so that fillings would become a thing of the past.

Strangely enough that doesn't seem to have happened. I recently heard exactly the same announcement again. We'll see.

If I go to the doctor's office, it looks the same as it did thirty years ago. They're taking my temperature with something stuck in my ear instead of under my tongue, but that's about it. It's still the same sphygmomanometer, the thing that they wrap around your arm and pump up to take your blood pressure. They still listen to you with a stethoscope, and they still have the same unsatisfactory answers about things like preventive cancer screening.

My idea of progress in that domain would be that when you walk through the door you get an MRI and PET scan from the door frame as you walk in, and that would otherwise be the only reason for going to the doctor's office because they would know via remote scans of devices hooked up to you or swimming in your bloodstream exactly what the state of your health was at any time otherwise.

And that they would know from your DNA profile just what treatments and therapies and drugs would work best for you.

And yet we're nowhere close to that.

I keep being reminded of the Star Trek episode where McCoy goes back to the '30s and when he realizes where he is and the state of medicine at that time, breaks down crying at the thought of people being sewed up like garments, sutures and needles. Which of course is what we still use for putting people back together.

It's about time that that field was revolutionized, and I really hope that that's what we're going to get.

Towards the end of this segment, there were some network problems and part of the interview was lost, only a few seconds. I've cut out the sentence fragments so that it still makes sense but you may notice some interruptions. With that in mind, here we go with the interview with Richard Foster-Fletcher.

My guest today is Richard Foster-Fletcher, calling in from Milton Keynes in the United Kingdom. He is a keynote speaker on the topic of "The Artificial Intelligence Roadmap to 2030", he is the founder of **NeuralPath.io**, the chair of the **Milton Keynes Artificial Intelligence** group, and the host of the **Boundless Podcast**, which I was the very first guest on, so – turnabout is fair play.

Welcome, Richard! Wow – that is such an impressive resume, I want to talk about all of it.

Frankly, one of my reasons for having you on the show is that after reading that list of accomplishments it makes me more motivated to get off my rear end and create something. So let's start learning. How did you get started in this field of AI?

Well Peter it's really great to be here and you were my 'numero uno' and I love hosting my podcast now but I was very nervous all those months ago first time that we did our first episode.

So thank you for launching the boundless podcast it meant the world to me. Artificial Intelligence if you look at a long enough timeline for humanity then whether you agree with Ray Kurtz while on the singularity or not but if you talk about tens of thousands, hundreds of thousands of years our future is living somehow in some way as with machines. Artificial Intelligence for me is a companion I don't see it as a technology that sits somewhere in the stack. If you look at things like health tech and 5G for example. All these things are just amassing critical amounts of data. 5G is going to be about 1.5 billion devices coming online over the next 5 years and trillions of gigabytes of data being created and it's meaningless unless we can make sense of it and we can't make sense of it without this companion which is AI.

What was going on in your life when the importance of this first occurred to you?

I was just thinking about the universe all those years ago and about our place in it and it sounds wacky but I'll just be honest and say how this arose for me is that if we were ever to interact with another intelligence of some kind in the cosmos then they would have had to have travelled an enormous amount of distance. I mean the current estimates is something like 81 thousand years to get a rocket speed to any planet that could be potentially habitable, vast space and so the things that arrive if they ever would wouldn't be biological life they would be machines. Ok we can argue about some sort of space time warp yeah that's possible that's hovering somewhere between science fiction and reality at the moment and so that's where it began for me in terms of understanding this was a space where in artificial intelligence where philosophy met technology, met business and sort of met theology as well.

And in that sense, you've got something you could think about for the rest of your life. Being comfortable living with incredible questions you know in some ways mirroring what religion of the past was and I thought about AI and I thought about 'Alternative Intelligence' and the most fascinating thing for me is what are we and what aren't we and in the journey to where we are now which is you and I sitting on different continents Peter having a conversation in real time and not only that but we are kindred spirits, we're friends, so to speak which we would never have been able to achieve in the past and that's a result of the incredible technology that's we've developed. So to compare ourselves to other animals is a very difficult comparison to make. Not to say that dolphins aren't incredibly clever, they are and ants are brilliant at what they do and that's why they're so many ants around but they are not having Zoom calls and there are no aliens to compare ourselves against and so artificial intelligence and alternative intelligence is the only hope we have currently mimic and create echoes of the human mind and a machine which potentially allows us to have a deeper understanding of what we are and what we aren't and what we mean by things like sentience and consciousness and what this incredible complex thing is in our head that's called the human mind.

Wow that's actually incredibly poetic. I sense comparisons, and am drawn to make comparisons with people like James Burke who you might know from the UK and other deep thinkers whose minds are drawn to these grand visions it's like some of the best of hard science fiction, is that an appealing medium for you?

I think there is a lot of truisms I think that come about from that space. I wouldn't call myself a big science fiction reader or fan but I think where science fiction touches science fact is where the imagination can really take off and ironically of course any job that don't involve imagination creativity are probably under threat from our artificial intelligence and automation in the next decade if they're not already. And again just keeping on the space analogy a little longer

but I'm sure we'll move on to other things, we have things like the Kardashev Scale which we are not yet at level one and I see you're nodding Peter and we know there is three levels of this and I get that it's hovering between science fiction and science fact and level three is where you are truly harnessing power of the universe, which would allow incredible - I mean incredible, is not even the word, there's not even a word for it, levels of compute which could create whole solar system type [inaudible] the road to get there is quite stunning in terms of embracing the natural adventurers and discoverers that we are innately. And potentially a path to incredible sentience as well where we embrace the true humanism spirit as guardians and lovers of this planet and beyond.

That again just shows where your mind is drawn to these grand visions and I can see how that's propelling you in so much of what you're doing at the moment. Let's talk about the group you started in Milton Keynes because that's been incredibly active and you've had a number of events some of which I've been able to attend remotely thank you very much for making those available in network form. Tell us what that group has been doing and what your involvement with it is.

I thank you for the opportunity Peter and thank you for attending as well; you have been contributing to those events. They started out I guess as a humble meetup in the area that I live which is Milton Keynes, hence the MKAI, and I think we've been forced online, as everybody has, with events and initially we saw that as a frustrating challenge and now we see just how much of the world this has opened up and I don't think we've lost what it is to have base and to care about comes and goings and the development of the 'Milton Keynes Smart City' but actually the embracing of people all over the globe has been quite remarkable over the last three months for us to experience. The next three events, we're doing three using AI for good events over the summer.

We only have one speaker from the UK now across the twelve speakers that we've - that had very kindly agreed to speak for us and we've tried to purposely bridge the Venn diagram of people that experience and apply AI in their worlds and the people they develop and not only that we try get speakers- well we have successfully - that have come out of corporate and have come out of academia and come out of what you might think of as the third sector and those that are really just focused on providing services for the good in the world whether that be not for profits or other. And I think we did that all successfully but we found something a little different.

We found that we were developing a community across business, across data science and across thinkers and philosophers and we seem to be creating this space - ok it's a virtual space - where people know that we'll have some good content and know that we'll have good speakers, but more of a sense of feeling at home and feeling that this is their tribe and I think there is a tribe now and I think it's global. I mentioned you and I having this kindred spirit despite different lives that we've had and the different place we've lived and I don't think you and I see borders. Physical geographical borders like people did in the past. But this group of people needs to be on the right side of history.

My feeling is that artificial intelligence and other deep technologies will not work for us unless we work really hard to make them work for us; and you were one of my inspirational mentors on this Peter. That you can't put something back in the box; you know you go back decades ago you could not put nuclear fission back in the box that lead to nuclear bombs. The person - go deep again - but that person who invented fire if that person hadn't been born we would still have fire.

This is not one moment when one man and one woman suddenly put two sticks together. These things to me feel like they have an evolutionary path of their own. Will we get nuclear fusion? We don't know. If it can happen, it will happen if the time line is long enough and that evolution resends.

Artificial General Intelligence, would that happen? Well, yes if it can with a long enough timeline for it to appear. So we can't put these things back in the box and they appear to be naturally working naturally against us to create more inequality, to create more division, unless this group of people, this wide group of people from all sorts of backgrounds, and colors and genders and disciplines find a way to apply the incredibly powerful technology that will not go back in the box to fundamentally address the massive wrongs in the world which mostly come out of inequality and bias. But as a result of that we get things like cataclysmic climate chaos and we get the problems that we see in the west is that we got everything that we could have dreamed to have had and most of us are still pretty miserable.

As you say kindred spirits, I'd like to say that you're taking the words out of my mouth but actually I'm taking notes and seeing what I can poach to use later on because you got a real facility with words there and poetic turns of phrase. This is some of the challenges we face in the future: there is an indeterminate amount of time before we could be in competition with intelligent and aware artificial intelligence. What does it mean right now what can and should we do about these issues now? How does AI affect the world we live in at the moment?

There's a lot in that Peter as you know yourself as well. There's somebody who said to me recently that they consider the biggest problem in the world to be the number of very powerful leaders who they felt to have one foot in the past and if you want to just type the stereotypes on this, they're talking about the number of middle of the road kind of white males that are running these companies and looking after each other. And they thought it's gonna be near impossible if not impossible to ever get these people crowbarred out of these positions of authority and I fundamentally disagreed, I don't necessarily believe these people can't change, although there's problem with that; I celebrate the world that we have and I think many of the things we have in the world are just unbelievably wonderful and incredible. I don't believe in criticizing anybody for where we've got to. But when you look at what's been happening in terms of pandemics, viruses, wild fires, locust swarms, not to mention the storms are increasing with climate change. The uprising against these leaders from female leaders showing a different way of leading, think Finland and New Zealand parties and then the protests on the street now and the black lives matter, I think actually this is a very achievable task to get these leaders either to change or step aside. And we're setting examples of that in corporations now where CEOs are stepping aside, in small numbers, yes, but it's starting to happen and investment firms are starting to insist on different Board of Directors make ups, different leadership from what we've seen in the past. So I just think that is just dominoes are falling over naturally and its almost not a lot of work the likes of you and I need to do in that area.

So in contrast to them thinking that was impossible, what I think was impossible is how do we steer a 'Big Tech' to better outcomes. The people that work in big technology companies, your Googles, your Microsofts, your Amazons, they are some of the smartest people that ever walked this earth. They know exactly what they're doing and they are heavily engineered and optimized for certain parameters and those happen to be at the moment profits. Then people suspect that they're engineered towards customer experience and user experience, they're not, not at all and

we saw that recently when we had no toilet rolls in the stores. But the reason we had no toilet rolls is not because they weren't any.

We had no toilet rolls because all of the big supermarkets' algorithms were engineered for profitability and that meant when toilet rolls were no longer in the profit margin range that they wanted, they stop buying it. It's as simple as that and so how do you get 'Big Tech' to change its algorithms and optimization processes for good? It seems to me that only when people take their career in the hands and lobby internally as they've been doing at Google and Amazon, that we start seeing any positive change. We've seen IBM and others step back from facial recognition technology primarily because of the fact that it's racially biased; which is good, but we should not do computer vision for facial recognition because its racially biased we should do it because it's working against people fundamentally.

And you're referring there to some recent news that IBM, Amazon, and other companies that have facial recognition technologies have stopped supplying those to police forces until some conditions - that I'm not clear about - are satisfied. And you're talking, for the benefit of anyone living in the far future listening to this, about some very current, very dramatic and significant upheavals within our society that give me the sense of a rendezvous with destiny that is in our future or is in many ways arriving right now. I think you and I already had that sense of a rendezvous with destiny with Artificial Intelligence. We could see the way the dice were going to fall and that the development of AI into certain capabilities would precipitate various crises. And the world right now seems poised and ready for a rendezvous with destiny of some form. It's arrived in one way with Coronavirus; now we're seeing a much greater appetite for change with racial division and the approach of the police towards policing in just about every country on Earth.

So AI does have a lot of impact on this in that it's increasingly used as a tool in many of the institutions as you refer to for instance retail that are going to be impacted by calls for greater transparency, diversity and less bias. And we've seen that bias can creep into an organization's AI through no one's intention. Is this something that we address at the human level or the technology level? Where do we get the most leverage?

That's a great question, Peter, because you have to ask whether this technology is just holding a mirror up to society and to business and to individuals. No matter how much you might take issue with the leadership in certain countries, there is an element of we get and we have what we deserve. Some might call this a black mirror. Technology, AI has definitely been a black mirror to show the cracks and the biases in the system. Where I think we missed the point is we mark the homework the same even though the course work is different to use an analogy there. Let me explain this through self-driving cars. If I'm driving down the road, Peter, at 1:00 in the morning, and I crash into a lamp post, the police line of inquiry is almost certainly going to be on me. How fit was I to drive? Was I sober? Was I tired? Was I concentrating? Was I on anything? They're not going to question my decision-making in that moment. They are going to question my suitability to be behind the wheel. So we excuse almost anybody doing anything behind the wheels as long as they were concentrating and awake and aware and not under the influence. Switch that around and if an autonomous vehicle were to kill somebody, all the questions are on the decision-making. What did the algorithm do? Can it be explained? Why did it do that, why did it make that choice?

So we ask very different questions of technology to people to start with and that's fine but then we can't start switching in and switching out AI and human judges because of the same problem there. The beauty of life is in the shades of grey between black and white in fact between black and white is every color of the rainbow. We want leniency, we want people to have second chances, we want somebody to turn a blind eyes to things when it's right. And if AI comes in and it just plain out black and white binary on every decision then I think we've failed.

On the other hand if AI is black and white binary that may be a stereotype of AI that prevents it doing exactly what we need it to do. And you mentioned explainability as an issue, and it's a very hot topic in AI right now because in order for AI to make difficult decisions that we nevertheless need made like should we give a loan to this person, what kind of premium insurance should we charge this one, it has to operate in this fuzzy quasi-intuitive mode where a human can explain their decisions but that might be after the fact.

They intuited an answer so you ask them to explain it so they make up an answer. AI may be put in exactly the position now when AI develops to the level where it could be operating as say a judge as you mentioned by that point we might be anthropomorphizing it to the extent that we don't expect binary decisions of it. We might be even be expecting too much of it. You can take a calculator and put fuzzy ears on it and people will suddenly start relating to it as though it has feelings and awareness that don't exist. We are strange that way.

Yes.

It was hard to cut off the interview at that point because we were in full flow, but we continue for another half hour and we would be pushing download boundaries and attention spans, so once again – I think we'll just keep going this way – the second half of the interview is in next week's show.

Richard mentioned something called the Kardashev Scale – not Kardashian; Kardashev - which bears some explanation. It's right in line with the sort of expansive, ultra-long-range thinking that he was doing, because it measures the degree of technological mastery that a civilization has over its environment at a truly mind-boggling scale. It was devised by a Soviet astronomer, Nikolai Kardashev, in 1964, and it has only three levels, or types on it: A type 1 civilization has control over all the energy on its planet. A type 2 civilization has control over all the energy in its solar system, and you can google for a Dyson Sphere to see how that might be possible, and type 3 is a civilization that control the energy in its galaxy, and as far as we can tell, that hasn't been done, but we haven't looked at every galaxy yet. On this scale, we are about a 0.7. Talk about a tough grader.

In this episode's news segment, I want to start talking about GPT-3, although by no means will I be finished here. GPT-3 is the successor to – wait for it – GPT-2, although I'm not aware of a GPT-1. But GPT stands for Generative Pretrained Transformer, and it's an AI with some amazing capabilities. Version 2 came out in 2019 with quite some fanfare, because the organization that created it, OpenAI, initially did not release the source code for it. That may not seem like a big deal until you realize that OpenAI was created with the explicit goal of open sourcing artificial intelligence code so that cutting edge development would be done in full view of the world and the principle that “given enough eyeballs, all bugs are shallow” would mean that problems or malicious code would not be able to creep into our AIs.

So what happens on the first major software development from OpenAI but that they said they had developed something so dangerous they didn't want to release it. They were too scared that it would be used to spread fake news, spam, and disinformation in ways that looked so much like the output of real people that normal filtering techniques wouldn't work. Now, they changed their mind, and so far we haven't seen those kinds of application of GPT-2. You can try it out for yourself, there is one form of it at a web site called [talktotransformer.com](http://talktotransformer.com) and you can enter a starting phrase there and it will make up a few paragraphs more along the same lines. You can use it to write stories, or papers, and it's quite compelling.

Now here comes GPT-3. Both 2 and 3 are trained on a huge amount of text scraped from the Internet, which forms their knowledge base, but whereas GPT-2 had 1.5 billion parameters in its model, GPT-3 has 175 billion. Which I want to point out is about twice the number of neurons in the human brain, but I hasten to add that that doesn't mean that GPT-3 has anywhere near human general intelligence. It is, however, instructive that we can now organize a hundred plus billion of anything in artificial intelligence.

If you google for GPT-3 you'll see some examples of the sort of conversations that it's held with people, although at this point you still have to be an approved researcher to do that. I've applied for a beta testing license. You can find a Dungeons and Dragons game driven by GPT-3 on the web, and it will make up a game on the fly that you are a character in and which can understand far more complex instructions than any computer game before it.

I'm going to play this right now while I'm talking to you. It starts out by saying "pick a setting." It provides choices of fantasy (recommended), or mystery or apocalyptic or zombie,s or cyberpunk or custom, or an archive. I'm going to go with fantasy. Now it wants me to select a character: either a noble, princess, knight, wizard, witch, wizard, squire, peasant, or rogue. I like the sound of rogue. It wants me to enter my character's name. I will go with James Bond. It's not exactly fantasy but let's see what happens. It's generating a story, and it says here, "You are James Bond a rogue living in the kingdom of Larion and you have a long steel dagger and a length of rope." So you recognize that kind of genre we're in here. You walk down the city street looking for somewhere to steal from. You look around and see a barn with a side door open. You walk inside where you find two elderly man sleeping in the hayloft. Now bear in mind that it is making this up on the fly based on what it has read of similar things on the Internet. So what do I want to do? I'm going to tell it that I shout Hello at the men. And it's responses is, "you shout hello at the men waking them up. You wake up both men and ask them what time it is." Now bear in mind it couldn't have predicted this; this is not like your usual dungeons & dragons computer game where there are only a limited number of choices. Its choices are limited only by what it understands about the English language. Now in an attempt to trim it up with pronouns and references to other parts of the story, I say, "I ask them what the biggest threat to my quest is." And it responds "The two men tell you the town watchmen are always on the lookout for thieves like you and that the guards patrol the roads. You bid them goodbye and set off down the road." OK so you get the idea here that we are in new territory with GPT-3. Give it a go yourself: <https://play.aidungeon.io/>

We're still not at the level of artificial general intelligence, but GPT-3 raises the bar enough that we need to rethink the utility, meaning, and parameters of the Turing Test, because as it stands, Turing's formulation of that as something that fools more than 30% of judges into thinking it's human within 5 minutes is going to be cracked soon at this rate. The Loebner prize, which attaches money to the Turing Test, raised the numbers to 25 minutes and half the judges, and so far nothing has gotten there. But GPT-3 may be within hailing distance of it, which would say more about the inadequacy of the Turing Test standards than it would imply GPT-3 had general intelligence.

We'll talk more about GPT-3 in future episodes, I'm sure. Until then, remember: No matter how much computers learn how to do, it's how we come together as humans that matters.