**Sue Nelson**
Hello, I'm Sue Nelson and welcome to another episode of the create the future podcast, brought to you by the Queen Elizabeth Prize for Engineering. Now, depending on what you read, artificial intelligence will either transform our lives, or destroy humanity. It's a subject that produces excitement and fear, especially within the fertile land of science fiction. AI, however, is grounded in reality and engineering, and this is the perfect opportunity to hear from two AI experts about its benefits, ethical issues, its future, and why we should sometimes proceed with caution. I'm joined by Dame Wendy Hall professor of computer science at the University of Southampton. She's someone who helps shape engineering policy and she was co-author of the government report "growing the artificial intelligence industry in the UK", which led to the creation of the office for AI, and a £1 billion deal for the AI sector. Thank you very much for joining us.

**Dame Wendy Hall**
Thank you, good to be here.

**Sue Nelson**
Later in the podcast we'll also be hearing from someone I know you're familiar with, its technology entrepreneur Azeem Azhar, so first of all I just want you to listen to how he defined AI: "artificial intelligence is this catch-all idea that somehow we're going to get computers to undertake some new tasks, tasks that we often thought were things that humans were exceptionally good at. The reality is that we should think about this as being an increase in our basic cognitive abilities. I think over the next few years, what we're likely to see is that, things that humans could do – like recognize what's in a picture or translate between languages or transcribe text – we're going to be able to get computers to do and that's what AI will mean for most of us over the next few years." Wendy does that tie in with how you view AI?

**Dame Wendy Hall**
Yes and no. That's a very academic response, and I kept all phrase it is because the meaning of AI has changed over the 50 years since it was first used. Ideas about machines that could think or outperform humans have been around in science fiction for a while, but the whole area started, really, with the study of [what] we now call AI artificial intelligence started with Alan Turing in 1950 when he published his paper which effectively was a paper about can machines think. He really got the whole idea going, because he this was before we had computers of course, and he was envisioning a world where we have computers that could think as human beings could.  Now that is what started the idea and the term was coined in the States in the 1960s, but today when we talk about artificial intelligence, we really mean machine learning, and they say this has come out of the neural network advances. It's all about machines being able to process and analyze a lot of data very quickly, to come up with answers, decisions, recognize a face – the things that are seen talked about. So it's moved that bit of AI; it's not about "can machines think like human beings", it's "can machines do certain tasks better than human beings faster and better", and it's interesting as Azeem used two examples – one is face recognition, now when I started in computer science in the mid 80s, lots of people said face recognition by machine was going to be impossible. One of the wonderful things about the brain is we can look at a picture and work out what it's about. People were saying in the eighties that machines will not be able to do this and today we have machines that, because they've been trained on lots and lots and lots of faces, and the algorithms are better the faster machines, the world has technology that can recognize a face sometimes even from just the eyes. So even if people are wearing a mask the machine can improvise a face which a human being can't do. Then what Azeem also talked about was natural language processing and speech translation, which is something human beings, if they're specially trained, can be very good at. That actually is easier for a machine, because it can be trained on lots and lots of data and can do that machine translation much faster than most human beings can, so that's an example of what machines can do better than almost most human beings on the planet.

**Sue Nelson**
That's right actually, because I use a program that transcribes radio interviews sometimes. Years ago they were useless, and although, yes, you sometimes get little mistakes, they're remarkably right, and they save time.

**Dame Wendy Hall**
I go to China a lot, and I'm never going to have the time, or the you know sort of ability to concentrate enough, to learn Mandarin, ever. But I can use a translation program like Google Translate to actually show someone a word which is in most circumstances enough to get you by. I'm looking for the station – translate the word station into Mandarin, and they can help you. That is amazing, and that was a very specialist talent before that we now have on our mobile phones. But face recognition – as I said, all human beings apart from the people who have memory problems or face blindness problems all human beings can recognize faces so you see my contrast here.

**Sue Nelson**
So is this why with capture, then, they make you tick the boxes with the traffic lights?

**Dame Wendy Hall**
Well this was what Azim said on the panel we had the Grand Challenges summit the other week. I'm so naive – I hadn't realized that when Google use capture to test whether you're a robot, and give you a set of images and ask you to pick out the traffic lights or the mountains or the road, I just assumed that was checking whether I was a robot or not. But actually, what Azeem told me was they're using it to train their algorithms for their automated vehicles. I should have known that, because people have used capture, which was something that's been around for a while – came out of research, Luis von Ahn invented it ages ago – to help do machine translation of the handwritten document scripts, so I hadn't thought that was what Google was using it for. And this is part of the problem, is that ¬–

**Sue Nelson**
It is, isn't it? Because people don't know that and you're not told why you're doing it.

**Dame Wendy Hall**
I'm developing a new seminar which I'm going to call "The future's fake", because actually we can fake anything now. I mean video, text, sounds, music, you really can fake anything. I think we have a serious issue in the future. What? Who? Which sources do you trust, and how do we get around this problem that anybody can fake anything? In all sorts of walks of life if that's going to cause really difficult problems.

**Sue Nelson**
I think before we get on some of the negative or more ethical issues of AI, it's good to hear something like translation or facial recognition. People can see how there's a very positive use for these technologies but perhaps not realizing that it is AI. Is there any other sort of area within society that people are perhaps using on a daily basis or a regular basis without realizing that it is AI.

**Dame Wendy Hall**
Well, Google. Just use a search engine. When you get that list, that recommended link, the list of hundreds of links – that's all done by using AI, really.

**Sue Nelson**
That could be a help or a hindrance, of course.

**Dame Wendy Hall**
Well, because it's a recommender system. They do use machine learning these days, of course, but when they started out it was a recommender system, it was an expert system and, basically, we actually train it. So every time that you use a search engine like Google, and you click on a link that's offered to you – and of course sometimes it's advertising, that's another whole story about the monetization – but you click on a link that's offered to you and then your response to that, whether you actually look at that page or not and stay on it, Google will register that and it will use that, what you do, to help the performance of its systems. So, basically, if you like something by clicking on it, that moves up the ranking.

**Sue Nelson**
Do you think people sometimes confuse AI with automation?

**Dame Wendy Hall**
Yes. Because often the word automation is often linked to robots and manufacturing and I think that, you know, for a long while we've had automation in industry that is not necessarily AI. I mean the last Industrial Revolution, for example, and the technical revolution generally, my father was an accountant and did everything by hand – dividing pounds, shillings, and pence by pound, shillings, and pence – all the ledgers had to be duplicated by hand. Hundreds of people managing the accounts for a big company by hand. By the time he retired they were using calculators, not computers now, of course, all that's done by computers – but we have more people working in the finance industry than ever before. So what that technology has done has enabled – the hard grunt work is done by machines – it's enabled the creation of lots of new jobs that leads us into all the things that we've had problems with, with the financial crash where we're not quite sure what the computers are doing. You always go from the "this is fantastic", so can we actually understand and keep control of what is being automated. So, yeah, I think automation is mixed up. They are of course synergistic, in a way, and robots will all use AI if they're doing anything other than just mechanical.

**Sue Nelson**
This is partly why AI is sort of catnip for the media and the coverage is not always helpful, I think, for the public.

**Dame Wendy Hall**
When you click on the terms and conditions to use something – we all know that we click on that because otherwise we don't move forward – we never read them. Then you get the cookies, and people just click say "yes, ok". Now, cookies use AI basically because that's really picking up your profile and using your profile to determine what information you get given. It's useful for us, but it also builds a profile in the company database and that's what we need to worry about what they do with it. On the other hand, face recognition some people don't know what... The point I'm trying to make is: people don't know what they're doing when they're clicking on a cookie.

**Sue Nelson**
It's true, but you don't have a chance – you don't get to the article if you don't click accept.

**Dame Wendy Hall**
We all see the scare stories about what happens in China with face recognition and a surveillance society, but increasingly we're using it. Our border forces are using it, our police are using it, because it's a really useful tool. When I come back into Heathrow, I put my passport in the machine, and it picks up on my face. That's what it's checking. So, effectively, we're building global databases of people's faces that the machines can recognize, and people will understand that. I think this is where we're going to start to see backlashes. Because people

understand that their face is being used to recognize them and that will frighten people. And [it] might move us into a more of a debate about who does what with that image of you and who are they allowed to share it with.

**Sue Nelson**
Well this is what Lord Browne who's the Chairman of the Queen Elizabeth Prize for Engineering Foundation, had to say about AI in our very first create a future podcast:

**Lord Browne [Audio Clip]**
"AI is lifted in a sea of hype and so people get worried about the hype leading to conclusions that may not be real, that robots will eat us all or we would all be extinct because machines will take over. This is far from the truth. AI is very important, and will be developed, I'm sure, for many applications. The idea of general intelligence being artificial will require us to first define what we mean by general intelligence. No one's got there yet, so we've got a long time to go – maybe [to] infinity. I like to think that a lot of these scares need to be put in place by people who understand a little bit more of the context."

**Sue Nelson**
Well I'm with someone who definitely – hopefully, I hope – understands a lot of the context of this. Wendy, what's your definition of general intelligence as opposed to artificial intelligence?

**Dame Wendy Hall**
You usually talk about general artificial intelligence, but maybe he's alluding to how to really define intelligence anyway, which is always the problem. How do we how do we define intelligence in terms of what we do as human beings? What makes us intelligent, and do we have additional features like a conscience that machines won't have? So, I have a lot of sympathy with what Lord Browne said, and we are a long way of having any idea how to build software systems, a machine that has any sort of general intelligence. AI is going to be, as he said, incredibly important. I wouldn't have done all the work I did with the review and the sector deal if I didn't believe that. I think AI can make our lives a lot better. There will be job losses, but there will be huge job creation; so many people can get involved in AI. You don't have to be a machine learning programmer, and it will help us in so many fields like health and education and energy and transport. It is going to transform our lives. But I do think we have to be wary of the future, because I think there are futures that could be unpleasant for us before we get to general intelligence if machines dominate in terms of decision-making without keeping the human in the loop. What's more, the issue of understanding why machines made a decision is here today.

**Sue Nelson**
And that's something that frustrates a lot of people is that you don't get to speak to a human being or because it's [not] a decision by a human being.

**Dame Wendy Hall**
Yes, and we're talking sort of about two different things but you naturally go to this. In my future is fake, I know you're a human being because I was sitting here looking at you, but generally if when you're talking to something you can't see, you have no idea if that's a human. In the near future you won't know whether that's a human being.

**Sue Nelson**
That's no different to the early days of the internet. You don't know whether the person you may be speaking to via – whether it's a dating app or just email, or somebody doing a chat bot for a store or a company – whether that's a real person.

**Dame Wendy Hall**
But we're gonna be getting a lot more of it.

**Sue Nelson**
Right.

**Dame Wendy Hall**
And when you're dealing with a company, you don't know if you're going to be dealing with a human being or a robot. Do we need companies – when they when they're selling stuff to us, or telling us why we haven't got our insurance or why our insurance gone has gone up, or why we can't book that ticket at that price – do we need the company to tell people "actually, this is a robot, not a human being".

**Sue Nelson**
Effectively, then, we've got a really good technology, a transformative technology, that as you say is going to make our future better. The difficulty is in information and data, in terms of consent, and this "what's fake, what's not?", so they're they sort of ethical questions: how do we deal with this? I know it's something that lots of people are thinking about. I suppose it's [a question of] how do we control it?

**Dame Wendy Hall**
First of all, ethics is something that is quite an abstract concept. No one taught me about ethics when I took my driving test, but we expect a car to have ethical principles in terms of the algorithms that is used for the car to operate on our streets, deciding which people it will kill in a certain circumstance. So, there's automated car, you've got a passenger in the car in the front seat, you're coming up to a pedestrian crossing, there's a mother and baby rush out on the crossing – do you kill the mother and baby, or do you drive the car into the brick wall at the side of the road?

**Sue Nelson**
This is like an episode of the good place.

**Dame Wendy Hall**
Yes, so you know this is the sort of decision that we make in a split second. Nobody talks about the ethical [considerations], you just try and do the least damage you can in that situation. A car will have to make a decision as to what to do at that point, and it...

**Sue Nelson**
Which brings it to the fact that it's only as ethical as the people who program it

**Dame Wendy Hall**
Potentially, or the people who buy it. Are you going to buy a car where the decision would be to save the mother and baby on the pedestrian crossing and potentially kill the passenger – you? This is extreme, right? It's extreme. So that's what ethics does, it looks at what are the boundaries on this, where are the big decisions going to have to be made. So I think that we will use automated transport in all sorts of controlled circumstances, but having automated transport on our streets with pedestrians and all the things that go wrong on an open street – with cars being driven by human beings and pedestrians and fans stopping by the side of the road and cyclists – I think we're a long way off that. But coming back to this idea of general intelligence, my point was there about ethics being something that is quite abstract and it's going to take us a long time to work this through, but we can't wait for the ethicist to tell us how to do this because a lot of it will evolve as the technology evolves. So, I talk about doing things in a socially responsible way, and often it's about

making sure there is a human involved in the signing off of the process. If you've got an AI system that's looking at, say, MRI scans to detect some sort of disease, tumors, cancers, whatever, then the decision's going to be made on the basis of the results of that scan as to whether somebody has treatment and what treatment they have. Then, I think, just like you do today, doctors have to sign off any decision, somebody higher up has to sign it off as that makes sense. There's been due process, this is the treatment we will take, and I think we've got to make sure we do that; we have to apply those types of processes as we develop the AI.

**Sue Nelson**
That's quite interesting because we've had a few Twitter questions for you and one of them was... let me just find it. Here, this is Amanda Groombridge. I'm sure she's submitted questions before to our other podcast, so it is good to know we've got regular listeners interested in this. Her question is: "Can AI be used to produce human disease pathways, using AI to collate available data and develop inclusive information sharing and personalized drugs programs?" This is a future she would like to see.

**Dame Wendy Hall**
Well there's an awful lot packed into that question. The answer, in principle, is yes. AI can do that and we talk about AI as being a way to save money in the health service, to increase the quality of service for less spend because there is not an infinite pot of money. Whatever colour your politics, there is not an infinite pot of money to put into a system like the NHS. But we have more people living longer, more sophisticated ways of treating diseases, and everybody should have a better quality of life as a result of that. AI can hugely reduce costs and do things that today has to be done by human being – it takes a long time and AI will do it faster and potentially more efficiently. The research shows that AI is as good as, if not better than, human beings, but there will still be mistakes which need to be checked and that's what we're talking about before.

**Sue Nelson**
Was that a Freudian slip there? Instead of good you said God.

**Dame Wendy Hall**
Did I? Laughs.

**Sue Nelson**
Here's another question from Avril Russell. She wanted to know what will happen if we build and unleash an intelligence that outpaces human evolution? Jeff Bannis also asked what do we do if AI machines tell us AI will ultimately eradicate humanity, and Rob Spence asked whether AI can ever totally take over a system? So this is control coming back again.

**Dame Wendy Hall**
Well, coming back to what Lord Browne said, potentially it is an infinite time before AI can take over doing everything that human beings do. But, as I'm making the point, along the way there are going to be crossroads at which, if we're not careful, will give more power to the machines than we want to. What Stephen Hawking said at the extreme and others have said, you know Elon Musk I think said that if we could achieve general artificial intelligence, then machines would potentially out-evolve human beings because we're biological and we evolve quite slowly while machines could evolve much faster. But I always like to point out, if you're Doctor Who fan, the problem with the Daleks where they couldn't climb stairs? There's a long way to go with it.

**Sue Nelson**
Not the most recent episodes though! They've learned how to do it.

**Dame Wendy Hall**
But that's the point, you see. I mean there are lots of things that could wipe us out on this planet long before we can build general artificial intelligence.

**Sue Nelson**
Well that's quite cheery.

**Dame Wendy Hall**
Hopefully, AI will help us sort out some of our problems like climate change, like better quality medical care for more people. There's an awful lot to be hopeful about, but my point is we cannot, early enough, think about how we do this in a socially responsible way. Because just like climate change, we've done so much. Look at the quality of life that the Industrial Revolution gave us in terms of heating, lighting, air conditioning, and all the other things that are now draining our planet of its natural resources. We have to make sure that we develop AI in a way that's for the good of humanity.

**Sue Nelson**
Well we heard technologist Azeem Azhar earlier with his definition of AI and I caught up with him in London before a meeting just a few days ago to discuss some of the issues surrounding AI, some we've touched on today, and to also ask him, like you, some questions that were put forward by people on Twitter. Now, I began by asking him how he felt about its perception and why some people view AI as a threat, either economically or in terms of our democracy.

**Azeem Azhar**
I am worried about the misperception, and I'm worried that the media benefits from the misperception because it creates more alluring stories. AI is a tool. It's a powerful tool, but it's a tool nonetheless; it's a hammer, it's an iron, it's a knitting needle, it's a blender – that's what it is. The risks of AI come from the ways in which companies in particular, but also governments, choose to implement it and what they do with the consequences of that implementation.

**Sue Nelson**
You spoke at the global Grand Challenges summit recently and you did say that the inability to consent to changes was an issue with AI. What exactly do you mean by that?

**Azeem Azhar**
Well normally we consent to the changes in our world by a democratic process; we vote for politicians, they make some laws and, if we don't like them, we vote them out. What we're starting to see is that the way in which we as citizens access different resources is now mediated, not by rules that we've agreed on, but by private companies: the Googles and the Facebooks and the Apples. Those companies in their design decisions are essentially creating or gatekeeping our access to resources, but we don't have a mechanism by which we can really say "we want to vote you out; we don't really like this." Now, the most striking example of all of this, I think, is something that's happened in the US. Amazon sells a smart doorbell – they sell everything. The doorbell in question has a video camera on it, so, when a visitor presses the button, it sends you a video feed of the visitor. You can very conveniently from your bedroom say: "I'm not in" and tell them you're not in, but it's very convenient, very helpful. In the US, Amazon has teamed up with more than 400 local police stations or police services to create a virtual surveillance network. So now, there's a police surveillance network running in those communities where the community members didn't get a chance to check into that, they didn't get a chance to vote for it, they didn't get a chance to agree the funding, to decide what kind of remedies there

should be – it's coming through the back door. How did they consent to that technology in that service being put into their lives?

**Sue Nelson**
And does that mean that their images are on that same database?

**Azeem Azhar**
In the case of these smart doorbells, what's happening is they're capturing the images of everybody who comes past and they're capturing the images of the people who approach the doorbell. In general, those images are being put up into the cloud somewhere, but what was particularly difficult I think with this surveillance network is that it is a network; it is aggregating the images and it's putting them in places where they could be inspected by people other than the homeowner.

**Sue Nelson**
Karl Kyrne on Twitter asked: What's been the biggest AI breakthrough that you hadn't expected?", which I thought was quite an unusual take on it.
Azeem: Those are always the hardest questions. I think the thing that's caught me most by surprise has been the way in which the AI ethics community has coalesced and started to be a force. Back a few years ago in 2015 or 2016, one of the things I was concerned with was the way in which AI systems were coming out with lots of biases – gender biases, race biases, age biases –lots of problematic ways in which they were being implemented from a purely engineering perspective, and that gave rise to lots of discussions around AI ethics. By 2018, so only a couple of years later, there was a very strong AI ethics movement. Large companies were saying we need to think about this, the researchers were thinking about this, there were programs in universities, and I think the activists, academics, and AI specialists in that field did a really tremendous job in not just raising awareness to the issues, but also galvanizing a meaningful response so that those considerations are now in the design and engineering modalities by which these products come to market.

**Sue Nelson**
So effectively you're saying that people responsible for AI didn't necessarily think through all the implications and, I suppose this might sound a bit rude, they're making it up as they go along.

**Azeem Azhar**
Well, yes. I think back in 2015 or 2016 they weren't thinking he through the ramifications, and I would have predicted that things would be much much worse and we would be having to shout much harder by 2019. The problem is not fully addressed, but we're in a much better position because the AI ethics community managed to come together and tell stories that people could understand and activate us to say: "Look, these are important issues, and we have to figure out how we address them." So, the good news is that while the battle is still to be fought, the troops have been assembled and are ready to fight.

**Sue Nelson**
Sarah Rose Gregory basically picks up that point. She asked on Twitter if AI is only as good as what the're programmed to be. What safeguards are in place for preventing replicating the structural limitations – which you sort of touched on – but not just that, she said supremacy and patriarchy as well.

**Azeem Azhar**
We decide the values – humans – and so we have to have the debate about the values that matter. While it's true that AI can absolutely reinforce the structures of the past, it can do something else, I think it's really interesting; it can lay bare, lay transparent, the framework or the architecture of the past, and we can expect

that and say whether we really like it or not. I'll give you a very simple example: some academics used some machine learning techniques to understand the relationships between words, as discovered in written English going back many many many years. They looked at millions and millions of words, and they discovered that there were relationships such as man to woman was similar to King to Queen but it was also similar to doctor to nurse. What we're seeing is, unimpeachably, that the corpus of English language documents captures this structural bias which suggests that if you're a healthcare professional and you're a woman then you're a nurse, and if your healthcare professional and your man then you're a doctor. The great news is that people used to argue about whether that really was a bias that was evidenced in the texts, and we now have evidence that it is a bias and we can move on and we can say “do we like that or do we not” and I think we can use that evidence to say “let's now fine-tune our systems to, you know, represent the world that we want to live in”.

**Sue Nelson**
That ties in actually with what Rob Spence wanted to know, which is whether AI ever totally take over a system, because he thought that a human surely must be necessary to prevent disasters. I know this leads into this ‘AI is going to you know take over the world’ aspect, but it is a genuine concern.

**Azeem Azhar**
It is a really valid concern, this question of control. We already have designed systems where humans have very little say in what actually ever happens. If you look at a company that does a bad thing – for example, they cut someone's line of credit off early – try to get that company to tell you which person was responsible for cutting that line of credit. They will not give you an answer, because they won't be able to give you an answer, because they've created a set of processes internally that remove human agency and individual human responsibility from the decisions that are taken. [Looking at] the way that stock markets behave, we start to see these accidental, what are known as ‘flash crashes’, when very simple trading algorithms go awry, are often out of control of human systems. So we've already started to build and live with systems where the individuals responsible say: “listen, I don't have that responsibility” and I think what we can do is we build AI systems into the world going forward is saying that, actually, it's really important to maintain some kind of thread to a human who is responsible and will say: “listen, the buck stops with me for this decision”.

**Sue Nelson**
Which leads us to an issue of what rights those of us who use AI actually have when we're dealing, like you say, sometimes with machines rather than human beings.

**Azeem Azhar**
I'm gonna take that question back away from the simple problem of machines. I mean, in a sense a company is also a machine, because a company is a set of processes that passes information around and comes out with particular outcomes. I mean, abstracted away, it's identical to a physical machine, and we've allowed companies to have rights and responsibilities that are often in conflict with individual people. Now, we can't extend that to AI systems. We shouldn't allow that to be extended to AI systems, and actually, we should go even further and say that if it doesn't make sense to extend to an AI system, then we shouldn't have ever allowed companies to have those rights and responsibilities visa vie our own individual rights in the first place. So again, I see the challenge with AI here as actually being an opportunity to go back and say: “look, we have allowed these non-human institutions like technical systems and companies to have too much power relative to the individual”.

**Sue Nelson**
We're in an era of fake news where people are worried about the algorithms that certain social media companies use because they affect social change, they affect democracy, they change people's views, so in that respect the AI is only as good as the morality of the people who construct it and who use it. Do we have or

should there be some sort of a regulator, I hate to say Ofcom or something like that, that actually says: “no, this is not right”, and who decides that?

**Azeem Azhar**
Well, we're in this odd situation today where it's being decided by very bright people in a handful of American companies, and one can just look at that and say that's not really the way that decisions should get made. Decisions like this should be made in the realm of politics, which doesn't necessarily mean politicians, it means people who are interested, who have different perspectives, who can participate, and where the goal is not simply one of profit maximization, where the goal is in some sense a societal goal that gets articulated through that. Now, whether we do it through a regulator or through some other mechanism – I'm reasonably open-minded – but it's clear that society needs to set the rules, companies need to implement those rules and have systems that ensure they're complying with them, and someone else has to check that companies are adhering to those rules and then be able to enforce any sanctions that are required. Maybe that's a regulator, maybe that's some other type of body, but we need to separate out who does what.

**Sue Nelson**
Basically we as a society need to engage more, see its benefits, but also see where we have to, when necessary, take control. Oh I hate to use that phrase ‘take control’, but to limit it where necessary so that we're collaborating I suppose rather than more of “it's a tool”.

**Azeem Azhar**
What's wonderful about the Royal Academy of Engineering is that it's about people engaging with these technologies which, as a whole, society has not done. Really, since the 1960s we have moved into a much more technocratic view of the world. We've said there are experts, you go off and figure these things out, and we as society are not going to engage with those experts; we're going to move into the Big Brother, Celebrity Love Island, Kim Kardashian economy rather than an economy of inquiry and examination. Technology is far too important to leave to technologists; consumers and citizens have allowed themselves to be lulled into this sense of immediate hedonism – press a button on your smartphone and the pizza appears within six minutes – and the problem is that there are no free lunches. The result of us not keeping an eye on what's going on and not being active is that we find ourselves in 2019 saying: “wait a second, I don't have the control that I thought I had. These things are being done to me and I don't seem to have a voice in it”. I don't think this problem is solved solely by slapping technology companies, which I think is something that they often deserve, it's also addressed by individual members of the public starting to say “I'm going to understand the detail, I'm going to understand the debate, and I'm going to participate”.

**Sue Nelson**
So, how do you see the future of AI? Do you see it as something that we will get right because we're having those conversations now, or do you see it as being on a sort of knife edge effectively where it could go down a slightly more sinister route if we're not careful?

**Azeem Azhar**
I don't think we've got much risk of the sinister, you know, glowing red eyes route that often is protect portrayed in the cinema. I think the real question is going to be about power, and during the industrial age we had industrialists who took power and we had industrialists whose power was kept in check by people and by politics and by activism. We're going through a similar transition now, and the question is not whether we'll have plenty of AI and whether it'll be wonderful for us so we will have plenty of it and it will have the potential to be wonderful, it will be really about the nature of how we divide up the rewards: financial, non-financial, and, ultimately, the power in that system. If you're an engineer who's building these types of things, that may feel

curiously divorced from your day-to-day but, in reality, it isn't. In reality, you build the system that is at the coalface that will be used in the case, for example, of a translation system billions of times a day by people from their smartphones. You have to think quite hard about the choices that you're making and the way in which you're trying to keep your organization in check as well.

**Sue Nelson**
Azeem Azhar, do you agree, Wendy, with Azim when he said that technology is far too important to be left in the hand of technologists?

**Dame Wendy Hall**
I think I said it in the first place. I quote Karen Sparck Jones, my mentor and friend from Cambridge whose work on information retrieval – she's one the women who's not cited enough in history – whose work on information retrieval is in every mobile phone, every speech recognition system you use. She had a phrase that computing was too important to be left to men and I remember wearing that t-shirt. That's not to denigrate men, it's to say it's too important – everybody needs to be involved. The phrase I've used for AI is that AI is too important to be left to the AI experts; we absolutely have to think of it in a sociotechnical way. Everything we do should be tested in terms of how socially responsible it is or may be.

**Sue Nelson**
Well Azeem just briefly did sort of pick up on that about structural bias and when you were both talking at the Global Grand Challenges Summit recently you were one of the speakers who was very strong on diversity and the development there. Why do you think it is so important to get diversity involved, particularly when most people think of the Facebooks, the Googles, the Twitters? It's all men [there].

**Dame Wendy Hall**
Well you've just answered your own question because of the nature of those companies reflect the nature of the people that founded them. They're not necessarily bad people, but that they have their way of thinking and that's how the company has evolved its way of thinking. It's important with any technology, particularly computing technology. We have so few women in the west involved in computing, and even around the world women involved in actually developing the technology. Loads and loads of people are trying to change that and get girls interested in coding and in working in the industry generally. It's so important, but for AI it's even more important because of issues of bias: because of issues of bias in data, bias in algorithm, and we are all biased in the way we think. Machines learn from what we do, so if they only learn from a small section of society, then they won't be fit. This development will not be fit for the whole of society. So even if we haven't got a lot of diversity in the teams doing the programming, we need diversity the teams doing the design, the testing, evaluating the behaviour of people when they use the system – we need that to be diverse. Not just in terms of gender, it needs to be diverse in terms of culture, race, ethnicity, religion, age, disability – so many different factors, and we've got to make an inclusive industry.

**Sue Nelson**
I mean I agree with you there, I made a Radio Four program a couple of years ago with Anjana Ahuja whose an FT columnist, and we were both shocked when we discovered that the early algorithms that Facebook was using didn't basically recognize dark skin. So, consequently, when there was an image of somebody that was black they were identifying them as a gorilla – which is shocking even now just saying that, I still can't believe it – but it does show you that the people who were programming and doing the algorithms were probably basing it on images of themselves, and I can only assume that it was images of white men. So how do you do this, how do you change it, though? It's easier, I mean engineering has done its best – is doing its best – to get more women involved in engineering, but it is easier said than done. So how do you get women into AI?

**Dame Wendy Hall**
I've been trying to do this for 30 years, getting women into computing, and I feel like I failed miserably. We are making some inroads, but we still have very few women coming to read computing at university and so the pipeline to go on to advanced courses in AI is tiny. There are so few women going on to those. There are many different ways of tackling it, there is no one answer and clearly nothing has worked so far. It's a culture change, and actually changing a culture, when it's so fixed, you have to got to change the children themselves and students, the teachers, the parents, the industry, the universities. I was in Southampton yesterday and we have a record number of students – it's scary – going to do computer science. Our AI MSC is bulging at the seams, we've taken two or three times as many students as we would normally, but I'll bet you the diversity is really low. But it's not the same all over the world. I go to India, I go to Malaysia, I go to Singapore, and there's lots of women getting involved in computing (and therefore will go into AI), but in the west particularly it's really bad. Now, what we've done in the sector deal, well in the review we did for AI, we actually put in [the need] to set up new PhD schemes and new master's programs for getting machine learning programmers into industry. We also put in a stream of conversion courses starting at master's level, but could also be continuing professional development, where we ask the universities and colleges to set up programs for non-science and engineering students to learn how to get into the work of AI

**Sue Nelson**
I think that's a brilliant idea

**Dame Wendy Hall**
... and we got 18 million pounds in Teresa May's last giveaway budget at the end of her term, and you'll [start to] see stuff coming out for these conversion courses to be developed at our universities and colleges, and I'm so excited by that. That scheme is going to develop over the next year, and hopefully the courses will start next October.

**Sue Nelson**
Brilliant.

**Dame Wendy Hall**
It's a drop in the ocean, but...

**Sue Nelson**
I know, but it's a good start. Now, I've got time here for a few more questions and they were asked by Twitter. It was great, actually; we got an avalanche of responses. One of them from Michael field was about how to remove bias from AI, which we've already addressed.

**Dame Wendy Hall**
Yeah, adversity diversity diversity.

**Sue Nelson**
Tere are a few that I put to Azim as well, the first one I put to him was from Claire Ainsworth who asked whether AI could ever become so sophisticated we had to consider its welfare as well as our own. She admitted that this was a speculative question, but I rather liked it.

**Azeem Azhar**

We consider the welfare of things when we start to get a sense that they may have a little bit of agency, they may have nervous systems, they may have the ability to plan, and the ability to feel pain and socialize, and that's why we start to consider the welfare of certain classes of animals. But less so with other types of living things, and so you could imagine that as AI gets more and more sophisticated, and it starts to have those attributes in ways that we can sort of tangibly put our arms around, we might start to think about that. But the other thing that we might want to consider is that, once AI systems become part of our ecology, part of our ecosystem, we might consider their health in the same way as we consider the health of the soil or the water runoff in a field – as an important part of that local ecology. So I could imagine we could consider the health of our AI systems and in a similar fashion.

**Sue Nelson**
Science is Simple said how do you deal with uncertainty in AI while we accept that the human mind is flawed? We are very quick to judge a machine that makes the quote 'wrong' decision.

**Azeem Azhar**
I don't know the extent to which I judge machines for making the wrong decision. I mean, judgment is something that you have to apply to things that are sentient, that have agency, that have decision-making capacity, so I don't judge my hammer when I miss the nail. What I can do is I can say the system has been designed poorly, and implicit in the word design is a human being or set of humans who have allowed it to get into the wild.

**Sue Nelson**
Jessica Tye asks are there any areas or task types etc where you don't think AI should be used?

**Azeem Azhar**
Well, AI today is full of the shortcomings of the people who designed it and I'll give you a simple example of that: you've essentially got an engineering system that you're testing against a bunch of scenarios to see whether it does a good job in them. Testing takes time, time costs money, so there's always a trade-off if you're an AI developer between how much testing do you actually do before you push something out into the wild. So, the more we care about the outcome, the more it deals with situations are not very resilient or people who are a little bit vulnerable, we have to think "has the appropriate level of testing been done to make this system perform in an appropriate manner?"

**Sue Nelson**
Now I got a question that, I must admit, I didn't understand myself because it said is Gödel incompleteness relevant to AI? So, you have to explain to me what Gödel incompleteness is.

**Azeem Azhar**
So, it's this idea that you can't define a system of logic that includes itself, so you always have to have some external reference. I think it is – I'm not a sufficiently strong mathematician to talk about that – what I do say and I think probably have come across in the conversation today is that AI is a really powerful and relevant tool, and our discussion should really be about how we put these powerful tools into the hands of enough people, and the right people, to get to the kind of society that we want to get to that addresses our sustainability issues and equality and distribution, and so on. There are a whole bunch of incredibly interesting philosophical questions, of which this is one, but in general not ones that will have an impact on where we go over the next four or five years.

**Sue Nelson**

Okay, right. There's Azeem under the spotlight there. Now some final Twitter questions for you. This was from Southwark Bell, as an AI professor she was curious to know whether you have Alexa in your home?

**Dame Wendy Hall**
No, and I was thinking about that last night and we haven't talked about much about privacy and data sharing. There was a previous question I was going to pick up on about how we do data sharing in the health service to enable companies to access to the data they need to build the algorithms. It's a really difficult problem and I'm quite wary of having one of those in my house because I just don't know who's listening.

**Sue Nelson**
That's safe, coming from an AI expert.

**Dame Wendy Hall**
Well I know how they're built, and I know that we have no control, where we are, over what the Googles and the Amazons do with their algorithms.

**Sue Nelson**
There was an interview with, I think it was Mark Zuckerberg, and they noticed that on his computer behind him he'd put a little bit of Blu Tack over his camera.

**Dame Wendy Hall**
Exactly! I don't know if there's time for this but there's a new thing all the companies are now offering us: two-stage authentication. The basic way they're offering this is you put in your password or whatever and then they send an SMS message to your mobile phone, which you then get a number to type in. It's a two phase or authentication. Now, that means they've got your mobile phone number, and we found out last week that Facebook are giving people those mobile phone numbers.

**Sue Nelson**
Right, yeah it's power and control, isn't it?

**Dame Wendy Hall**
Yes, and all the security experts say that, actually, this is about control, because they just got another piece of information about you. So that is not the way we should be doing two-phase authentication, and that was a lesson for. And Alexa and Siri, I mean.

**Sue Nelson**
You know I've almost got material for a whole new podcast on this actually. Here's a question from Debbie Jones which fits into this – she asks whether it's AI that we have to fear, or the motives of the investors and those who own the payroll?

**Dame Wendy Hall**
What's so interesting is that we need companies to sign up to certain ethical standards for AI, or socially responsible stands we might start off with, because we need to understand there has to be transparency in what they're doing with our data. In Europe we've got GDPR. In Europe a company would have to tell us what they're doing our mobile phone number. That has to be explicit. In America, they have to say they're compliant, but Facebook and GDPR are miles apart. For example, and I mean GDPR is actually, probably too stiff a regulation to enable innovation, but we need to mandate some regulation in this space about data, across the globe, and I would include China in that. China has got to be at the table. There's so many people living there,

and we're going to increasingly buy AI services and technologies from China and India.  I think the key thing is, though, it's easy to pillory Facebook and say it's all their problem. A lot of people use Facebook and love what it enables them to do, and Facebook has to make money. It's a company, it has shareholders. It's no good saying a company shouldn't make money, but it is about signing up to some ethical framework, and we have to agree there are some emerging global ethical proposals for global ethical frameworks which are very embryonic, will be very simple to start with, but I really think our governments have got to push for this.

**Sue Nelson**
Professor Wendy Hall, thank you so much for joining me. I really enjoyed that, and my thanks also to Azeem Azhar, the technology entrepreneur and producer of the Exponential View newsletter and podcast. Thanks for listening, and do join us again for the next episode of the Create the Future podcast.