# Science and Technology Committee

## Oral evidence: Robotics and artificial intelligence, HC 145

## Tuesday 28 June 2016

Ordered by the House of Commons to be published on 28 June 2016.

Written evidence from witnesses:

- Professor Rose Luckin
- Deloitte
- Innovate UK
- Microsoft
- ABB

Watch the meeting

Members present: Nicola Blackwood (Chair); Victoria Borwick; Jim Dowd; Chris Green; Dr Tania Mathias; Carol Monaghan; Graham Stringer; Derek Thomas

Questions 83-159

Witnesses: **Professor Rose Luckin**, Chair of Learning with Digital Technologies, UCL Institute of Education, University College London, **Angus Knowles-Cutler**, Vice Chairman and London Office Senior Partner, Deloitte, **Dr Michael Osborne**, Dyson Associate Professor in Machine Learning, University of Oxford, and **Paul Mason**, Director of Emerging and Enabling Technologies, Innovate UK, gave evidence.

**Q83 Chair:** I welcome the panel to our second evidence session on AI and robotics. We had a fascinating discussion last time with some very real dilemmas and questions being raised by our panellists. Can I start by asking the witnesses the question I put to the previous panellists? Paul Mason, how close do you think we are to seeing intelligent, general purpose machines that can handle complex problems?

> ***Paul Mason***: It depends on what you define as a complex problem and generally how widely available things are. There are systems in manufacturing now that have complex control systems that are quite widely understood. If you go on to the internet and engage with pop-up chat advice, and you think there might be a person at the end, quite often it is a computer. You think you are interacting with a human being but you might well be interacting with a chatbot. Sometimes it gets the point of your question; sometimes it does

not. Machine translation is quite common in internet searching. Some of it is there, but some of the things you see in science fiction films are probably still a long way away.

***Angus Knowles-Cutler***: I concur with Mr Mason. We are moving apace towards that, but the impact of technology on the UK economy and UK jobs is already very significant. Sometimes we are very interested in AI and advanced robotics, and there is a risk of not seeing what is happening in the workplace when tasks within jobs are being automated and replaced. That is happening now. It is better not just to look at complex AI and very advanced robotics but to understand that what is happening in the workplace is almost invisible and has been happening substantially over the past 15 to 20 years and is gathering pace.

**Q84 Chair:** That is a very helpful comment. Professor Luckin, perhaps we can look at this through the lens of how close we are to developing these systems, but also what will be the greatest potential benefits for society over the next decade as a timeframe.

***Professor Luckin***: We are very close to developing effective AI systems that do specific tasks. We are a long way from developing AI systems that have general intelligence. They are two very different things. In order to make sure that the specific tasks AI is capable of achieving are the ones society needs, we have to address the human needs before the technology. We should be driven not by what AI can do because it is smart but by the challenges we face as a society and the way in which AI can help deal with some of those specific challenges, such as those in education.

**Q85 Chair:** I think it is called AIED, isn't it? Dr Osborne.

***Dr Osborne:*** On the general point of how far away we are from general purpose AI, I completely agree that it is a very distant prospect. The advances we have seen in the recent past have been very much about the narrow definition of a task to which an algorithm can be applied, so the progress we have seen in AI has been less about fundamental advances in the algorithmics and much more about working out ways to apply within a narrow, confined domain the technologies we already have. This discussion is as much about how to restructure work and highlight the kinds of tasks Angus already mentioned and to which AI can be applied as it is about the advances in the field of machine learning and artificial intelligence itself. Part of that is about the increasing availability of data. The best way crisply to define a task such that it could be addressed by an algorithm is to gather a sufficiently dense and comprehensive dataset. We have already heard the example of speech recognition and machine translation. These are excellent examples where for a long time the problem was seen as beyond the scope of automation. The complexities inherent in translating text from one language to another were seen as being beyond that which could be incorporated into an algorithm, but once we started getting sufficient data, and getting enormous corpuses of text pre-translated into a multitude of languages and lots of recordings of speech whose meaning was exactly clear, we were able to make progress towards that.

**Q86 Chair:** If we take a step back for a moment and look at all the invisible ways in which AI is already impacting on our workplace and daily lives, to what extent are the opportunities

in robotics and AI being fully capitalised in UK industry and Government? Paul Mason, are we making the most of those opportunities, or are we risking losing market share?

***Paul Mason:*** It is a good question. I would probably separate the physical from the cyber in terms of technologies and applications. In some areas, the UK, certainly in advanced robotics, is quite strong as a user of such technologies. Modern car manufacturing facilities, the nuclear reprocessing industry, the defence sector and those sorts of things are pretty well advanced in using those technologies, and they are very productive. In other parts of the economy those technologies are not so much deployed, and they are slightly newer. There are opportunities and we are good in places.

In terms of investment in the future, outlook for the UK is a more mixed picture. The number of robot systems being installed in the UK relative to other nations around the world is quite dramatically less. The numbers for last year show that installed shipments of robots in China were 75,000 compared with 2,400 in the UK. When those things roll through to manufactured products, and products become more reliable and less expensive, it creates a threat to our industries that have not invested in that way. We need to make sure we adopt these technologies to protect our existing industries, as well as having the conversation about the impact on jobs and about the growth of new industries that these technologies also offer. I could probably say much more but I will stop at that point.

***Angus Knowles-Cutler:*** Talking broadly of technology, automation, robotics and AI, at Deloitte, in collaboration with Mike and his colleague Carl Benedikt Frey, we have taken a look at what has happened in the UK in the last 15 years and can probably draw some conclusions from that. The conclusions on the impact of technology on the UK workplace are very positive. We looked in detail at 360 job types and at ONS Government data. The conclusion is that while technology has destroyed about 800,000 jobs in the UK in the last 15 years—half the secretarial, travel agent, counter clerk and librarian jobs have gone—technology has, directly or indirectly, created another 3.5 million jobs in their place. The good news is that each of the jobs that has been created pays on average £10,000 better than the jobs that have gone. If you look at it across the country, every nation and region of the UK has had a net benefit from technology. The technology-driven jobs that have been created fall into four categories: technology jobs themselves; creative jobs; professional business service jobs; and a lot of caring jobs. Therefore, we can draw some comfort from our recent and long-term history.

The work Mike, Carl and I did tried to project that looking forward. Trying to look 10 to 20 years down the road is difficult territory, but our estimate was that about 35%, or a third, of today's jobs in the UK could be automated in the next 10 to 20 years purely from a technology point of view. It does not factor in social and political resistance to that change, the ease or difficulty of implementation or the cost-benefit of human labour versus investment in the technology in the first place. Those would be brakes on getting to that 35%, but if those brakes come off it will accelerate, so there is a lot of potential. One comment to make on the UK economy related to the latter point is that while there is a ready supply of human labour it slows down implementation of technology, but if human labour becomes in shorter supply or costs more, it could well accelerate progress in adopting those technologies.

**Q87 Chair:** We are legislators and we are writing recommendations to Government. Should we be seeking to intervene to govern this technological revolution in any way? Should we be getting out of the way? What role do the Government have in trying to prepare society for these changes as they come along, Professor Luckin?

***Professor Luckin:*** Before I answer that question, I would like to say something about AI. AI is not just a technology; it is a way of thinking and looking at the world and a methodology for developing systems that can behave in an intelligent way. It is more than machine learning. Machine learning is just one way of implementing artificial intelligence, as are robotics and autonomous systems, so we need to think of AI more broadly. It often gets talked about as either robotics or machine learning.

In terms of society, there is a role for legislators. Going back to your original question about losing market share, we are at risk when it comes to the educational technology business specifically in relation to artificial intelligence in education. That is not the same as educational technology or learning technology; it is a very specific use of artificial intelligence, whereby AI systems can explain the decisions they are making. Other countries are heavily investing in this area, particularly in Asia and the US, and to a certain extent in Australia. I am not aware of any significant funding for artificial intelligence and education research and development in the UK. We have small amounts of money from Innovate. The funding that the research community has taken advantage of to hold its position internationally has all come from the European Union, and that is a serious issue for us at the moment.

We have the most vibrant educational technology start-up community in the UK, in particular in London. This community is crying out for advice from researchers. They understand the technology of AI but not the pedagogy and learning. Every week a new company comes to my door to ask for advice, help and collaboration. We need some way to support those collaborations, because these are the companies that will help to feed the multi-billion-pound educational export business that we should be driving.

***Dr Osborne***: If I narrow the focus of the question to considering the relative success of the tech sector in the US and the UK, one of the big differentiators has been the enormous advantage of big US firms—the Googles and Facebooks—in the enormous data sets they have gathered. I brought up data before as being at the heart of the success of machine learning, but it has also been at the heart of the success of all those firms, not least in enabling them to deploy machine learning. The problem is that, now these US firms have this enormous first-mover advantage in having gathered that data, including data on a lot of UK citizens, it is very difficult for UK start-ups to get a leg up on them. I am not sure what the regulatory solution to that is, but it is certainly an issue I want to raise.

**Q88 Chris Green:** There have always been concerns about the impact of automation or improvements in technology. It will always have an impact on the workforce. To what extent have recent advances in robotics and artificial intelligence presented a new and different challenge to the workforce, Dr Osborne?

***Dr Osborne***: This is a question we get a lot when talking about automation. My first answer is always that, even if all we see this time is similar to what we have seen in the past, it would still need to be an issue to take seriously. Let's say, for example, that is no

bigger deal than the changes associated with the deployment of technology in the industrial revolution. Of course, the industrial revolution was associated with a whole host of new societal institutions that had to be introduced, for example to tackle child labour, and we would certainly want to be aware that similar changes might need to be put in place to respond in particular to the inequality that might be exacerbated by the trends Angus mentioned before, where lower-skilled jobs were being automated away and higher-skill jobs were relatively secure.

My second answer is that the reason this kind of technology looks different from what we have seen in the past is that it is much less clear what remains as the preserve of human labour alone. Now we are seeing algorithms that can substitute for human cognitive work rather than just manual work, as has been the case in the past.

**Q89 Chris Green:** In the past, blue collar workers have been more vulnerable, and in the future perhaps more white collar workers might be.

***Dr Osborne***: Exactly. We have already seen the example of algorithms substituting for the work of paralegals and junior lawyers, accountants and auditors—sorry, Angus. Once those possibilities open up, it is much less easy to see where the role for human labour remains. With that in mind, it is something we need to consider quite seriously.

**Q90 Chris Green:** Putting it crudely, we have seen the impact of technology on blue collar jobs in the past and we know how to handle that and we know how it will go, but for white collar jobs it is completely unknown territory. That could result in all kinds of different problems. Mr Knowles-Cutler, do you want to pick that up?

***Angus Knowles-Cutler***: I suppose the starting point is to understand what machines can increasingly do—the skills or tasks within jobs that are liable to automation. They can do more and more, as I think we are all aware. The stark number that came out of the work we did with Mike and Carl was that, from the quite detailed model we built, the conclusion was that jobs that today in the UK pay £30,000 or less are five times more vulnerable to being automated than jobs that pay £100,000 or more. As we said, in the past 15 years it has been a very positive story in the UK economy. Looking forward to the next 15 to 20 years, we may be in new territory. Yet again, we went from looking at jobs to looking at industries and sectors of the UK economy. The conclusion was that no sector is immune to these changes, including professional services and business services, but some are more vulnerable than others. Our conclusion was that, purely from a technology point of view, there were perhaps 2 million retail jobs at risk from automation, probably 1.5 million jobs in transport and about 1 million jobs in accommodation and food services. You can see the sort of jobs where the risks are and what some of the regional implications might be.

Perhaps I am jumping ahead a little bit, but the answer seems to be, first of all, to know what is coming and take it seriously, which I am sure is why we are here, and to find national solutions in a competitive world in skills, education and knowledge. You cannot go far with this conversation without its being clear that in business and, I am sure, in public sector organisations we know what sort of people we need today and in the next year or two, but we need to stretch ourselves and answer the big question as to what skills,

knowledge and type of people we are going to need in 10 to 15 years' time, because that is the life cycle of people coming through school, university or retraining.

**Q91 Chris Green:** One of the concerns is what kind of transition period you have. Are technologies going to come along where whole sectors of jobs might disappear and are no longer needed, or is it more likely that aspects of people's work will change fundamentally? Therefore, many people can retain their job but they just go about it in a different way using new technologies.

***Angus Knowles-Cutler***: Implied in the question is a very important observation, which is that we often think of this as binary, so it is a human versus a machine or a robot. It is subtler; it is tasks within jobs. One of the reasons I became interested in this a couple of years ago was that I mapped out my working day when I first started as a graduate analyst in 1985. In my 12-hour day I went to a physical library and looked at things called microfiches. Some people may know what a microfiche is, but it is a very generational divide. Six librarians would help me. We did not have a computer, so I worked out everything on a hand calculator ledger. That was a 12-hour day. Last year I got one of my graduate colleagues to do the same 12-hour day. I was never the swiftest, but she did it in 40 minutes. The very mundane, repetitive tasks had gone out of my day, but the reality is that it did not destroy my industry, and in fact my industry is much larger than it was back in 1985, so there is a subtlety there that is very important. These are tasks where technology is enabling us to be more effective as productive workers.

**Q92 Chris Green:** As is highlighted, there are many millions of jobs under threat. At the moment, are there any particular industries on the verge of getting to the tipping point, where something is coming along which means that a particular sector is under threat?

***Dr Osborne***: It would be ridiculous not to mention autonomous vehicles, which we have not yet discussed. There is excellent work going on in the development of such vehicles in the UK, witness the Mobile Robotics Group in Oxford in particular. I see these vehicles on a daily basis navigating the streets of Oxford without too many difficulties.

**Q93 Chris Green:** After black cabs we saw Uber come along, and now we should be thinking about this.

***Dr Osborne***: It is no coincidence that Uber is one of the biggest investors in autonomous vehicle research. Those technologies are not a million miles away.

Another consideration is that many of the technologies will not be able to do everything that a black cab can do currently, even with a 10 to 15-year horizon, but they might be able to do enough to substitute for a lot of the work. Even as we divide up a job into tasks, it is important to recognise that we need not complete all those tasks, effectively, to automate away the job. If I go to the example of the typing pool in the 1950s, it was not the case that the word-processing software that replaced those workers was able to make small talk with colleagues, make coffee and do all the other things, but ultimately it was able to do enough to get rid of the job. In a similar way, autonomous vehicles might not be able to make conversation with passengers or load and unload bags, but it might still be

enough, if the technology is sufficiently cheap, to replace many of the jobs, and even if not in taxi driving, certainly in the driving of forklifts, agricultural vehicles, mining vehicles and inter-city trucks. There is a whole host of occupations associated with these kinds of technology that will imminently be under threat.

**Q94 Graham Stringer:** Mr Knowles-Cutler, you have been through some very interesting statistics, and you have answered many of the questions I was going to ask. Why are we not in our national statistics seeing an increase in productivity? One of the problems with the UK economy is that productivity is not increasing, yet to listen to you, we should be massively more productive than we are now. How do I put those things together?

***Angus Knowles-Cutler***: It is a big question that a lot of people better qualified than me are trying to answer. Looking at it from the point of view of the work we have done, while there is a plentiful supply of labour that is cheaper than implementing technology, we probably are not implementing some of that technology. It is the cost-benefit between the cost of labour and the cost or difficulty of implementing technology. On the productivity question, it probably suggests that there is technology and emerging technology that is not fully implemented that would increase productivity. If the cost of labour goes up, or the labour supply becomes less plentiful, the technology is there to increase productivity. That is the closest I can get to it, Mr Stringer.

**Q95 Graham Stringer:** I can understand that in a particular industry. That has been part of the debate on the EU. Overall, you have given us some excellent examples of where investment in computers, AI and robotics has increased productivity, yet that investment is not coming out in our national productivity figures. Why?

***Angus Knowles-Cutler***: Does anybody else want to answer? *[Interruption.]* Saved by the bell.

**Graham Stringer:** You can think about it.

**Chair:** I suspend this session for 15 minutes or until we are quorate. My apologies.

*Sitting suspended for Divisions in the House.*

*On resuming—*

**Chair:** Can we resume the session, with apologies to the panel? Thank you for your patience in waiting for us. On this Committee, occasionally we find that democracy intervenes.

**Q96 Graham Stringer:** I was looking for an answer to the productivity conundrum. I do not know whether you have found an answer while we have been away.

***Angus Knowles-Cutler***: We have done our best, Mr Stringer—we passed the word. It is a big, complex question. The first thing we noted is that changes in productivity have been very different across different industries and sectors of the UK economy both since the financial crisis and over a longer period of time. The flipside, the other way of looking at the number of jobs that have been automated and that could be automated, is that, from a business or public sector point of view, there is productivity potential. In sectors where we have seen a lot of jobs go in the last 15 years—300,000 in retail and 100,000 in manufacturing—the flipside is that those are the sectors where we have seen productivity improvements. The UK car industry is a very successful one. We are one of the leading countries in the world, but with many fewer people working in car manufacturing than there were in the 1970s. Some sectors, such as manufacturing and retail, have adopted technology and therefore become more productive. Other sectors, particularly some of the service areas—business services, professional services and financial services—have been fairly anaemic in their productivity gains since the financial crisis, but there is technology there that is yet to be implemented. That is my attempt.

***Professor Luckin***: Part of the problem is that whether your job has been replaced and you need to retrain yourself or to be retrained for a different type of occupation, or whether, as we have talked about already, certain parts of your job are automated, that blended approach is going to become much more prevalent. You need different skills. I do not feel that at the moment we are equipping either students in school or workers in the workforce with the requisite skills to know how to adapt themselves to use the automation they are being offered to best effect. We need to take that on board and make some changes to address it.

**Q97 Graham Stringer:** Dr Osborne, you heard what Mr Knowles-Cutler said earlier about more jobs being created than lost. That is a very optimistic view, admittedly based on historic data. Are you as optimistic as Mr Knowles-Cutler, or do you think he is being over-optimistic?

***Dr Osborne***: I do not think he is being over-optimistic. There will be a host of new occupations generated by these new technologies. My concern is that the occupations that are generated might not be sufficiently well paid to substitute for those that are automated away. If a truck driver loses their job due to the introduction of an autonomous vehicle, they might be able to get a job in some kind of service occupation that paid a lot less than they earned in the past. The interpersonal skills and social skills that are most difficult to automate are also somewhat a commodity and hence not associated with very high rates of pay. I suggest it is that which might lead to exacerbation of inequality as a result of automation.

**Q98 Graham Stringer:** Has any work been done on the likely impact on income disparity as well as regional income disparity as we get more automation?

***Dr Osborne***: We have done an analysis of how the probability of automatability relates to income, and the trend we saw was a clear negative one: the more you earn, the safer your occupation is from automation. The regional story in the UK was exactly as you might expect. London was the most secure, and regions outside the city were much less secure.

**Q99 Derek Thomas:** Professor Luckin, can you say a little more about how artificial intelligence technologies could be used in education to equip future employees?

***Professor Luckin***: There are two things I want to stress. First, we need to change some of the things taught in schools. Of course, students need to understand the basic subjects, but the very things on which we focus our education system are the routine cognitive skills that are the easiest to automate. We need to make sure that students also gain what are called 21st- century skills. They change, and they will change, because as the workforce changes, the skills we need will change, but they are negotiation, communication and being able to synthesise multiple sources of information. It is probably also artificial intelligence and knowing how to make the most of the kinds of automated systems that will be in the workplace when you get there.

We need to think about what we change about what we teach. We also need to think about the teaching workforce. Lots of people talk about workforce replacement. Michael and I have talked about educators probably not being the most at risk of having their jobs automated, although there is some fear that that could be the case. But there is absolutely the possibility, almost the certainty, that there will be some automation of the more routine skills teachers have to use: record-keeping, marking and analysing data. Delivering some of the routine cognitive subject matter can be done very effectively and efficiently by artificial intelligence technologies. What we cannot do is deliver exactly the 21st-century skills that we need learners to have: socio-emotional intelligence, negotiation and so on. Who is going to train the teaching workforce so that they know how to take best advantage of the sort of AI technologies that can be for them the saviour of the situation, in terms of helping them make sure the future workforce is skilled to the appropriate level? I have not said anything about the current workforce and how we also need to deal with that issue.

**Derek Thomas:** I believe that covers my questions.

**Q100 Carol Monaghan:** My question is to Mr Mason, mostly. The RAS special interest group was due to review the RAS 2020 strategy. Is that right?

***Paul Mason***: It has published a strategy document called RAS 2020.

**Q101 Carol Monaghan:** I believe that was due to be reviewed at the end of 2015.

***Paul Mason***: Pass. I don't know.

**Q102 Carol Monaghan:** You are unable to give me any information about that.

***Paul Mason***: I can follow it up for you, but I am not quite sure of the timing of it. Quite a lot of the recommendations are not yet enacted or in place, so our position is that it is probably still quite a valid document.

**Q103 Carol Monaghan:** As far as you know, there is no talk of a refreshed strategy or any adjustment of the 2020 strategy on the table just now.

***Paul Mason***: That is a strategy published by the Knowledge Transfer Network special interest group. It is not an official Government strategy. It would be worth looking at a more overarching strategy document that might bring about strong Government action in areas like regulation, standards or procurement, or the other levers Government have to try to help industries to grow. Probably some stronger detail could come out in terms of particular technologies or sectors to invest in, because that strategy speaks generally to things like skills or being challenge-led, which Rose mentioned earlier, with which we would completely concur, or the importance of testing assets and facilities. I think that is the position.

**Q104 Carol Monaghan:** If you are able to get back to the Committee on that, it would be helpful. Innovate UK has never run a dedicated funding programme for RAS technology. You can correct me if that is wrong. Would that not be considered a priority within Innovate UK funding?

***Paul Mason***: In the area of robotics and autonomous systems we run a number of programmes. Over a five-year period, we ran a programme called ASTRAEA, which was in autonomous aviation vehicles. We have been running a competition, recently announced, with CCAV in autonomous vehicles for about £20 million or so. We have run a couple of smaller programmes, which is what we have been able to afford, in analysing and using the big data that underpin some of these technologies. We have made small-level investments and have a current commitment of about £20 million to £30 million, but it is quite small compared with what some of the major robotics companies globally are investing in R and D.[1] I think ABB put €1.15 billion into R and D in 2012. That is not all robotics, but it makes £20 million to £30 million for a whole country look relatively small. We, Innovate UK, have to take account of all parts of the economy, including this—healthcare, infrastructure systems, biosciences and all sorts of different areas. Basically, we have done what we think we can afford to do in this space.

**Q105 Carol Monaghan:** I am looking at the figures. They say that RAS technologies will have an impact on global markets of between $1.9 trillion and $6.4 trillion per annum, and Innovate UK has said that a sensible share of that would be 10%. Surely, to get that share we need a clearly defined funding line that people can tap into.

***Paul Mason***: The $2 trillion to $5 trillion is the 'economic' impact, not the market. There is a big difference. We would be talking about the market opportunity. All robotics and automated systems would probably be $20 billion to $50 billion by 2020, depending on whose market research report you read. I am not sure we can say that we would get 10% of the $20 billion in five years' time, but that aspiration to grow billion-pound industries is pretty important. The UK could invest much more than it does and it would get a return. To give a couple of examples of investments we have made, in 2009-10 we invested £15,000 in two artificial intelligence PhD graduates at the University of Cambridge.[2] They were founding SwiftKey at the time and they have just sold out to Microsoft for

[1] The witness later clarified that the precise figure depends upon exactly where you draw the boundaries of what constitutes a "RAS" project, but it is in this range.
[2] The witness later clarified that "We did invest £15,000, but then made a second investment of £50,000, making £65,000 in total."

$250 million. We put some money into a Smart Award with a company called Magic Pony, which has just been acquired by Twitter for an undisclosed sum rumoured to be about $150 million. They are quite big returns for people exiting, and that is specifically an AI space. Service robotics, difficult environments and mobile production are areas where it is not too late for the UK to gain market share, but we have to invest in order to do it. We would recommend a Catapult and a broader investment programme to help companies develop products and services they could sell.

**Q106 Carol Monaghan:** Mr Knowles-Cutler, you look as though you want to come in.

***Angus Knowles-Cutler***: No. I will leave it to the experts.

**Q107 Carol Monaghan:** Philip Nelson from RCUK talked to us. He said he had had discussions with Innovate UK regarding the establishment of an RAS institute. Where are we in terms of progress with that? How do you envisage such an institute facilitating the growth of emerging RAS industries?

***Paul Mason***: We think three things would be required: first, investment in a catapult; secondly, a distributed innovation programme; and, thirdly, investment in the science base. EPSRC already has some centres of excellence in RAS. Their take is that bringing those together in a more robust and stronger way with a long-term commitment in an institute would be required for the next generation of science and discovery in that space. We trust their judgment. It sounds like a very credible activity. It is not the same as an innovation programme that takes existing knowledge and tries to turn it into money, effectively—products and services that people can sell. The institute could work very closely with that, but it would probably not do the job of helping companies to grow very directly. It is a second sequence of events, but we work very closely with EPSRC on the scope and where the exciting areas are. Our thinking and theirs is very aligned on the opportunity areas and the types of technology that need to be nurtured.

**Q108 Carol Monaghan:** Have you talked about the timescale for such an institute to be set up?

***Paul Mason***: My understanding is that there is no funding either to us or to them to do that. Obviously, that is not something in our gift. Investment by Government is a ministerial decision, not an Innovate UK or EPSRC decision. We will continue to make the case and, as things are affordable, we hope to see them coming forward, but public finances are public finances.

***Professor Luckin***: I understand the importance of robotics and autonomous systems, but I raise two issues about the proposal. The first is that it is once again focusing on the technology and not on the problems it could solve. It would be good if it could be more challenge-focused. What are the problems we are facing, and how might robotics and autonomous systems solve those? Secondly, it does not take into account the broader picture. We are acknowledging that we have a problem. We have a workforce that does not have the right education and skills, yet nothing in that would help that situation at all. There would be no investment in robotics and autonomous systems.

***Paul Mason***: That is certainly not the case in respect of Catapults. The whole point of a Catapult is to grow a critical mass of equipment, skills and expertise under one roof. It would not do large-scale teaching of 10 million full-time learners in the UK, but in terms of the skills required to develop and deploy technologies in an industrial setting it would do quite a lot in that space. In terms of developing new products and services there is quite a lot of learning by doing. You try things and see if they work; you test them out in real environments, and they would be challenge-led. That is the thinking around our Catapult.

***Professor Luckin***: I acknowledge that and I look forward to seeing it, but in the past that has not been the case. I am talking about not just skills training but the research and innovation needed to develop artificial intelligence systems that can be used in our education systems to train the teaching workforce and help the blending of human and automated education within and beyond the classroom. That is where the investment is essential and where there is no investment at the moment.

***Paul Mason***: That is not within our scope, but it is a fair comment.

**Q109 Carol Monaghan:** Do you not see that sort of thing in the BBC micro:bit?

***Professor Luckin***: There is no artificial intelligence in a BBC micro:bit. There could be. We could develop a system that would blend the micro:bit with physical computing, and we are already developing a system that uses artificial intelligence with physical computing. Mr Mason mentioned EPSRC, and that is where the problem lies. The Engineering and Physical Sciences Research Council funds technology development and the Economic and Social Research Council funds education, but neither does the crossover. If they do, they say, "Well, it's got to focus on technology, but you can apply it in education if you want," or, "We'll look at education with technology, but we're not going to pay for the development of the technology." The truth of the matter is that it has to be truly interdisciplinary. The learning sciences have to inform the design of the technology; otherwise, it is not going to work.

**Q110 Dr Mathias:** This may be a simple question. Dr Osborne, if artificial intelligence can plan, learn and adapt based on new information, how can we verify that it is still working as originally intended? It is the Hollywood question.

***Dr Osborne***: I do not think it is a Hollywood question. It is a critical question for the field of machine learning at present, and is at the heart of a lot of research that is currently under way. Part of the problem is that machine learning up until now has been focused on achieving slightly better performance on well-defined problems, such as the classification of images or the translation of text. The interpretation of the algorithms that are produced to achieve those goals has been left as a secondary goal. We are not where we would want to be in ensuring that the algorithms we deliver are completely verifiable and validated, but we are making a lot of progress, and there are approaches that will enable that kind of verification.

**Q111 Dr Mathias:** We are not there yet, but you think we can get verification.

***Dr Osborne***: We can get there. There are some classes of techniques that are better at enabling that kind of verification than others. I am a personal critic of deep learning, which is one of the techniques that has achieved most success in recent years, exactly because I do not think it is as successful in enabling that kind of verification and validation.

***Professor Luckin***: I would say three things in addition to what Michael said. First, there are other types of AI systems in addition to machine learning, and some of those are absolutely transparent. Part of their advantage is that they are transparent, in that they present what we call open learner models. They show teachers and learners exactly what has gone on, how the system has made its decisions, and why, as a support to the learning process. They are transparent systems. Secondly, public engagement is essential in order to get the public to understand what AI is and what it can do, so that they recognise when things are not behaving in the way they should, and when they are being sold a load of old baloney as well.

**Q112 Dr Mathias:** How on earth would you get that to people who are not in the field?

***Professor Luckin***: There are certainly ways of engaging with them. We do it all the time. I always use participatory design methods. We always engage with our stakeholders. Sometimes it is really hard, but it is essential if you are really going to develop the field.

The third thing is about educating people in schools about what artificial intelligence is, so that they have a much better engagement with the technologies when they meet them and much better ability to do some self-verification to complement the kind of verification Michael is talking about, and the sort of transparent systems we could also use.

**Q113 Dr Mathias:** You are talking about the people verifying it rather than another system.

***Professor Luckin***: It has to be a combination of both, in the same way as we are talking about blending automated systems and human systems in the workforce. In a way, that verification has to be a blending. Humans need some kind of understanding; otherwise, how do they know whether or not the verification they are being offered is genuine and how do they still trust the system?

**Q114 Dr Mathias:** I think we are expecting a lot of all the users, and that is quite scary. Do you want to add anything?

***Angus Knowles-Cutler***: Perhaps I could make a broader point to pick up some of the earlier questioning. When I talk to my colleagues at Deloitte who are at the heart of how technology is affecting business now and in the future, in their digital and technology strategies they are looking—perhaps it is an elusive quest—for purple people. What they mean is that, if blue is knowledge of AI technology and systems, and red is knowledge of business, people and markets, they want a mix of the two. When we talk to a lot of clients right across UK businesses—I am sure it is the same for public sector organisations—that is the fundamental skill set being sought. People who are just technologists or AI specialists and do not understand current and future applications are not what is needed,

nor do they need business people or creative people who do not understand the application of technology. It is a blend for the future that we are looking for.

***Paul Mason***: I agree with everything that has been said, but to add to something Angus said, quite often we find that those skills are not imbued in the same individuals, certainly not at the cutting edge of science or business. Bringing companies and universities together to work on collaborative projects is a good way of ensuring output overall. They learn from one another as they do the work. The vast majority of our investment goes into that kind of programme, and it is very fruitful.

**Q115 Dr Mathias:** Is it possible to build intelligence aligned with human values?

***Dr Osborne***: The difficulty there is understanding what human values are.

**Jim Dowd:** Don't ask them.

**Dr Mathias:** People are vulnerable today.

***Dr Osborne***: In machine learning, progress has been made by carving off very crisply defined problems for which there is a clear sense of what is right and what is wrong.

**Q116 Dr Mathias:** You evaluate that.

***Dr Osborne***: That is right, and that is how we have advanced such as we have. It is much more difficult to imagine an algorithm that can understand all the things that a human wants to achieve and make progress towards that.

**Q117 Dr Mathias:** You cannot have Tibetan AI.

***Dr Osborne***: Not yet at least.

**Chair:** I thank all of you for the evidence you have given today. I think that all of us now aspire to Tibetan AI. We are very grateful for the time you have taken. This is a complex subject, and we are going to do our best to produce a robust and helpful report that will not just inform the Modern Transport Bill, which is coming forward to produce regulation on driverless cars, but will also provide an evidence base for this House, which is obviously grappling with a lot of the issues we are talking about today. It is important that we have as much information as we can get at this stage to do that. Thank you for your time. I hope that, if we have questions as we move forward, you will be kind and send us answers.

## Examination of Witnesses

Witnesses: **Dave Coplin**, Chief Envisioning Officer, Microsoft, and **Mike Wilson,** General Industry Sales Manager, ABB Robotics, gave evidence.

**Q118 Chair:** Thank you for coming today. I saw that you were listening intently throughout the previous discussion, so I will jump right in with the impact of AI and robotics on the workforce. You heard some different views from the previous panel, and we heard some other evidence in our last session. Where do you stand on this? Do you think it will eliminate jobs? Will it create new forms of employment that will compensate? Mike Wilson, what do you think we need to do to prepare society to respond effectively?

***Mike Wilson***: I look at this from a manufacturing industry perspective, because that is my background. To date, robots and automation have taken over the dirty, demanding and dangerous jobs, generally the ones people do not want to do, and they have added value to the manufacturing workplace. I think we will start to see increased use of robots throughout manufacturing. It has already been said that sectors like the UK automotive industry are a lot more productive than they used to be because they have employed a lot of automation. Likewise, they have also employed thousands and thousands more people who otherwise would not have been employed.

Earlier, it was touched on that generally across UK manufacturing we do not use as many robots as our major industrial competitors. We use about 31 robots per 10,000 employees in manufacturing, whereas Germany is up to about 160, so we are well behind. The challenge we face is that we need to accelerate that; we need to put more robots into our manufacturing facilities so that our workers are given the tools to be more productive. Ultimately, it is all about making workers work smarter rather than work harder. We should push for greater use of automation; ultimately, that will make businesses more competitive and therefore they will employ more people. We will get more jobs out of it, but we'll get better jobs.

***Dave Coplin***: I worry about the line of questioning, not just today, but across society. The conversation is all about humans versus machines, and that sets us off on the wrong foot. Technology has never been about humans versus machines. The story, certainly from our perspective, in the personal computer revolution is about how we augment humanity and how we enable human beings to rise up and achieve more than they could on their own. Rather than talk about automation and labour saving, we should talk about creativity. We should talk about how we can help people to achieve more than they could on their own—your point, Mike. If you start in that place, the way forward starts to emerge as a very different path. It is not about massive disruption in the workforce but about how we create new markets and new opportunities. Then we move very quickly away from the regulation of robots and AI to how we create the pipeline of skills to feed that. That is where we have to go with this stuff.

**Q119 Chair:** To follow exactly from that, and from Professor Luckin's comment about the issues in the UK, which we have just reported on in our digital skills report, we have a STEM skills crisis in the UK, but we also have issues with basic skill levels. How do we address that? What kinds of skills do we need to bring into the workplace to deal with blended tasks, and with automated tasks not just for basic uses but for the creative and innovative uses we would all hope to see?

***Dave Coplin***: For me, we have to try to stop separating them. We talk a lot about STEM, which is really important—don't get me wrong—but, without art and creativity, innovation is dead. We could have a bunch of scientists, which would be brilliant, but their

ability to be creative in the future world of work is the thing that makes them successful. We have done a great job in the UK. We should be proud about how we put computer science in the curriculum, but we have to take the next step, because we should be looking to ensure that every member of our society is equipped to be able to use technology to deliver a positive outcome in their lives, regardless of whether or not they want to be a computer scientist.

How do we do that? We have to inspire people. It goes back a bit to the conversation about the rhetoric and the imagery we have for computers and robots today, which is a dystopian view. Where are the good role models and the things where we can say, "Look how creative this individual has been as a combination of their human skills—empathy and creativity—and their ability to manipulate the technology to deliver a great outcome"? That is the narrative we need in education, and that is why we need to inspire kids. We do not need to frighten them off with a bunch of science; we need to show them how creative they can be and how it is a blended world. If we can get that right, we will solve all sorts of things, like the gender issues we face in attracting people into science and technology; we will extend resolution of some of the problems we face because we will have a more creative approach. But I fear that, as long as we talk about STEM in exclusivity of the other skills that will be required, we limit our ability to be successful.

**Q120 Chair:** Mike Wilson, this is an argument for the STEAM agenda—STEM plus the arts—which as a musician I can only endorse. Where do you stand on this? As a manufacturer, in your industry how much creativity do you want to see in your workforce and how would it help you to deliver your agenda?

***Mike Wilson***: Creativity is very important, particularly when you are looking at the application of automation, because you need people to come up with the right solutions and new ideas about how to apply this kind of technology. Therefore, creativity needs to be very much part of the agenda.

In terms of exciting and stimulating schoolchildren, I do not know whether you have come across the Lego robot schools competition. A sizeable number of schools are involved and young children are getting very excited about technology, and it brings in creativity at the same time. The challenge we face is not so much that that is not successful but what happens after that, because it drops off and there is not a next stage in the education process to continue that excitement and maintain the interest of those children in robotics, automation and artificial intelligence. We need to look at further education, in the universities and technical colleges, as to how we encourage the children who got excited about Lego robots to move into engineering and continue their creativity and education in that role.

**Chair:** We could talk about this all day, but I have to move to Committee members.

**Q121 Graham Stringer:** Does the robotics and AI sector have difficulty accessing investment finance?

***Mike Wilson***: I am not sure I can answer that because it is a little outside my area of expertise, but my conversation with banks at present is that they have lots of money they

are willing to lend people. If people have the right business case, there is a reasonable chance that the investment would be there.

**Q122 Graham Stringer:** That closes down a whole line of questions. We have heard different views about regulations across robotics and AI. Some people think there should be very light regulation; some believe more regulation will help investment because it will give security. What aspects of robotics and AI might benefit from regulation?

***Mike Wilson***: From a manufacturing perspective, the normal standards have to be applied in terms of the use of robots and automation. A new generation of collaborative robots is coming through now and there is a certain amount of confusion about how they can be applied, and what is considered safe to be used. That will continue. The pace of development continues, and the standards and the legal frameworks around them are not keeping up with the development of the technology. That certainly needs to be addressed to ensure that people have a clear picture of where the standards are going to be.

**Q123 Graham Stringer:** Does that mean you think the regulatory framework should be product by product rather than across the industry?

***Mike Wilson***: Not product by product, but possibly technology by technology, because there are different issues depending on specific things within certain technologies.

**Q124 Graham Stringer:** Who is developing the regulations, and who should develop them?

***Dave Coplin***: I think you have to do it in combination with a number of different parties. We have to do it with people like yourselves; we have to work with a number of NGOs, in particular the privacy groups. You ask about where we need some regulation. If AI is going to work successfully for us as a society, we need some intelligent privacy and we need to figure out how to do that. We need to hold the people who write the algorithms accountable for them, so how do we make that happen? We also need to make sure that we are able to minimise the biases that will be imbued in the algorithms. In AI every time an algorithm is written, embedded within it will be all the biases that exist in the humans who created it. How do we spot that? How do we become a diverse society to prevent that? These kinds of conversations will help us resolve it. It is for the tech industry, the Government, NGOs and the people who will ultimately consume the services. We have to find a way of convening those four parties together to drive forward that conversation.

**Q125 Jim Dowd:** My question is principally for Mr Coplin. Is it right that your title is chief envisioning officer?

***Dave Coplin***: It is, sir. I am immensely proud of that title. I created it myself.

**Q126 Jim Dowd:** I can understand you would need to be. Are there any junior envisioning officers? How many do you have?

***Dave Coplin***: I hope so. I am raising a 10 year-old one right now.

**Q127 Jim Dowd:** What is one?

***Dave Coplin***: The reason I created that position is that I was fed up with my industry talking about the technology and ignoring the human beings who use it, so I created a position inside my organisation that works more broadly across societies to think about and have a conversation about what we want human beings to be able to do in the future. If we know what we want humans to do, we can be a lot more mindful about the kind of technology and policy they will require to support that. I figured that, if I had a job that was essentially about the future of humanity, I deserved a job title that had a certain degree of pomposity about it, hence the choice of job title.

**Q128 Jim Dowd:** There is no lack of scope.

***Dave Coplin***: Absolutely.

**Q129 Jim Dowd:** Can I come to Microsoft's experience with Tay? What was the point of it?

***Dave Coplin***: The point was twofold. There was a technological side. Could we create a technology—an AI bot—that would be able to converse with a human being? The second part of the experiment was a social one. What would a society do with a bot that you could converse with? Unfortunately, to our regret, we learned exactly what a specific society could do with a bot you could talk to. You could fool it into saying the most terrible, awful things that Microsoft in a million years would never want to be associated with.

We see the sad thing about that example if we look at the application of a similar technology in a different society. We have another chatbot very similar to Tay called Xiaoice in operation in China. It has had over 40 million conversations with Chinese teenagers about celebrities and gossip, and not once has it been bullied into doing things we did not want to do. The lesson for us with Tay was, first, what can we do with the technology, and, secondly, what are the implications for a society of this kind of technology? Look at one of the great areas of disruption that is about to come. The call centre industry is going to be transformed by the prevalence of artificial agents. As a customer of one of those call centres, do I need to know that I am not talking to a human being? It is an open question. These are the things we are going to be challenged with, and Tay was one example of flushing out some of those issues.

**Q130 Jim Dowd:** You said it was a social experiment. It was not a commercial activity.

***Dave Coplin***: No, not at all.

**Q131 Jim Dowd:** You just thought it was something that would be worth doing.

***Dave Coplin***: The regret we as a company have about it is that we should have had the foresight that people would want to be naughty with it, and we didn't. We were a bit too naive about it, so we deeply regret what happened.

**Q132 Jim Dowd:** It only lasted 16 hours, didn't it?

***Dave Coplin***: Yes. Absolutely, and how terrible is that—that the first thing we would want to do is to game the system? As a pure technologist, I am slightly depressed by humanity.

**Q133 Jim Dowd:** Aren't we all?

***Dave Coplin***: Yes.

**Q134 Jim Dowd:** A large part of the problem with it, allegedly, was the "repeat after me" function. Was that inbuilt or acquired?

***Dave Coplin***: It is a blend of both things. The way artificial intelligence works is that it learns based on repetition. In the previous session we talked about how big the corpus is—the sample of data used to train the data. All that was happening was that people were training Tay to say the wrong things, in the same way as a teenager might train their kid brother to say the wrong words at a dinner table. Tay was kind of doing what it was supposed to, but we did not build in the controls to say, "Do you know what? There are some things you don't say, not even at the dinner table."

**Q135 Jim Dowd:** Microsoft employs some of the smartest people in the world. It didn't become one of the biggest global corporations without them. Did none of them spot that that would be likely to happen with the kind of people who inhabit cyberspace?

***Dave Coplin***: On one level it is a great question. On another level, you have to understand that, if you are a computer scientist creating the ability for a machine to be able to converse with another human being, you are naively wrapped up in the amazing technological progress and potential of what you have created. You don't for a minute think that there are people out there who want to pervert it and turn it into something destructive.

**Q136 Jim Dowd:** Don't you?

***Dave Coplin***: Not if you are a good computer scientist. Remember, this is a research project rather than a commercial one. We are creating it to create the aspiration to see what will happen in a society. We put it out there to see what happens. Therefore, it would not go through the same rigour as a product because it is a research project; it is asking questions.

**Q137 Jim Dowd:** Is what you have just described what is meant by the "critical oversight" mentioned in the explanation by the company?

***Dave Coplin***: Absolutely. With all these projects, we sometimes need to take a step back. We are the technologists. Just because we can do something does not necessarily mean we should. If I had my time again with Tay, I would do exactly the same thing, except that I

would build in some filters to prevent it from saying certain things. You could still get it to appear foolish and say stupid things, but not the hurtful, terrible things it said.

**Q138 Jim Dowd:** Are there any broader lessons? Is it your intention to bring this back at some stage?

***Dave Coplin***: Absolutely.

**Q139 Jim Dowd:** Therefore, there are lessons you have learned.

***Dave Coplin***: Absolutely.

**Q140 Jim Dowd:** What is the point of bringing it back?

***Dave Coplin***: The point of bringing it back is that we believe as a company that many of the conversations we will have in the future will be had by artificial agents. We have a principle called conversations as platform, which enables the system—technology—to do things on your behalf. To take a trivial example, if today you want to book a holiday yourself you have to search on destinations, airlines, hotels and what you are going to do when you get there. In a world of agents, we know who you are and we know the things you like to do. Give us the destination and on your behalf we will negotiate things you might like to do and give you a human outcome. That is where we think the future is going. That is why we need bots, artificial agents, that can converse with human beings, that cannot be tripped up and that deliver value to the individuals we are looking out for. We think this progression is fundamental to where we are going. The lessons we learned from Tay are fundamental to enable us to do the right things in the future.

**Q141 Jim Dowd:** But this is all processing, not understanding.

***Dave Coplin***: Absolutely. For me, even the term "artificial intelligence" is bogus at this point. There is no intelligence; it is simply processing against an established pattern.

**Chair:** Victoria, will you process a new pattern?

**Q142 Victoria Borwick:** It is very interesting. Listening to you, I am interested in how we are going to drive forward what we are doing in the UK. Do we have a lower level of automation than most of our industrial competitors? What is holding back automation in the UK?

***Mike Wilson***: We certainly have a lower level of automation than most of our industrial competitors, and that has been demonstrated by industry statistics. There are probably a number of reasons behind that, but one of the fundamental ones is lack of awareness within general manufacturing. The automotive industry knows what it is doing and what its competitors are doing. If we step outside automotive, a lot of manufacturing businesses are not aware of how far behind they are. They are fearful of taking on the technology; they fear that they lack the skills and they think it is too expensive. We need to provide

that awareness and education to ensure that our manufacturing sector at least questions whether it should be using automation rather than employing people to do those jobs. In parallel, we need to put something in place to prime the skills channel, as it were, and make sure that we have engineers coming out of colleges with the skills to implement those kinds of systems. At the moment, we really struggle.

**Q143 Victoria Borwick:** My personal view is that we do not start teaching those sorts of subjects early enough, but that is a slightly different question from what we are asking today. You state in your written evidence that the solution is automation, so what is your vision of how we would make that a reality?

***Mike Wilson***: In terms of manufacturing, most of the technology is available today. It is not about developing new things and new approaches to solving problems; it is about utilising existing technology. Some of the initiatives, such as the Manufacturing Technology Centre and what they are doing, are helpful because they provide access to businesses, but generally I find that SMEs do not approach the MTC; they think it is for big companies rather than for smaller businesses.

**Q144 Victoria Borwick:** Do you think that is because of the initial cost?

***Mike Wilson***: It is partly the initial cost and partly a perception. Because it is a very impressive building with a lot of expensive equipment in it, they feel it is not appropriate for them and they think it will be expensive. The first step is principally about awareness, and Government and politicians can potentially help by highlighting the companies in their constituencies, in their areas, that have used automation and been successful, and use that as a vehicle to encourage other manufacturing businesses to look at that kind of technology.

**Q145 Victoria Borwick:** You also commented that the Robotics and Autonomous Systems special interest group has had little impact on driving forward this agenda. Do you want to make any comment as to how we could improve that?

***Mike Wilson***: From a manufacturing perspective, it did not and does not have any particular representation from manufacturing and it is driven largely by academic interests looking at the more advanced robotic technologies.

**Q146 Victoria Borwick:** Do you think there should be practical users, or end users, on the group as well?

***Mike Wilson***: Exactly. It needs a significant manufacturing and industrial robotics content.

**Q147 Victoria Borwick:** Do you mean an industrial manufacturing user or somebody who manufactures the robotics?

***Mike Wilson***: Manufacturing industry, but also some of the automation suppliers, because they can add value to those conversations as well. They know what is available and what is being done elsewhere in the world.

**Q148 Victoria Borwick:** I wanted to make sure I understood, because there are obviously those who make the robotics and those who use them, and they are not necessarily the same.

***Mike Wilson***: It needs to be both: those who use it and have had good experience, and those who supply it.

**Q149 Victoria Borwick:** That is helpful. A year ago the Government agreed to establish the RAS leadership council, but we have heard that the idea has since been dropped. Do you have a view about that council? Is it necessary or unnecessary, or what do you think would be better?

***Mike Wilson***: It is appropriate to undertake that activity because we need leadership, but to go back to what I said about the strategy, a significant element of that needs to be focused on manufacturing, not just on the advanced robotics technologies. If it is to be put forward—I think it is worth doing—maybe it needs to be framed in a slightly different way.

**Q150 Chair:** Dave Coplin, what kind of leadership do you think we need in this area?

***Dave Coplin***: Stewardship is the kind of leadership we need in this area. We need help in opening some doors. The real challenge faced by Government is that we need to show people the potential of what all this technology enables, and I do not think we do enough of that. We need to say what is on the table. We then need to create a world that is frictionless, or as friction-free as possible, where people can embrace that technology and create new businesses and opportunities around it, but we need a safety net from you to help us with things like accountability for the algorithms, privacy and the issue around bias, so that people can be held to account in how we build those tools. There is a very delicate balance to be struck, but it has to start with something that is encouraging.

I believe as an individual and we believe as an organisation that artificial intelligence in particular is probably the most important technology on the planet that anybody is working on today, not just us at Microsoft, but Google, Facebook and everybody. We need to create a world that understands it, can embrace it and get the most out of it while minimising the risks it brings to our society. That is where we need the help. In that sense, true governance is to create the opportunity for that and help us to minimise the risks.

**Q151 Chris Green:** Mr Wilson, how do we make sure that an autonomous system is safe, and that unwanted behaviours do not emerge as it operates in the real world?

***Mike Wilson***: The technology being employed today is not autonomous, in that it is not making its own decisions. It is outside my area, but in things like nuclear plants and so on there is a significant level of automation, and fail-safe mechanisms are built into that. From a manufacturing perspective, all robots are programmed and controlled by people,

and the systems are set up to ensure they are safe to operate with and, provided people adhere to the safety standards being applied, everything is fine.

**Q152 Chris Green:** There are fail-safe systems within incredibly sensitive industries such as the nuclear industry, but if we look at autonomous cars, there is a completely different scenario. You always have the risk versus safety idea, but with autonomous cars you might have the safety versus speed idea. The stats from the car will not be just how fast you can go from nought to 60 but how fast you can manoeuvre round the city of London and get from A to B. Can you get there faster in one make rather than another? That may lead to some manufacturers optimising speed at the cost of risk. Large organisations might be quite protective of their software, algorithms and everything else. We have seen with what happened with Volkswagen and diesel that they can get up to some mischief. There are all kinds of problems as to what may happen when autonomous vehicles are unleashed, so who should be responsible for the verification and validation of controls?

***Mike Wilson***: Driverless cars are outside my area of expertise, but it is important that appropriate regulations and bodies are put in place to ensure that in the use of autonomous vehicles and automation, where it interacts with the general public in that way, every step is taken to make sure that it is safe and the risk is negligible.

***Dave Coplin***: All these technologies, not just autonomous vehicles or robotics but artificial intelligence, will bring with them some significant ethical challenges to our society. That is where we really need some help. We have seen this movie many times before. I used to work for Bing, which is Microsoft's search engine. For me, search is one of the big areas where we have seen machine learning in particular come into our society and present challenges. I give you one that we had a few years ago. You are familiar with the auto-complete function of search engines. As you are typing in your query, the search engine does some maths and says, "Based on the characters you have just typed in and the billions of other queries that have also been typed in today, the statistical probability of your completing a sentence like this is 99%, so I will suggest that as the next option." A few years ago we were presented with a moral dilemma by people typing in, "How do I commit suicide painlessly?" As a search engine, morally what do we do with that question? How do we answer it? Do we serve the answer, or not? It is not illegal to answer that question, but is it problematic? Yes, it is. We have to have a conversation as a society about what we want that answer to be. We have seen it more recently, in the context of terrorism, with the question, "How do I join ISIS?" Is it wrong that I should type that in? Is it wrong to answer it? I do not know; I cannot answer that from a technology company perspective.

**Q153 Chris Green:** Facebook has been challenged with that. Certain groups that you would not like to encourage have been closed down, but others have not.

***Dave Coplin***: Exactly. This is where we need to come together. These are the questions that the technology will bring to us. The classic theoretical example with autonomous vehicles is the no-win scenario. If I have to crash into one of two groups of human beings, how does the algorithm decide? These are the really difficult legal, ethical and moral discussions we have to start thinking about? We move away from the "us versus them" mentality. What do we do with that? It brings us back to speed versus safety. The

Government will have to set a baseline of law that says these are the things that are legal for you to do.

**Q154 Chris Green:** Should it be just Government? In your book "The Rise of the Humans" one quote that might resonate quite a lot at the moment is "take back control of the potential that technology offers our society." What kind of involvement should people have more widely?

***Dave Coplin***: I firmly believe that the conversation needs to be with Government, the tech industry, NGOs, where relevant, and the people we care about most, namely those we are trying to deliver services to, but the challenge we have—again, I come back to this—is that the starting point has to be a naively optimistic view: "Let me tell you about the great things that are going to happen, it's going to be unicorns and rainbows." Understand that it is naive, but that is what is on the table. When we understand that, then and only then, can we work backwards and say, "How do we manage the risk of this being used badly?" If we do not do that, I worry that we end up in a downward spiral of, "Oh, you can't trust technology; you can't use it," and the end state of that game is that we turn off the technology. As a technologist, that is game over for me—we have missed the point—so we have to start from a different perspective.

**Q155 Chris Green:** We need more people envisioning for the future.

***Dave Coplin***: I think so. My point is not that you need more people with ponytails and beards, but that you need to inspire lots of people to think very differently. For me, it is about creativity. We should not be having a conversation with kids about robotics and AI but about what they can make with it, what they could do with it. When you do that and you see their lights go on, they do not care whether they are using robots or AI; they want to engage with it. If we can do that, we create a positive influence and a positive aspiration about how we can lift humanity further.

**Q156 Chair:** Mr Coplin, I would like to reassure you that we make no judgment of the fashion choices of any of our witnesses. I would like to take you back to some of your comments about ethics. Obviously, this is a central strand in the questions we are asking our witnesses. You raised the scenario of the driverless car that has to make a decision in a no-win situation, but surely a human driver in the same situation would have to make the same decision. The only difference is that it is perfectly clear who is responsible, and there are not questions about whether the system has in some way malfunctioned. It is a question of responsibility rather than a question of decision making alone. It is not clear to me at this point whether the proposal is that tracing responsibility should go back to the programmer or to the person in charge of the AI, and if the skills level of the person in charge is not high enough to understand the programme they are using, how that responsibility could be given to the person in charge of the robot, driverless car, search engine or whatever other system they were using to make an impossible decision.

***Dave Coplin***: There are some fundamental building blocks that will help us with that problem. They will not themselves ultimately solve it, but they will give us the tools to be able to solve it. The building blocks are things like the way in which we create the

algorithms. They must be transparent. I must be able to see the pattern or rules that have been used to create the outcome. As a human I need to be able to inspect that, as much as the algorithms need to understand what the humans may choose to do with that information.

We need a level of accountability for the algorithms. The people making the algorithm and the AI need to be held accountable for the outcome. We also need to do lots of work around bias. Bias is really important. A relatively new level of understanding for us in the world of AI is how we as humans imbue human bias in artificial intelligence. If you will humour me, I will give a trivial example—something that occurred to me while I was in Canada a few weeks ago. If you were to ask somebody to write an algorithm to make tea, you should give real thought as to who you are asking to do it. If you asked a British person to write that algorithm, there would be a warm pot, tea bags would be scalded by boiling water, the pot would be left to stew for a few minutes and then the tea would be poured into a warm cup. If you asked a North American to write the algorithm for making tea, there would be a cup of lukewarm water with the tea bag and milk already in it, and it would be served up. Both of those are valid cups of tea. *[Interruption.]* I share your view, but as a consumer of that cup of tea, being transparent about who wrote the algorithm is really important to you, because it affects the outcome. If you will forgive that really trivial example—

**Chair:** I fear you have spoken controversial words.

***Dave Coplin***: It shows you that we now have to be mindful of the philosophies, morals and ethics of the organisations that are creating the algorithms that increasingly we rely on every day. That is where this conversation needs to go. We have to look into that and establish those things. As a company, we know that the answer to many of these problems lies in things like diversity and the right kind of internal regulation and processes, so that we do not end up in a situation where we have unintended consequences, or, if we do, it is transparent enough that we can look back to see where the problems occurred. This is the learning we must have as a technology industry working with you in Government and with policy makers and the public at large as to the outcome of these things.

**Q157 Chair:** In order for that sort of accountability and transparency to be effective, the verification and validation systems would have to be able to respond to active learning systems, and the people who are analysing those validation and verification systems would need to be able to understand the outputs of those systems. We have discussed how far advanced we are with AI and robotics. How far advanced are we with the verification and validation systems? We have not had positive comments so far on this particular area of research. Where is Microsoft on this?

***Dave Coplin***: It is not an area I am that familiar with in terms of the specifics.

**Q158 Chair:** You have given us the trivial example of biased tea-making, but bias in AI weapons systems might not have such a trivial outcome. Are there any areas of AI or robotics that you think should be off-limits?

***Dave Coplin***: That is a wonderfully moral question. There are some baselines. We should not be doing anything that is illegal. As a society, we have a duty of care to consider the implications of how the technology might be used to destroy life rather than to give life. My frustration with where we are with this conversation, not in this room but more broadly as a society, is that I think this technology will help our society with some of the greatest challenges it faces, from diseases like cancer right through to how we feed poorer nations. The technology will help us with all those things, and that is where we should be putting all the focus, rather than on the things that will destroy life or take life away. We need regulation to prevent that, as with any powerful technology. I do not see this as being any different.

**Q159 Chair:** Mike Wilson, before we close the session do you have any key points on the direction we are going in with ethics and the frameworks we need to put in place in the light of the comments we have just heard?

***Mike Wilson***: I agree with Dave's comments. I should stress that we are probably a long way from the reality of these kinds of technologies getting to that point, but we should be thinking about it now so that we have had the thought process before we get to that point.

**Chair:** I thank both of you. It has been an extremely thought-provoking session and extremely helpful for us in developing our thinking for our report. We may well have questions to follow up with, and I hope you will be kind and write back to us. Thank you very much.