Prof. Nick Jennings CB

Interviewed by

Elisabetta Mori

12th March 2020

At

Imperial College, London

Welcome to the Archives of Information Technology. It's the 12th of March 2020, and we are in London at Imperial College. I am Elisabetta Mori, an interviewer with Archives of IT.

[00:13]

Today I'll be talking to Professor Nick Jennings. He is an internationally recognised authority in the areas of artificial intelligence, autonomous systems, cybersecurity, and agent-based computing. He is the Vice-Provost for Research and Enterprise and Professor of Artificial Intelligence at Imperial College in London. He was the UK's first Regius Professor of Computer Science, a post bestowed by the monarch to recognise exceptionally high-quality research, at the University of Southampton, and the UK Government's first Chief Scientific Adviser for National Security. With more than 80,000 citations, he is one of the world's most cited computer scientists. He was awarded a Companion of the Order of Bath in the Queen's New Year Honours List in 2016 for services to computer science and national security science. He has received a number of international awards for his research including the Computers and Thought Award, and the ACM Autonomous Agents Research Award. He is a Fellow of the Association for the Advancement of Artificial Intelligence, a Fellow of the Institute of Electrical and Electronics Engineers, a Member of the Academia Europaea, and a Fellow of the Royal Academy of Engineering.

[01:34]

Welcome Nick, thank you for being here today. Let's start with where and when you were born.

I was born in London on the 15th of December 1966.

Would you like to describe your parents? What were their occupations?

My father was a prison officer. He is now retired. He was born in 1947. My mother also worked in the prison service, on the administrative side, and she was born in 1946.

Have you got any siblings?

I have two younger sisters, Rachel, who was born in 1974, and Gail, who was born in 1970.

Where did you grow up?

I grew up on the isle of Portland, in Dorset.

What was your family life like?

We were a very happy family. We were quite close-knit as a family. We would discuss things, and really be involved in one another's life.

Who were the important influences on you in your early life?

I think in my early life, my parents were, and my friends that I had at school. I was also very keen on sports, and so some of the people who ran our sports teams were a really important influence on me early in my life. And then, actually, some of my teachers, particularly those who helped me with some of the sciences and some of the computing.

[03:03]
What schools did you attend?

I went to Weymouth Grammar School, there was still a grammar school system in place at that time, and so I went there in 1978, and I stayed there, I did my O Levels there, and I also did my A Levels there.

What subjects did you study at secondary school, and which did you most enjoy?

I definitely enjoyed the sciences. So, I enjoyed maths, and I enjoyed computing, and I enjoyed physics and chemistry, I think those were the subjects I enjoyed most. I hated some subjects, so I really hated art, and, and subjects that were creative in that sense, [laughs] because I really struggled with them. And then at A Level I took the very science-focused route of doing maths and physics and computer science.

[03:57]
What was your first computer?

My first computer that I used was a BBC Micro, and the first computer I owned and bought was a Dragon 32.

And what are your memories of the time?

So, at that time, I was very keen on sports. I loved playing football and cricket at that time, and I think that's what I used to spend a lot of my spare time doing, and really, really enjoyed those. I always enjoyed doing things in teams, and, sort of, as part of teamwork. And then, sort of, gradually I got interested in, in computer science, and sort of, programming my own computer to, to do, to do things. I programmed some simple, simple games, and, and really trying to explore what could be done with a, with a computer at that age.

Can you make some examples of the games you programmed?

I remembered programming a computer to play pontoon, so, twenty-one, blackjack. So that was good. And then also, some more, very basic games where you were able to program some things visually. So, I particularly remember, sort of, trying to replicate some of the Space Invader games that were just coming onto the market in those days. So, sort of, being able to have an animated computer screen, which was, when I was doing this was a big deal. It's of course not anymore, but, actually, at that stage, having some graphics, and being able to, to display things on them, was a, was an amazing revelation and insight for me.

Did you have any friends doing this with you, or were you…?

Actually, [laughs] most of my friends were not the slightest bit interested in, in computer science, and I'm, I'm not entirely clear how I got interested in doing it. My friends were, wanted to, wanted to have fun, wanted to play sport, and were far less academically interested than, than I was.

[05:59]
Which colleges and universities did you attend?

So, after sixth form in Weymouth I went to the University of Exeter to do computer science. So, in 1985 to 1988.

Who and what were the important influences on you at this time?

So, this was the first time I had been away from home, and actually going to Exeter was going to a big city for me. So, Portland is a, is an amazingly small place where actually, it's quite a tight-knit community, lots of people would know lots of people, so, you know, wherever you went, you would come across people you knew. And so, actually going into a bigger place, and going into a city, and of course Exeter, on reflection, is, is not the world's biggest city, but it was from, for me as someone who was going there from, from Portland, was an amazing experience to sort of, go out and, and meet people from other parts of the, other parts of the country, and other parts of the world, and… So, Portland was very insular in, in nature, a lot of people lived there, had always lived there, that's where their parents lived, and actually meeting other folk with different views and different accents and different perspectives on life was a, was a really interesting time for me.

[07:17]
Is there any particular event that shaped you during your years at the University of Exeter?

One of the most important things was that I met my, I met my wife. So, we met as undergraduates in Exeter. My wife studied biological sciences. And we got together through some, some mutual friends. So that's a, that was an important thing that happened to me at Exeter. And also, I met many of my best friends, so, sort of, friends in my, in my peer group. In fact, a number of us still meet up. We were in some small hall, so, which was a small house, and, actually many of those people I still meet regularly. My best man was, was one of, was one of those people, and I've been involved in, in the weddings of, of the others there as well.

What's your wife's name?

My wife's name is Joanne, Joanne Jennings as we is now, Joanne Smith as she was then. And, she is also a scientist, so, she did her PhD here at Imperial College London, so she is actually a proper alumni of the college. I just, I just happen to work here, as she reminds me of. And so, she is a, she is a scientist who now works in the broad area of patents.

[08:41]
During those years, did you travel abroad?

As a student and growing up, we hardly ever, I hardly ever travelled anywhere. And so, I had a very limited [laughs] world view I think at, at that time. When I, during my PhD, I was also employed as a research assistant on a, on a European grant, so an ESPRIT-2 grant, called ARCHON, which was one of the first big projects in the area of multi-agent systems. And actually, was a, at its time was quite a big project, so it had, I think fifteen partners spread out around Europe. And actually, that was the first time I had really got into travelling and to going to different places. So actually, it was a fantastic experience for me to, to go out and to see the world much more. And now it's one of the things I, I really enjoy. So, one of the, one of my ambitions in life is to go to more countries than years old I am, and, and I am ahead, so I am, I am ahead by about ten or so countries at the moment. So, I really do enjoy seeing the world and, and exploring it. And I think that love sort of, came from those, from those early trips on my science career.

[10:05]

When did you first hear about AI?

So, I first came across artificial intelligence when I was an undergraduate. We… Exeter was an amazingly small computer science department; I think there were about 20 to 30 of us on the course. It was so small that, actually every undergraduate had their named individual pigeonhole, so that gives you a feel of the size of the

department. It really was tiny. But actually, one of the strengths of the department, and one of the things that a decent number of the staff were interested in, was artificial intelligence. So, so I think we came across, basic introduction courses to AI in the second year of undergraduates. And actually, I was then fascinated by sort of, the idea and the concept of making the machines, and making machines do smart things. And so, sort of, that's really where, where my interest was, peaked, in thinking, this, this is a really interesting topic in itself. At that time, AI was, was quite a niche thing to be doing, it was very different to its, its wide recognition that it exhibits today, but sort of, you know, then it really was, it was quite niche.

[11:25]
What brought you to Queen Mary in 1988?

So having been introduced to AI as an undergraduate, and also, I had worked, I worked during my summers as an undergraduate for the Ministry of Defence on Portland, and so I had had a bit of exposure to AI in some practical applications at that time, I wanted to find out more, and for me finding out more meant doing a PhD. So I was, that's what I, that's what I wanted to do. So, I looked around, and actually there was the great project, the ARCHON project, that was really exciting, so it gave me the opportunity to both be a student and get paid a moderate amount of money. So, I remember my first salary was £10,000 a year, you know, that's what I, that's what I earned as a research assistant, doing my, doing my PhD. And… But it was also great in the sense I was doing a PhD, working on a project that had both industry and academia, and I think sort of, that, that really shaped a lot of my feelings and, and the way that I've worked since then.

[12:40]
And you got your PhD in Artificial Intelligence in 1992. Can you tell us more about your research at the time?

So, the project was about multiple-agent systems. So, in AI, AI is very much concerned with building one smart artefact, one intelligent system that's able to do something interesting or something intelligent. And actually, I always felt that, OK, if you can build one of those, why would you only ever want to build one of those? You

would have a collection of those, so you would have a number of interacting intelligence systems. And that's what multi-agent systems was about. And so, my PhD thesis was in particular how you could get a number of these intelligent agents to cooperate with one another. So, how they could work together to solve a particular problem. And, we had a number of problems in the project that were identified by industrial partners. So, one was in the area of electricity management, and one was actually with CERN, in the particle physics place. And so, I used those as a, as examples to illustrate my thesis and its application.

[14:02]
What were the influences on you at this time?

I think sort of, scientifically, I was very much shaped by being in a big consortium where there were lots of different views, and lots of different perspectives. And I was really sort of, taken that people from different backgrounds were all able to contribute something different to that problem. And that's something that I think has again stayed with me throughout my career. You know, I had come, at this point I was, I was doing this project, you know, I was kind of, in the first year of, of having graduated, so, you know, I had a very narrow base on which to build. I had been trained as a computer scientist. And actually, all these other people were interested in, how people behaved, and how people thought, how, how you could build systems and connect them and get them to communicate with one another. And I, I found sort of, that, that combination of the different backgrounds and the different disciplines really interesting. Because I personally learnt so much, because at that point I knew so little. And then also that mix of, how we could have academic thoughts and come up with some nice academic models. And then you had the industry partners who were actually faced with solving real-world problems for which they wanted some technical solution. And sort of, that fusion of, how we could have interesting scientifically challenging, advanced ideas on one hand, and actually use them to solve real-world problem on the other hand, was, was really a significant influence on me. And I think sort of, that combination of great basic science, and using it for real, in the real world, to make a difference, is something that's shaped my whole career.

[15:56]

About your career. At Queen Mary, you went from PhD to a postdoc, to professor in chair, and you, you became chair at the age of 31, which is kind of, early.

It all did seem to happen quite quickly for me. So, so I thought there were elements of things that I was just fortunate about. So, a year into my being a research assistant, a lecturer vacancy came up in the department, and, amazingly they asked me to apply for it. So, you know, this is one year after being an undergraduate, this is, with no papers, no grants, you know. This would never happen today. You know, it was a huge, a huge gamble, I think. You know, hopefully, I think I repaid the gamble. But, it really was that opportunity to, to get on and, with the, with that first academic post, which many people find, you know, really, a really hard step, and often involves lots of studies, postdocs and all those things, and it kind of just fell into my lap, just because of circumstances, for which, you know, which I am very grateful. And then sort of… So, I finished my PhD while being a lecturer, and working on the, on the project. So, you know, that's quite a lot of complexity to, and time to manage. But, the project, because it was a sort of, flagship European project in, in this nascent area, so multi-agent systems when I was doing it, was really a, a small area. So, there wasn't a sort of, main international conference. You know, there were 20 or 30 people round the world who were, who were interested in it. And so actually, having a, a big European-scale project, gave us quite a lot of prominence in terms of, of the area. And so, ARCHON really was a, sort of, stood out and really gave me that, that foundation on which to, to work. And at that time, the size and makeup of the community, not many people were interested in real-world applications of it. So actually, I got, I got to build up an international reputation very quickly because, because the area was so, so small. There wasn't lots of, there wasn't lots of other people to, to compete with and to get your views out there. The challenge for us as a community was actually being noticed as a very small bit of, of AI, which over time has been rectified, you know, now, now it's a well-established area of AI, but when I started in it, it really was a niche part of a niche area.

[18:42]
Have you got any sons or daughters?

I have two children. I have a daughter called Anna, who was born in 1997, and a son called Matthew, who was born in 2000. So, Matthew is at university, he's, he is doing civil engineering at Bristol, and my daughter did English literature at Cambridge and is now working in local government.

[19:07]
And in 2000 you also moved to the University of Southampton. What brought you there?

So, I felt I had done a lot at Queen Mary, so, I had gone from being a PhD student to a professor in, in about ten years. And, I kind of felt that I had done as much as I could do at Queen Mary at that time. And also, I wanted… So, I wanted to move to a bigger stage, so, Southampton was a, a really exciting and emerging Department at that time. There was a whole load of great hires that they made at that time. And it really felt like a, a vibrant place to, to be. So… I also wanted to go somewhere where the, where lots of the staff couldn't remember me as a PhD student [laughs], and actually just go in as a, as an established professor. So… Also, for family reasons, that's around the time my, my kids were born. We wanted to, to get out of London, and sort of, be in a, in a, a slightly better environment for personal life, and a place where you can have a garden, and you can go out for walks and things like that.

[20:24]
What was your role in the Southampton's Agents, Interaction and Complexity Group, AICG?

So, when I went to Southampton, we had a, it was to the Department of Electronics and Computer Science, and we had, we had quite a traditional sort of, group structure. So, we were… It was a big department, it was kind of, 100 academics, and so you need to have some group structure in that. And so, we got together a group of people who were initially interested in sort of, AI and the Web and multimedia. So that was the first group that I joined in Southampton. And when I joined, that was headed by Dame Wendy Hall. She moved on to become the Head of Department, and then I took over the group. So that group was called Intelligence, Agents and Multimedia. And this group was a, was amazingly successful at the time. So, it had Wendy Hall in

it, it had Professor Sir Nigel Shadbolt in it, and it had, myself in it, and a number of other really prominent computer scientists. And, actually, it was a really impressive period where there was lots of action. I felt we were really at the centre and really driving what was going on in computer science in the UK and, and internationally. But actually, as we were so successful, our group became an ever-bigger part of the, of the department, and it sort of became slightly unbalanced in terms of the size that it was in the sense that our group kind of dwarfed some of the other more traditional groups. And so, to rebalance things and make sure that, sort of, all the strengths of the department were represented, the Intelligence, Agents, Multimedia Group got split into two. So, one got split into a, to a Web-based, so World Wide Web-based group, and then one became an agents' group, and we merged together with some people who were interested in complexity science, and, and interacting systems. And so that's what the, the Agents, Interaction and Complexity Group grew out of.

[22:45]
So, at Southampton, you got several grants. And, can you guide us through your grants from ALADDIN to orchid?

So, as I said, actually it was a very, very exciting period, and, I really felt we were at the forefront of, of computer science. And so, actually, we were able to get a good series of grants that, that let us push forward this agenda. So, ALADDIN was the, was a, funded by BA Systems and the EPSRC, so the Engineering and Physical Sciences Research Council. And actually, that's the first big project that, that I led. So, so ALADDIN was a, was a multimillion-pound project. It was looking at agents, which was mainly what we were doing in Southampton. It involved some machine learning, which involved some colleagues, some, what are now colleagues here at Imperial College London, and some people who were interested in decision-making under uncertainty at the University of Oxford. So, so it was a broad collaboration between the, between those universities. And again, it was about fusing together different approaches and different methods to be able to build particular systems. And so, we worked closely with BA Systems. We chose to illustrate many of our concepts through the area of disaster response and disaster management, because it's, it's an important domain itself, and its indicative of the kinds of applications where BA Systems could see themselves deploying these sorts of technologies. So,

ALADDIN was the first, my first real chance to sort of, lead a, a good size, a good size project.

[24:44]
Where did ALADDIN stand for?

So, ALADDIN stands for Autonomous Learning Agents for Decentralised Data and Information systems. And that's the amalgam of what the academic partners were in, made into some acronym. We had a series of related projects that, that flowed from this. So, looking at that broad set of methods, so, AI, machine learning, multi-agent systems, looking at particular domains. So, some were in supply chain management, some were in smart energy systems, and also in disaster management. So, again, that sort of broad suite of application areas. And then sort of, I brought those together into, into a programme grant called ORCHID. So, this was a EPSRC programme grant. It was a big five-year grant. And what that did was, we took a number of the partners, so we took the Southampton and Oxford base. And actually, the thing that we did differently in ORCHID was, we introduced some humans into our multiple-agent systems. So, up to now, most of the, most of the systems I built, all the agents, all the actors, were bits of software. So, it was a software representing a particular task, or particular function. And, you can only… You can only go so far with that, because actually, what we found is, as we wanted to deploy these applications, actually, people didn't want to employ them in a completely autonomous fashion; they wanted to have humans involved in that decision-making process. So, in ORCHID, for the first time I had started building systems that had users and people as part of the multi-agent system, and that involved bringing in the University of Nottingham and Professor Tom Rodden, and that expertise. So, ORCHID was the first project where I was interested in, in what we termed human-agent collectives at the time, so, humans, people working together to solve particular problems. And our main application area that we worked on was in disaster response. So, following on from ALADDIN and some of the interim projects that we had had, we, we really built up expertise with that.

[27:18]

And you worked with Rescue Global, deploying a system of machine learning, information fusion and coordination algorithms to inform their response to the 2015 Nepal earthquake. What did you learn in the project?

Rescue Global were great collaborators with us. So, they're a disaster response charity. So, their aim is to be some of the first boots on the ground when a disaster happens. They fly there, they, they have their own jets, they go in, they help stabilise, they help, understand what is going on on the ground. And actually, they were amazingly progressive in terms of what they were willing to try. So, they were willing to try out some of our prototypes and some of our methods, alongside some of their traditional ones. Because, you know, when they're deployed, they're doing it for real, they're saving people's lives. So, so they want to make sure that they are able to fulfil that function, but they're always continually trying to learn and get better. And so, they use some of the ORCHID technologies to help them respond to a disaster situation. And in particular, they, they use some of our inference technologies and some of our learning technologies to understand the situation of what had happened in Nepal, in particular around Kathmandu, and one of their first challenges was to deploy water filters. So, they had the problem of figuring out, where should they put these water filters, and water supplies, to, to best respond to the disaster? So, to respond to as many people as possible, and save as many lives as they can. And so, they had a particular process for doing it, they, they did some manual calculations, they took some maps. And in parallel to that, we ran some of the ORCHID technologies. And we identified quicker a number of the areas that they identified, but also, we identified some areas that they didn't, and when they went there, they found that actually these were good areas to deploy in. So… So, it was a great illustration of working closely with an industrial partner to, to take state-of-the-art research and to see it deployed for real-world benefit, yeah, that helps save people's lives.

[29:40]
So, talking about the collaboration with industry, what about the establishment of start-ups like Aerogility?

So, I've been interested in working with industry throughout my career, as I've said. A lot of the early career I would say was working with big industry. So, I worked with

companies like BA Systems, British Telecom, Hewlett-Packard, those sorts of big companies. But, I then also started, and sort of, it was a time when the start-up scene was starting to emerge, so actually, when I moved to Southampton, some people from a predecessor company of Aerogility got in touch with me. So, they were quite interested in multi-agent systems and in some of the work that, that we had done. And so they approached us about trying to work with them. And, for me, again, it was a real learning experience. I had not particularly worked with, with start-ups before. But it was, you know, a very interesting pace of work, the things that they cared about, the, the speed with which they had to get things done, the, the focus on really getting customers, and seeing your work deployed, was, was really vital for them as they really sort of, lived from day to day and week to week and, and month to month. And so, I became very interested in, particularly in technology start-ups. So how, how work can go out of university labs, so, from my lab and, and other labs, and really be the basis of a, of a new company, something that's being started up, that, that's helping to create jobs, that's helping to create a place for itself in the world. And I really hadn't come across those sorts of companies at such a personal level until, until I started working with Aerogility. And I've worked with them for, how long? For nearly twenty years now. And we're still working and refining our product. They've got a range of good customers, so, companies like EasyJet and Rolls Royce and BA and Lockheed Martin are using their, their product. And they've gone through a whole evolution of the dotcom boom and the dotcom bust, so, you know, they, they expanded a lot, they've also shrunk a lot, but they've, they've kept going, and it's been really interesting to sort of, see that journey.

[32:20]
What are you working on at the moment with Aerogility in particular?

In particular, we're looking at, at finessing their product. And their product is an agent-based simulation of aircraft and aircraft engine in order to try and optimise the, the repair processes. So what you want to be able to do in that industry is, you want to keep aircraft out working, doing their jobs, in the air, for as long as you can, and you want to bring them in just before any failure might occur. And if you think about doing that across a whole fleet of aircraft, so, you know, a, a company like EasyJet has many of those aircraft out there, and they want to optimise which ones they bring

in, at what time, in what place. And so, trying to solve that problem is a very complicated problem. You have lots of different agents representing the individual planes, representing the operators, representing the people who run the repair shops where, where you want to be able to take them. And so it's a big complex multi-agent system, that we're really trying to optimise the performance of, of the whole system. And what's interesting in the simulation part is, actually you can run things forward and say, what happens if we, if we introduce another factory for, for repair in a particular country in the world, how would that affect the, the amount of time that you could leave people, you can leave planes in the air, and you can move people around. And so, it gives you that management information, and also a, 'what if' simulations about, about the future and the environment.

[34:16]
Let's go back to 2004, when you built a sensor network system to gather information on the impact of climate change in the Briksdalsbreen glacier in Norway.

It was a very interesting project. It was again sort of a, an interesting domain for me to be able to work in. I had not particularly worked in the environmental space. I had worked much more in the industrial space. And the team at Southampton had been working on placing sensors in glaciers to be able to track what was happening deep inside the glacier. So, you can see what's happening on the top of the glacier quite easily, but actually they wanted to be able to figure out what was going on inside it. And so they had spent a lot of time constructing these individual sensors that you could blast into the glacier and they would sort of, flow through the glacier, they were like big eggs, big plastic eggs if you like. And sort of, they would track data, give you velocities, temperatures, moisture, and those sorts of things. And so that gave you a sensor network. But what they didn't quite know was, how to, how to control all of those, and, and coordinate their behaviour, so that you can get maximal coverage of the glacier for as long as you possibly can with the battery life that you have in the sensors. Because of course you can't plug them in to anything. So, it's a, it's a optimisation problem, again, on battery life and communications and coverage. And each one goes in on its own, but you want to be able to coordinate them. So… So I came into the existing team that Kirk Martinez led, and sort of, provided some algorithms that helped them do this coordination. And it was a, amazingly interesting

project, but actually, in the end the, the sort of, the result of this was that, the, the glacier really eroded very much quicker than, that we had hoped, unfortunately. So this is a, this is a global warming, climate change story. So, there was a plan to run, look at the glacier and continually track its progress and performance over a ten-year period, but actually by halfway through the project we, there was no glacier left. It had really melted away. And so, we, we had, the team transferred it to a different glacier in a different country. But again, sort of, that really brought home to me in a very direct and personal way sort of, the impact of, of climate change on our environment.

And what year was this, like, 2009?

Yes, that sort of time.

[37:04]
Between 2010 and 2015 you were Chief Scientific Adviser for National Security for the UK Government. What can you tell us about your role there?

This was a, a great job. I really enjoyed it. So most government departments have a chief scientific adviser, and they are often sort of, senior academics who are brought in to, to provide advice across a whole range of scientific disciplines. Now there wasn't one of these in the national security space, so I was, I had the privilege of being the first one of those, which is both good and bad. So the good bit is, you can shape it very much into what you want it to be. The bad bit is that, there is quite a lot of structures that you need to put in place. You need to worry about resources, and, and influence, and how you do your role, when there, when there hasn't been one of you in the national security space before. So, so that was, that was fun, setting all of those things up. I really enjoyed working in national security. So, the science challenges of national security are, are many-fold, so, some of them were from my own areas of science that I understood well, so, cybersecurity, big data, and analytics to understand what's going on in a particular situation. So that was great. I was on firm ground there. Other areas of science that I had to understand and to grapple with were things like, home-made explosives. So, how, how people can make bombs at home, in their factory, and, what worked and what didn't work, and what you might

try and do to, to mitigate some of those. Some of the human and behavioural sciences, about, what makes people do terrorist activity, why, why they behave in a particular way, what interventions can you make to try and stop those things. So, so my role was to come in and oversee all of that science. There was some great science that went on, but it hadn't necessarily had a, a senior person leading it and coordinating it. So there were lots of, there were a decent number of scientists, but they didn't really talk or interact with one another. And also they didn't really have a longer-term agenda. So, so many folk in the national security area are very good at solving a particular problem that needs to be fixed today, or next week, or, long-term might be next month. And that's, that's fine when you're responding to operational requests, when you're trying to disrupt a terrorist organisation, or you're trying to infiltrate a terrorist organisation, or to, or to spy on a, on a foreign government. But actually, there's also a long-term element that you need to be able to take into consideration. So, so I set up a, a science function that helped us prepare for those longer-term perspectives.

[40:10]
What do you think you learnt from this experience, like, five years' experience there?

I learnt a lot about trying to be successful and get things done outside of a university environment. So, that was, I think, the big thing that I, that I learnt, is that, not everyone in the world has the same set of values and the same set of things that we do in, in universities. So up till then, you know, I had worked solely in a university context, I had worked with businesses large and small, but I hadn't worked in a government organisation. And so, understanding that science was an important aspect to it, but it wasn't the only thing that was taken into account, was very important. And the value systems, so, you know, the fact that I might publish in a particular journal, or, or have a particularly exciting paper, actually, everyone was highly indifferent in, that, you know, they really didn't care. They were all interested in what you can use science for. And so, trying to influence and persuade people and get things done, required a whole different set of values and a different set of arguments to get things, to get things done. I mean, the, the community were very receptive to it, but it took a while to, to figure out how to be effective and what to be effective about. So, so that was very interesting. As part of the network of chief scientific advisers,

we all got together on a regular basis. So, I was responsible for the national security areas, but we did things as a network of chief scientific advisers, under both Sir John Beddington, who was, who was the CSA, the Government Chief Scientific Adviser, who appointed me, and then under Sir Mark Walport who, who was his successor. So, I spent three years broadly working under each of those. And the network of chief scientific advisers, we would tackle things that would be of interest to, to broader government. And so, I learnt how government work, I learnt why government works, and sort of, really, helped my understanding of how the country is governed.

[42:20]
You often refer to the relationship between humans and AI systems as a relationship between partners. Would you like to say more about this?

So my model of AI systems, and particularly multi-agent systems, was very much shaped by what we did in the ORCHID project. So in ORCHID, we were interested in the, how humans and software systems could work together. And humans are, I think humans are amazing problem-solvers, and, really good at certain things that machines are simply not good at, and I am unconvinced will ever be really good at, or as good as humans at. But having said that, there are things that computers do that humans are, that are way better than, than us as humans. And we see that through, through, we've seen that throughout computing, you know, computers can add up better, faster, more accurately than people can. So, then you have things like spreadsheets that come into operation and automate some of those tasks. I think as AI progresses, some of the, some of the tasks that are better suited to computers and machines will be taken on by AI systems, but there will always be activities that are best done by humans. So if you think about our disaster response application in ORCHID, one of the tasks that we needed to be able to do was to plan a schedule of which, which rescue teams would go to which bit of Nepal at what time, with what resources. And previously Rescue Global did that as a hand calculation, they would sort of say, 'Right, we're going to send some over here, and we're going to send some over there.' Now that's an optimisation problem. You have a set amount of resources, a set amount of times and priorities, and actually machines are very good at solving those sorts of things. And so we automated that. But actually, what we found was that when we had just automated the solution on its own, the human problem

solvers would always want to be able to say, 'Actually that's nearly right, but I want to tweak this bit, and I want to send this team to this area, and this team to this area,' for reasons that I understand, or have intuitions about that are not in the data or are not, the machine's unaware of. And so they wanted to be able to tweak things. And so, that model of humans and machines working together as partners where some of the tasks are done by humans, some of the tasks are done by machines, and there's a flexible interplay between that, has really shaped my view of AI. And as AI has become sort of, increasingly into the public consciousness, I think sort of, this, this mixed systems of humans being in the control, and sort of, being supported and augmented by smart computer systems, is, is starting to gain prominence. A lot of the early debates were about full automation where you take humans out of the loop, and you just automate everything, and… And while that might work in some cases, I think the dominant paradigm for where we're going to see AI deployed in the future is in these mixed human AI partnerships.

[45:45]
This leads me to ask you the question about, what's the future of jobs in this vision?

So, I am… There is a whole spectrum of thoughts in my field about what this will mean for jobs. So I am, I am an optimist, which means that I think that humans do have a unique place in society, and will always have a unique place in society. I think there will be shifts of what people do and what machines do, and that some of the things that people do today I think will be automated away in the future, but I think there will be new jobs created and new opportunities created by AI systems. And so, I am confident that there will be significant employment for humans for the, for the foreseeable future. And that our jobs will be made better and easier by AI systems.

[46:50]
And what about decision-making?

I'm very much of the view that, that variety and heterogeneity in decision-making is a really good thing to be able to do. And so I like, I like sort of, to see AI systems as, as a challenger or as a checker for human decision-making. So, if you think about some of the medical diagnoses and medical applications of AI systems, I think, actually, a

number of, a number of the protocols involve two humans making decisions. So, if you are looking at, at radio, at radiograms and cancer images and those sorts of things, you, you waste a lot of time by having two people looking at it. I think it's much more sane to have a, a person looking at it and a machine looking at it. They look, they detect different things. And sort of, challenging one another. So the machine saying to the human, 'Why have you passed this one?' and, 'Why do you want to, why do you want to focus attention on this?' So, I think they're good cross-checks, and good counterbalances of one another.

[47:55]
Thank you. Let's go back, let's go back to your career. After sixteen years in Southampton you moved to Imperial College in 2016. So what brought you there?

So, I wanted at this point to… I had really enjoyed being a chief scientific adviser, that was a, a great experience, and sort of, working at strategic level of science leadership across a broad range of disciplines. So, I think that's what I really enjoyed about being a chief scientific adviser. So, as I said, you know, some areas I know a lot, I knew a lot about; some areas I knew less about, was just intrinsically interested in and curious about. And so I was interested in a role that had those sorts of characteristics. And the university role that has those characteristics is often the pro-vice-chancellor for research, or, as we would call it at Imperial, the vice-provost for, for research. And so that gives you a strategic view over all the research that's going on, and sort of, the ability to influence and to lead it. And so I, I wanted that kind of role. And timing was again on my side, the role came up at Imperial, and if you want to be a pro-vice-chancellor for research, I think Imperial is a fantastic place to be that. It's a world-class university with great science going on. And so that was sort of, the, the big opportunity for me. So it was very much the, the role and the place that I wanted to come to.

[49:36]
Can you talk a little bit more about your roles here at Imperial College? Like, what are your, what you are doing, what's your strategy?

Yup. So, a lot of the role of what I do here at Imperial is, I tend to look at things that cut across disciplines, so multidisciplinary research. We have people, we have heads of department who run individual departments, we have deans who run groups of departments, so we have a dean who looked after all the engineering departments, a dean who looks after all the medical departments. And so I tend to look at the things that cut across those things. So when our engineers want to cooperate with our medics, or, people in natural sciences want to collaborate with our business school. And so, many of today's big societal challenges are not within one discipline, or even within one group of disciplines. They are, really cut across those things. And so, multidisciplinary, transdisciplinary research is really important for the future for tackling big societal problems. And, and that's a key bit of my role. So, figuring out what the strategy is, what are the key areas that we want to address, to really be able to tackle those science, big societal challenges that can be underpinned by, by science. So again, for me, the role is very much about, having an impact, doing some great science, bringing together different disciplines and seeing what difference it can make to the world. Which is a very Imperial, very much at the heart of what we do at Imperial.

[51:15]

And the other part of my role, so I'm the Vice-Provost for Research, and enterprise is, the enterprise part involves leading engagements with industry, with our spin-outs and our start-ups, for both staff and students. So again, that sort of plays to my, to the experience that I have in working with industry, something I've done throughout my career, and start-ups and spin-outs. And so, that's, that connection from our research to our spin-outs and our corporate partners is, is really important, and exciting for me.

[51:51]

And, talking about collaborations with industry, you collaborate with several companies. Can you tell us about your collaboration with Crossword?

So, I've worked with a number of smaller companies and start-ups throughout my career. We've spoken about Aerogility. So, Crossword, in fact the, the founder of Crossword, Tom Ilube, was the founder of the, the company that became Aerogility, so Tom is someone that I've known for, for 20 years. And he was very interested in setting up a, a company that was looking to take research from universities, in the

broad area of cybersecurity, and make them into products. So, so it brought together many of the things that I was interested in, so that sort of, engagement piece academia into companies. It brought in the experience of cybersecurity that I had learnt as national security adviser for, for cybersecurity. And so, brought those, brought those things together. So I am an adviser to them. I talk to them about how they can engage with universities, what kind of products are likely to be interested, in, for the, for the national security community.

And do you have specific collaborations with Imperial College for Crossword? Or does it…

Not spe- So, my, my connection is an advisory one. There's no, there's no direct research that I do, or anyone does at Imperial, for, for Crossword.

[53:30]
OK. And what about Darktrace?

So Darktrace is a recent collaboration of mine. They're, they're a start-up but they've, they've, start-up, and really grown to, to over 1,000 people. They're a unicorn, so, you know, they have a valuation of over a billion pounds. And, they again sort of, bring together some of the things that I'm most interested in. So they're very much an AI company, deploying a system to protect networks and infrastructure. So that combination of AI and cybersecurity is, is two of the things I, I know a reasonable amount about. They've been amazingly successful and grown really quickly. And they're at the stage where they're starting to, to really use and push the envelope around, around AI and what, what kinds of products they want to do next. And so again, it's an advisory role, helping them understand what the art of the possible is around AI and cybersecurity.

Who approached you from Darktrace? How was the collaboration established?

I knew a number of people who founded Darktrace from my time in the national security community. So, many of them went to form Darktrace, so many of the people who worked in some of our more technical areas formed Darktrace. So I, I

had known them for a long time. Obviously while I was in Government it wasn't appropriate to, to be working directly with them, but when I, when I left, it became a much easier collaboration to do.

[55:06]
So in 2014 you became Regius Professor of Computer Science at the University of Southampton. And what are your memories of that moment?

So I was really proud to be the country's first Regius Professor of Computer Science. So, I think it, it meant a couple of things for me. So, firstly, the Regius Professor is a, is a title that's created by the monarch, and they've existed throughout history, so some Regius Professors in this country date back three, four hundred years. So, they tend to be in more established subjects, so, often around theology and philosophy and things that were around then. The Queen created a number, a number of new ones in relation to her Diamond Jubilee, so, I think ten or so were created. And I thought it was great that actually the first one of these was, or one of these, was in computer science. So, it was great to have, for my discipline, as a computer scientist, to have that recognised by a Regius Professorship. So that was great. It was then fantastic that actually the, the place that was chosen to hold the Regius chair was the University of Southampton, it really brought home and sort of, endorsed many of the great things that had happened in Southampton over, over a twenty-year period, really building up computer science, really recognising it as a top place to do it. And, and so sort of, the, Southampton getting the possibility to host one of those was fantastic for, for the discipline and for the university. And then to be able to, to be chosen to be the first one was a, was a great personal honour for, for me.

[56:56]
And in 2016 you were also awarded the Companion of the Order of the Bath, in the Queen's New Year's Honours List, for your services to computer science and national security science. And what are your memories of that day?

So I think there are a couple of days to pick out. So the first is when you get that letter. So, before it's all public. So, the, the Queen is minded to bestow this honour

on you, which was, which came completely out of the blue, and, you know, was a letter, you think, I'm glad that I didn't get lost in the post. [laughs]

What does the letter look like?

It looks like, it looks… It looks quite a grand letter, but, you know, it, it says, 'The Queen is minded to…' It doesn't say, she will. So, you know, there's some doubt in it. And, and you have to say… 'If you were to be asked, would you be willing to accept?' And so you have to reply, so, again you have to hope that actually, your letter saying yes doesn't get lost in the post, is where, you worry slightly about our, our postal system. So, there's the first bit about the letter that you get. Then there's a couple of months of silence. You, you've received a letter, you've said, 'Yes of course I, I would be deeply honoured.' And then there is a period where you know, and your wife and family might know, but of course no one else knows. And then, sort of, there is the, mine was in the New Year's Honours List, so sort of, just before the New Year it's publicly announced, and then sort of, there is a, the fantastic celebration, and lots of people getting in touch with you to say congratulations. So, not only people that you're in contact with at the time, but also have been in contact throughout your career. It's amazing the number of people who re-established contact and said, 'Fantastic.' So that's another really great day. And then the third great day for me was the actual day where, where the Queen gives you the award. I went with my family, so my wife and my two kids. My parents and my in-laws also went to Windsor for the day, and so, we had a, a lovely day in Windsor celebrating it.

[59:16]
OK, so, if you look back at your whole career, or at your whole life, what would you do differently if you had your time again, and why? If there is something of course.

I think the, scientifically, the thing that I would do differently, or, or scientifically that I wish I had spent a bit more time understanding, was, is about human beings. So, I think I spent a lot of my time in my career focusing on software and automation and less time thinking about people. And sort of, I wish I had spent a bit more time in my career looking at the impact of societ-, the impact of people on technology, and the

impact of technology on society. So I do more of that now than I had done, and I wish I had got into that earlier in my career.

[1:00:07]
And, what are the proudest achievements of your career?

So, some of the things that I'm most proud about are actually the people that I have worked with during my, during my academic and, and government career. So, in terms of my university life, I've, I have now graduated over 50 PhD students, and that's, I'm really proud of them as a group. Some of them have gone on to do great things in, in academia, some of them have gone in to do great things in, in industry, and some have gone on to do great things in other walks of life. So, so that's sort of, 50-odd PhD students. I've also had the pleasure of having about 50-odd postdocs, so in total, that's about 100 people I, I can claim some responsibility for in terms of their, of their careers and, and contributions, and, and some of those have really gone on to shape the field. I left a great group of people when I moved from Southampton to Imperial, and they continue to flourish there. So I, so I think, sort of, the people impact is something that I've been very proud of. And then also, I think the, the other thing I would highlight is the, is the fact that you can do good basic research, and do research that actually has an impact on the real world, and that those are not mutually exclusive. So many people tend to do one or they tend to do the other. And I think, I've always wanted to do both, and I think, I think what I'm proud of is that I've illustrated that, that kind of connection within the same group, within the same team, is entirely possible.

[1:02:03]
What do you think are the biggest challenges and opportunities for AI in the next ten years?

AI at the moment is a, is going through a, an amazing time of interest. And, I'm slightly worried that, that we're going slightly over the top in terms of what AI can do for society. I think it will do a number of great things for society, and I'll talk about those in a moment. My slight fear is that we're going to go, we're going to go down the, the, towards the trough of disillusionment in terms of it as a, as a field, and that

people, we will not be able to deliver in the timeframes that we think we will, everything that's being promised at the moment. So I see lots of naivety around AI. So I see lots of people who have got really difficult problems that they have had for a long time, and, and they say, 'And we're going to use some AI,', or, 'We're going to use some machine learning,' and magically that will solve all our problems. And I think those people are going to be hugely disappointed in, in what's going to happen. But I think where the… So I think that's a challenge. And I think the challenge is around managing expectations, about getting people to understand what, what AI is, and what it can do, and what it is not and what it can't do. So I think that sort of, much greater realism on AI is, is really important that we get right. It's a powerful computing technology, but that's, that's what it is. It is not magic. But, having said all of that, I think, we really will see a number of transformative applications appearing in the coming ten years. So I think we'll see great applications of AI in healthcare, I think that's a, that's the most obvious place where we're going to see some of these things. Already. So almost on a daily or weekly basis, you see yet another application of AI dong something really useful, really important for society. So, so I think medicine is a, is a real example. I think sort of, energy, the future of energy, and climate change, is really important, to use AI to help manage resources, to help us discover new ways, new science, new ways of doing things in the energy space, is really important. And I think in education, I think, the way that we educate people throughout their lifetimes, both when they're at school, and when they're at university, and, and throughout their continuing professional development, I think all of those can use AI to really help tailor that learning experience to you. So, you know, if you're fortunate enough to work in very small groups with, or have an individual tutor, then, then there's lots of evidence that that's, that's a good sway of learning. AI can help you tailor those experiences, can help probe, and test things that are specific to you in a way that a good teacher and a good educator can.

[1:05:06]
How do you see the role of the Alan Turing Institute in, in this country, for AI?

So I think the Turing Institute is, is really important. So I think it's great that we have a national institute for AI, that really brings together some of the best researchers in the area, gives a concerted mass, and gives industry and government and society a, a

single place where they can go and find out about what's real and what's not. And so I think the authority that comes from having a national institute is, is really important.

[1:05:47]
What advice would you give to someone willing to pursue a career in AI today?

I would say that, you've made a good choice. AI is a fantastically exciting area. It really has opened up opportunities to work with a whole variety of people from a whole variety of areas of science and society that I would never have envisaged when I was training as a computer scientist. So, it really is an underpinning technology that's really important and really impactful.

Thank you Nick for your time today. It's been a real pleasure talking to you.

It's been a pleasure. Thank you.

Thank you.

[End of Interview]