Capturing the Past, Inspiring the Future

Dr Catherine Breslin

Interviewed by

Mark Jones

11th March 2019

At the

WCIT Hall,
32a Bartholomew Close, London, EC1A 7JN
Kindly provided by The Worshipful Company of Information Technologists

Welcome to the Archives of Information and Technology. It's the eleventh of March 2019 and we're in the London offices of the WCIT. I'm Mark Jones, an interviewer with the Archives of IT. Today I'll be talking to Dr. Catherine Breslin. Catherine is a machine learning scientist specialising in speech and language technology. She has a distinguished academic background with a degree from Oxford and post-graduate qualifications from Cambridge. Her career has been about the application of speech and language research to real-world problems. She has worked for Toshiba and then Amazon where she played a major role in developing Alexa onto the other platforms. Today she works for Cobalt Speech advising clients on their own specific speech and language technology issues. As well as talking about her own career to date I hope she'll be able to share her insights into the possibilities for others wishing to make a career in the machine learning and voice technology space.

Morning Catherine, thank you very much for coming to talk to us.

Thank you very much for having me.

So, in the usual way, if we could, we'll perhaps start with your schooling and education and have a walk through that timeline. Can we start with your secondary school, up in Yorkshire?

That's right, yes.

You moved there just before you started secondary school, didn't you?

Exactly, SO that's why I don't have a Yorkshire accent actually [laughs] which people comment on.

So, secondary school - I see you ended up doing A levels in maths and physics and STEM subjects. Were you always that sort of pupil or did that come halfway through school?

I was always interested in maths and science, I think, as far as I can remember. Maths was probably my favourite subject all the way through school, I remember maths back

in primary school and learning about maths and finding it very exciting to be able to deal with these exciting different problems through my school career, and so it sort of developed over time, and as I went through school that just got stronger. I was also interested in the sciences and, in particular, physics which I think is a very applied mathematical subject as well, and so as you get more into physics and to GCSE and into A level you start to bring maths into physics and I always enjoyed that, the application of the maths that was learned in the maths classes to the physics that was learned in the physics classes.

Right, and at school were you given a lot of encouragement to get into sciences?

Yes, I think my teachers were very encouraging through the time. I think I had a natural inclination towards sciences, they were my stronger subjects anyway, so it was kind of natural that I would go into those subjects, and when it came to choosing A levels I had very supportive A level teachers in maths and physics who helped me choose the right modules to take and the right choices to make there.

00:02:57

Was it unusual for people from that school to go to Oxford and Cambridge?

There were not too many of us from that school who went to Oxford and Cambridge but my physics teacher had been to Oxford which meant that he knew about the system and he was very encouraging, so I remember, we went on a school trip to, I think it was the Oxford Open Day, the Oxford Science Open Day, while we were doing A levels and three or four of us from my year applied to Oxbridge and two of us got in, which I think is unusual for a school in Yorkshire but I think it was encouraged by both the teacher who went to Oxford and the other teachers who were teaching at the school who were also encouraging us to aim high.

I think that's one of the important things, isn't it, people's school where there's no track record of people going to Oxbridge, and I know staff who don't have that track record either, they perhaps just don't become aware of the possibility of university education generally and perhaps Oxbridge in particular.

Yeah, and I think taking us to Oxford for the day was a great thing, just to show us that it wasn't this mystical place where we wouldn't fit in but that there were real people there and that they were doing really exciting stuff. Obviously they took us to other universities as well, York being closer to where I lived and I remember going to York and to Durham and a few other places, but they did really make an effort to encourage us to see Oxbridge as an option was a good thing.

So, with your A levels you could have done maths, I guess, but you chose engineering. Why was that?

So like I said, I was always interested in physics as well as maths and I really liked taking maths and applying it to things, so when I was doing A levels, for example, my favourite parts of A level maths were the mechanics modules where you could really see the maths and see how it coincided with the physical world. So you could take things in the physical world and model them. I never actually considered doing maths at school, I was torn between physics and engineering and, in the end, I decided that engineering was the more practical of the two, of physics and engineering, and so that's where I decided to go and focus.

00:05:05

That's interesting because I did maths at A level too, but pure and applied, and by the time I'd read maths at Cambridge for two years I'd found it really hard, actually, and the pure maths just left me cold and I found that more and more abstract and more and more difficult.

Yes.

It was the applied maths I was interested in but you couldn't just do a degree in applied maths.

Yeah, I think that when I talked to the mathematicians, when I was at Oxford and obviously was studying, I found the same. I was not that excited by this pure maths

that they were studying, the abstract theoretical principles. I much preferred what we were doing, when we applied the application of it to the real world.

Yeah, I think maths at university should come with a health warning, actually, 'cos I found it wasn't just an extension of A level, it was like a different subject really and if I'd known what it was like I wouldn't have started doing it, if you know what I mean. [laughs]. Ok, so off to Oxford for engineering, which college were you at?

I was at Trinity.

Right. Did you pick that specially or that's happenstance?

I think it was just a coincidence. They sent us a booklet when I was applying, with all the colleges in, and I just looked through and Trinity seemed like a nice looking college. I didn't really have a lot to go on. Obviously I hadn't been round any of the colleges or seen what they were like so it was just pure luck.

So it all worked well, you enjoyed Trinity, you enjoyed Oxford?

Yes, I did.

So, you said in the notes that you were sponsored by GCHQ – how did that come about?

So when I was doing my A levels, and I thought about where to go to university and was applying for engineering, I knew that sponsorship was an option and lots of companies did, and I hope still do, offer sponsorship for undergraduates through their studies. And this was about the time that the government had brought in tuition fees for university so how to fund university courses was topical and on a lot of people's minds. So I thought that, as I was going to study engineering, I ought to find sponsorship and so I wrote to a batch of different companies who offered sponsorship and GCHQ was one of the ones who came back and interviewed me and offered me a job.

So did you actually work at GCHQ as well?

I did. So I worked for four summers at GCHQ, so the summer before I went to university and the three summers in between my degree years. We did short projects over the summer so it was essentially like a fourth university term and I felt like there was a great community of other students there as well so we lived together and we socialised together and we worked on projects.

Interesting. So were you right inside all the security stuff or were you a bit on the edge of it?

We were working on projects within the technical part of GCHQ and what was great about this, and all internships I think, is that the chance to take your studies and apply it to actual real world problems is a great way to learn more, to sort of reinforce what you've learned in university but also to see why it's useful to learn most things and how you can use it in the real world.

What's GCHQ like as a place to work in? Because I know some people who have been there, some are really intrigued by it and enjoy the work, others find it a little bit intimidating, scary.

So I wasn't actually there for very long, in the end, I just did these four summers, but like I said there was a great community of other students and they really looked after us and so I met a lot of other great people and had a great time, and also got to work on the things I was studying about so I found it a great place to do that.

Interesting, ok. So, back to Oxford then, engineering. Did you find that easy, challenging, what you expected?

[Laughs] It was quite challenging I think, if anything. University is a big change, obviously, from A levels, and you had to learn to be a bit more motivated for yourself, to make sure you get things done. You don't have teachers keeping you on track all the time so learning how to get stuff done at university was a challenge. And also the topics get a lot harder, you get to learn a lot more stuff but also a lot more in-depth.

So this was also your introduction to computing as a sort of discipline, was it?

Exactly, yes.

How did that come about, just part of the course was it, there were some modules?

Yes, so I actually, throughout my childhood and throughout my school, I'd used computers, like we had a computer at home that I used, I did a little bit of programming, a little bit of writing web pages about the time that the Internet was just taking off. So I was a bit familiar with computers but I'd never, nobody had ever suggested it to me as a topic to study before and I didn't really realise that you could. So then I went to university and in my first year of Oxford one of the courses you take is a computing module and that was the first time I realised that this stuff I'd been doing as a side hobby as a teenager was actually something I could study and so I decided ... at the time Oxford offered a joint honours in engineering and computing science which you selected at the end of your first year and so that's what I decided to do, go and do a joint honours.

Ok, so did you enjoy the computer bit as much as the engineering bit?

I did. The computing side, I think, bits of it veered a little more on the theoretical side like the maths, maths and computing are quite strongly related once you get into the theory of computer science. But I took some of the applied computing courses and some of the more information courses from the engineering degree and between them I covered quite a lot of interesting computing theory and applications.

Do you think your degree, is your degree something you use regularly, so to speak, no,w or is it just a stepping stone to what you now do?

I think it is something I use regularly because a lot of the theory of what I do now, the first place I encountered that was in my degree course, and a lot of the computing that I learn is typical of a lot of the applied maths I learned there is applicable too.

That's interesting because I find the same really, with my computer science degree. It was quite theoretical and every so often it's good to have that theoretical underpinning because it explains why something works, doesn't it, or how it works.

Yes, I think that's true and understanding computer theory and some of the ... I remember doing algorithms and data structures and things like that. They're really important when it comes to building and designing computer systems to do things and so having that theoretical background to fall back on is really useful.

00:11:46

So, you took a machine learning course towards the end of the degree, is that a random choice or was there an inkling that it might be what you'd be interested in?

Well, like you said I was interested in the information. So engineering tends to be split up into different strings, so you have mechanical, civil, electric – I forget the others – and probably lots of people ... [laughs]. One of the ones I was most interested in was the history of information in engineering which covers things like signal processing, machine learning, control theory, those sorts of computer related kind of thing and engineering subjects.

00:12:54

Ok, so what sort of things did you learn there, is it just the basics of speech processing?

In Oxford the course focused a lot more on medical imagining rather than speech processing. So I learned some signal processing at Oxford, some control theory and then a lot about graphics and imaging and computer vision. When I did my fourth year project – everybody does a research project in their fourth year – mine was focused on medical imagining, so using ultrasounds to detect breast cancer. So, that was my first introduction to using computers to do these sorts of vision tasks.

So did you have any idea at this point of what career you might end up doing?

No. [laughs] No, I didn't, I was interested in the topic and taking a machine learning course opened my eyes to the kind of methods you could use to get computers to do some of these tasks. So things like looking at an image and seeing what's in that image is something that people find very easy, most children find very easy, but getting a computer to do it is really hard. And at the time there weren't good systems that could do this – things have changed a lot since then – but I was really interested in how you could build these systems which would do these tasks which people could do but we didn't know how to get computers to do them.

00:13:51

So intrigued ...

Yes, intrigued.

Ok, so after Oxford – oh I see you did a bit of rowing at Oxford...

I did, yes, well I was cox.

Oh, very good, that's the less strenuous bit! [laughs]

Everyone has to try it out but I didn't like the mornings very much. [laughs]

I'm with you there!

The early mornings on the river but sometimes it was quite beautiful down on the river when the sun's coming out, once you've dragged yourself out of bed – a bit of mist and a few swans on there, it can be quite nice.

So then it's off to Cambridge for an MPhil. How did that come about?

So, like I said, I was interested in the application of computers to some of these problems and I looked at vision in my undergraduate course, and specifically medical

imagining, but then I thought, well what about speech, how does that work, how do you get computers to understand what people are saying, and so that was the main thing, just curiosity, and Cambridge happened to have a course focused on speech and language processing which was almost a direct continuation of my studies.

Is that the engineering department or physics department?

The course that I did at Cambridge was jointly run between the computer department at Cambridge and the engineering department.

Alright, ok. I see that at that time you also did some work with Toshiba.

Yes.

How did that come about?

So Toshiba have a history of a relationship with the engineering department here in Cambridge and so, firstly, they were looking for interns and when I was doing the MPhil course, it was a year-long course and I knew I wanted to go on and do a PhD and do some further studying and so, that gap between the MPhil and the PhD, I had taken an internship at Toshiba and then they actually funded me for part of my PhD as well.

So is it relatively easy to get funding for PhDs, or was it then, shall we say?

I think it is not easy to get funding for PhDs, and there are lots of different sources but I was funded by the EPSIC mostly, so I had one of the EPSIC grants and, like I said, Toshiba topped it up so I ended up with more. And I think it's got more competitive recently, especially around AI and machine learning which are very popular subjects right now.

Right, so to get funding these days you've got to be doing one of a relatively small number of topics within the whole thing, have you, have you got to be doing something very specialist to get funding?

I think I'm a little out of the loop on PhD funding now [laughs], so I'm not sure I'm the best person to answer.

00:16:34

Yes, my daughter has a PhD and it took her some effort to get funding, I have to say, so I don't think it's got any easier, that's for sure.

You have to have a good idea of what you're doing and where the funding sources are, and have help I think.

So, I wouldn't naturally associate Toshiba with things like speech recognition and language work, but is that actually a large part of what they do, or a significant part of what they do?

Toshiba has a research lab here in Cambridge which is the group that had this relationship with the engineering department in Oxford, and they do a number of different things. Toshiba is mainly, sort of here, for TVs I guess, but actually in Japan they have a much larger business. They do a lot more things and speech and language is definitely one of the things that they're actively looking into.

Interesting. Ok, so then onto the PhD, how did you come to choose your topic?

So my PhD topic. Actually I had a different topic in my first year, I was looking into something very specific and mathematical in my first year to do with speech and language processing, but my results were not that great. The topic I was doing didn't seem to be working out very well and it ended up being very computationally intensive to do any experiments on it so it was still theoretical and we didn't see it going anywhere over three years, and so my adviser suggested changing topics at the end of the first year.

That must have been quite disappointing, was it, after a year's hard effort which I guess must happen in research, that's the nature of it, isn't it?

That's the nature of research, you try something and you see if it works and sometimes it doesn't, despite your best intentions.

So best to switch topics while it's not too...

Yes and, of course, a PhD is not about proving that you've built something that works really well, it's about advancing the state of knowledge, and so if you have results which aren't working it's still valid research.

That's right, it's still a result in a way, yes. So I found your thesis on the Web, I confess I didn't read it all, I saw the title, 'Generation and Combination of Complementary Systems for Automatic Speech Recognition', is that right?

Yes.

Perhaps, just as a good point to dive off sideways, can we explain to listeners what speech recognition is. Is it maths, is it physics, is it engineering? Let's talk a little bit about how it works.

Ok, so speech recognition as a task is quite well-defined. You've got audio and you want to guess the words that were spoken in that audio, so it's a mapping from audio to words. Then you have the method in the middle which processes the audio and works out what those words are. There are three different parts to a speech recognition system. So the first one is the acoustic model. The acoustic model is a model which will take small fragments of audio, so maybe 25 milliseconds, very short audio segments, and it will make a good guess as to what sound is being spoken, or what phoneme the person is saying so you can think of phonemes as being the building block of words. So things like 'a', 'k', those are all phonemes. So we have this acoustic model which is telling you at each point in time in the audio what the most likely sound is that the person is making. The second part we have of the speech recognition system is the lexicon, so the lexicon tells you how words are constructed from phonemes. So, for example, that will tell you that the word 'cat' is 'c' 'a' 't', or the word 'read' is 'r' 'ee' 'd'. And then, if you have multiple pronunciations of the

same word, like ‘read’ and ‘read’ [pronounced ‘red’], you’ll have those in the lexicon. The third part of the automatic speech recognition system is where you’ve got the language model, and that predicts what sequences of words are likely. So, if you have a sequence of words which begins, ‘Hi, my name is’ we can predict that the next word is likely to be ‘Catherine’ but it’s probably not likely to be ‘goodbye’. So there are certain sequences of words which are more likely than others. And so we put these together: so we take the acoustic model to go from audio to sound, the lexicon is from sound to words and the language model is from individual words to sequences of words. So putting them together we can build up a model, a probablistic model of what the likely words are in that particular piece of audio.

00:21:14

Ok, and how does the learning bit fit into that?

So we build speech recognition systems from core data. So we might train the acoustic model, for example, on maybe – it depends on the size of the system that you’re training – but anywhere from fifty to thousands of hours of audio you might use, to learn the parameters of that model. And the same for the language model, that will rely on text so we might have thousands or millions of words of text and we can learn from those all sorts of patterns that are spoken.

Interesting. So I saw in your thesis, there’s quite a lot of maths, there’s a fair degree of physics I thought, I guess that’s the nature of it, that’s part of the attraction maybe, it’s the combination of things.

It’s a combination, yeah, it lies between computing and engineering and maths, it’s the part in the middle that covers a lot of different areas.

Ok, so back to the PhD then, at this point were you perhaps thinking of an academic career?

At one point I was and then I thought that, actually, my inclination is not towards necessarily research on its own but more applied research, so taking what we know about how things work and using it to build things in the world.

Did you supervise students while you were doing your PhD because I know I was supervised by a PhD student?

I did a few maths supervisions for some of the undergraduates at one point and, of course, supervising practical courses. So engineering students do a lot of practical courses as part of their studies and so supervising those as well.

How did you find that?

I enjoyed it. It was nice to ... it helps you to understand things sometimes, when you're explaining it to other people, because you have to think of different ways because, often, the way I understand something, if I explain it to somebody it's not the way they understand it. And so it forces you to be more creative about the way that you explain things and cement them more in your mind. And it's also good to see the undergraduates themselves learn and understand after you've explained things.

So, round about this time, you published the first of what turned out to be quite a long series of papers at various conferences and so on, is paper-writing and publishing a really important part of learning, do you think? I know one does these things for conferences, some of them for marketing purposes, but is it also really to improve your own understanding of the subject?

I always think that writing things down is a good way to improve your understanding of things, because you may have ideas in your head but I think until you sit down and actually put them on paper in words that other people can understand, sometimes there are things in that thought that aren't clear enough, that when you try and put it on paper you realise that there's a gap in your understanding of something, and so I think writing is very useful for that kind of purpose and also you can put things down on paper and start to see the links between different things more easily. So a big part of paper writing is literature review, so understanding what people have already done

and writing down how that links with your current work and writing down the motivation for what it is you're trying to do. So I think all of these things are important for cementing what it is that you're doing.

So we're up to about 2008 now, and you've finished your PhD.

Yep.

So you're sitting there, one month from the end, wondering what to do next, or did you have a plan? [laughs]

I had already taken a job offer with Toshiba at that time. As I said, they funded me for some of my PhD and I'd done an internship with them, and they had an office in Cambridge at the research lab. so I decided to go and work there. At the time, so this was 2007/2008, there was the financial crash so there was a freeze on hiring in a lot of places and AI machine learning didn't have quite the image that it has now, so there were not that many companies actually building machine learning and Toshiba was one of the ones who kept their research going and were very interested in building on it.

So what sort of work did you do?

So that was ... I carried on doing research, some applied research into speech recognition systems there. I did a bunch of different things, one of my favourite projects there was when we built our ... we built something which is a microphone array but sits on there like a table, and we used it to record our meetings. So we'd have group meetings and we would record them. And then we got our own speech recognition system which transcribed those meetings.

Instant minutes!

Yes, so we were going towards this idea of indexing meetings so that you could look back on them and see what was done and see what was agreed and see how the meetings went, have records. And also it was sort of a platform for testing out our

research so much more than taking .... there were lots of data sets that we do speech recognition and research on which are created and built by different agencies, different universities, and they are very fixed, and this was the first time I worked on real-life data sets that was being created as we were using it.

00:26:35

Did that project ever come to be part of a product, or was it just pure research?

No, it was pure research, it was our internal research platform and I think it was quite ahead of its time because it was – this was what, eight or nine years ago - so it was before the current wave of voice products.

Ok, so several years at Toshiba, four years was it, something like that?

Yep.

And then back to Cambridge University.

Exactly.

What happened then.

So after working on research on speech recognition for a few years I thought it would be great to expand the topics I was working on and the position came up at the university to move to a slightly different group that was working on dialogue systems. So that's the system that will have a dialogue between a human and a computer. In order to do that you need to do speech recognition to understand what the person said, but then you need to interpret it and then understand what the computer should say back to the person and keep that iteration going for a few steps so that you've got an actual conversation going back and forth between people. And so it was a great chance to learn how that technology worked and to sort of link it up with the speech recognition.

So was that a fixed term couple of years?

Yes.

Right, just one project really was it, was that why you went?

So I worked on a couple of different projects but sort of post-doc positions tends to be fixed term, yes.

Ok, and then in 2014, Amazon.

Yes.

Was that an advertised post, or did someone poach you, how did it come about? [laughs]

So this actually, I think, demonstrates the importance of having a good network of people you know. There was a particular conference, a speech related conference that happened in the summer, which I didn't go to but a friend of mine had gone to, and he'd been talking to some of the people at Amazon and some other places because he was also finishing his PhD and looking for jobs, and he came back and he said, 'You know, Amazon told me that they might be opening an office in Cambridge and I think you should talk to them.' And so I did, I got in touch, and I also knew somebody who was working there in the US, so I got in touch and they arranged interviews for me.

So it's actually quite a small world of specialists, or at that time was it quite a small world of people, so most people knew everybody in the field, did they?

I think, yeah, there is a small world, and I think definitely a lot of people have strong networks in these fields. And speech and language is still – it's growing fast – but it's still a relatively small field.

So what's Amazon like to work for?

I enjoyed working at Amazon. I think it had a few great things. So one is the chance to work on products that people are using, that lots of people are using. There are very few places that you can go in the world and work on a product that lots of people are using. One of the products I worked on was the voice box for the Fire TV. The Fire TV is a box that you plug into your TV that streams video from the internet and lots of people have one, so the chance to build systems which help all these people is something you don't get in many other places.

That's interesting, I hadn't thought of that, so there's got to be a more disciplined approach to releasing a product that's going to be used by millions of consumers, I guess.

Exactly, you can't break what you've already built and released, you have to be very careful about building and releasing, but you have to keep updating and innovating because you've got to build the next generation of what people are going to use, but also not break what they've already got, and fix all the issues that people have. There are things that don't work that should work and you fix them as well, so it's a careful balance of all of those things to keep products.

00:30:35

I think the impression most people have of Amazon would be that it's obviously a very large company, very marketing-oriented and very business-focused, but this must be proper engineering.

Yeah, Amazon has a large number of engineers working around the world on various different products. It's quite an engineering focused company, at least to me, I was working on the engineering side so to me it was very engineering focused.

So you said this was a start-up office in Cambridge, I think, so how many people did it grow to while you were there?

The office was created when a Cambridge company called Evi was acquired by Amazon and I was the first person to join the office who wasn't part of that Evi team.

So the Evi team did the factual question and answering behind Alexa and I joined to work on the speech recognition side to start with. The office grew quite a lot over the time I was there. So I was there for four years and when I joined there were about thirty people in the office and now they have an office in Cambridge which, I think, has enough room for 400 people.

So, at a guess, how many people do you think worked on the Alexa, do you know, hundreds or two?

So I think the latest figure I saw was just before Christmas, I think, there was an article about how there's maybe 10,000 people working on Alexa. Yes, Alexa is a large project and, if you think, you're not just talking about the core speech and language technology but everything Alexa does and responses. So there is a team, for example, looking after all of the different interactions that you can do, so if you're shopping there are various people thinking about shopping, how it should work on shopping, and there's marketing people and there's project managers, and it's very important to tie all these different things together. So it's not just technical roles like Alexa, but going from there there's lots and lots of different roles and lots of people involved in different ways.

I'm amazed that it's that number, I mean, I thought you might say hundreds, you know, 500 or something, but not thousands. It's amazing that anything gets done, to be honest. [laughs] So did you enjoy the time at Amazon?

I did, yes.

And what did you do, what roles did you do?

I joined Amazon as a scientist, so I was a speech recognition scientist and I was working on the speech products. I joined in the summer of 2014, just before Alexa launched, so after Alexa launched I started to work on Alexa and was working on how we take Alexa from the Echo to other devices, like the Fire TV. Then, over time, I started to work more on the language understanding side of Alexa, so if you take a system that's going to talk with somebody, like Alexa or Siri or Cortana or Google,

you have a number of different technologies in there. So you start off with speech recognition and you talk to Alexa and it understands what you've said, and then it has to work out from what you've said what you actually wanted, so it has to understand the intent behind what you requested because there are obviously many, many different ways to ask for the same thing. So that's the language understanding. And then there's the understanding and answering side as well, so if you've got a factual question that you want the answer to you have to be able to take a lot of different knowledge sources and find an answer to a question. And then you have the voice of Alexa. So I worked, when I first joined, on that speech recognition side, the input to Alexa, and over time I moved more towards the language understanding side, taking the output of the speech recognition system and decoding the intent and what people are asking for.

00:34:25

Did you become a manager?

I did, yes, so I grew the team in Cambridge while I was there.

Did you have any formal training actually, for that management position?

I did yes, Amazon has a few different training courses.

Did you enjoy managing people?

I did, yes, I think it's nice to see people ... people grow as they start to learn things and to be able to deliver stuff themselves and to get more confident, and that's one of the nice things.

So if you look back on your time at Amazon, are you particularly proud of anything or do you just enjoy being part of this whole big enterprise?

I think it was great to be part of Alexa at the time it launched because I think it was a much bigger success than people realised it was going to be before it launched. And just seeing how it grew and being able to contribute to that was great.

What do you see the future of Alexa being and products like it – I'm not asking for insider information here, just generally! [laughs]

I have no insider information on Alexa, but if you look at the general market we have a lot of systems out there now, like I said, there's Siri and Alexa and Cortana, Google Home. They allow you to talk and to do some commander control kind of things so you can say to ... you can ask Alexa to do things or ask one of these other devices to do things, and they will by and large interpret what you want and act on it, which is one of the first steps. But there are a lot of other things that they could do so they're not so good yet at understanding context, and I know that's something that a lot of people are working on. So if you say something like, 'play Frozen', now Frozen is both a song and a film and so understanding the context in which somebody is asking that and being able to do the right thing is, I think, a big topic, and that's a relatively simple example but there are plenty of other examples where what you say depends largely on where you're saying it and the history of what you've said. So you've got the physical environment that you're in, so if I'm talking to the TV I'm much more likely to want the film than I am to want the song and vice versa if I'm talking to a speaker. But there's also personal knowledge of mine, so if I'm asking for a particular thing and I've asked for it many times before, sort of learning from my personal preferences is another context and also the context of the dialogue that it's happening in. So if you've said something to the computer three steps ago in a dialogue, and then you refer back to that, we're not so good at tracking those sorts of contexts yet. So I think being much more aware of preferences and context is one thing that's going to happen in the near future.

00:37:17

A different question slightly, that leads to some interesting thoughts about privacy, doesn't it?

Yes.

You know, some people really like the way that Google looks at their emails and gives them adverts recommending what they wrote but other people find it deeply intrusive and offensive. It's probably more of a challenge with something like this, isn't it?

Yes, and that's where I think we're starting to see companies changing the way they're doing things. So some of this is driven by law, GDPR for example. In Europe that gives consumers particular rights over their data. One of the things is transparency and 'opt in' so by and large you have to provide enough information so that consumers can opt in to how you are using data about them. The companies themselves are thinking about this, even before, without these laws, about how to be transparent because in order to be used by people, and you want them to be useful for the people you're building, but in order for that to happen the consumers have to understand about the system and so, therefore, you have to get them to trust you, to trust that you're going to do the right things with the data. You have to be transparent with people about what you're doing.

Yes, interesting. Ok, so you had, what, four years at Amazon give or take?

Yep.

Up until back in the last year I think, wasn't it, so why did you decide to leave?

Working on something like Alexa, which is a great product that a lot of people use, it made me realise that there's a lot of companies who would like to make use of voice and language technology but it's difficult because, firstly, the sudden increase in the amount of people trying to use AI means that there are not enough people in the world yet who know much about the technology, enough about the technology to build systems. And that's changing, but right now there is a lack of people to do this, and I think there are just a lot more different applications out there which people could be looking into if they had the right people working on it and so I'd really like to help companies and other businesses who have applications for voice and language technology to build them.

So you're with Cobalt Speech now.

That's right, yes.

Is that a relatively young company?

Cobalt Speech has been around for about four and a half years. It was founded by a guy called Jeff Adams who is in Boston in the US and he has worked for many years in speech and language technology, he worked with some of the big companies, Nuance and Amazon and as a small start-up, and he founded Cobalt for the same reason that I just said, about helping companies who want to build speech and language recognition but don't have the way to do it yet.

So is that part of your network working again or is it a coincidence that he was at Amazon as well?

I think ... I haven't applied outright for a job for a very long time, I think networks are really important when you're working in specialist subjects so yes, it was a network.

So what sort of work are you doing now?

I'm working directly with companies who have ideas for how they want to use speech and language technology and helping them convert their ideas into the actual technology.

I see from the website that Cobalt is an entirely virtual company.

Yep.

Is that a big change for you?

It is, yes, I'm working from home now.

Good and bad?

[laughs] It is good, yes.

Yeah, I've started working from home and I've found it was great in terms of getting lots of work done but I did kind of miss the social side of it a bit, actually.

Yes, so I have to force myself to arrange lunch and coffee with people I know so that I have reasons to get out of the house.

So you spend a lot of time on audio conferences and video conferences, do you?

Yes.

Do you get out and about visiting clients?

I do, yes.

Do you enjoy that?

I think it's good to go and talk to people in their offices and to see how they ... hear their ideas.

And have you carried out doing some research stuff in parallel, is research in some way part of the job as well?

No, I haven't been doing any research for a little while, I'm focusing on taking what's been done in the research world and applying it.

Do you think you might go back to pure research at some point?

I don't know. I think there's some good things about research, that you get a lot of freedom to look into interesting things and to try interesting techniques and to draw

up your own agenda. That's very tempting for a lot of reasons but right now I think I'm happy where I am.

That's good. It seems to me that there are some quite exciting possibilities for applications of speech recognition that are probably only dabbled into but I read a reference about how there's a theory that it's possible that we may be able to detect the onset of dementia if Google speech patterns change. I mean, do you think that socially significant things like that are possible?

Yes, so there are definitely people looking into how you can use voice to detect medical conditions because some of them do affect your voice in different ways, and if these diseases are coming on slowly, coming on so slowly that people who are being affected either don't notice or the people around them aren't noticing because they're not seeing rapid changes, they're just seeing slight changes over a long period of time, that technology can start to detect things like slurring in voice and some of the other ... it's not my area of expertise so I don't know too much about how it's done ... but definitely things that affect your voice and the way you speak and the slurring and the words that you use and things like that can be picked up more by computers over a longer period of time.

That's exciting, I think. So do you think with a company like Cobalt, which I'd probably characterise as a small to medium tech company, do you think there are opportunties there for newly qualified inexperienced people or is it more for people who have been around a bit? Because I know it's pretty interesting, I think, for students to think of working for a small start-up directly but opportunities might be limited.

I think there are a lot of start-ups and a lot of people are trying to build start-ups in not just voice and language but all AI related technologies, I think a lot of people are very interested in working in a small company, and I think you shouldn't necessarily be put off because of it being a small company and you having no experience, but I think you have to look at yourself and what it is that you want from your career because, if you are not the kind of person who copes very well in a chaotic kind of environment and you need more guidance to get started then maybe a start up is not the right place. But

if you're happy with that level of ambiguity and that level of chaos then it can be a great place. And if you need more guidance to get started and sow a career path then the big companies ... I don't think I would have been great working in a start-up when I was just starting out for the reason that I think I needed a bit more time to get started, but that's not to say that there aren't people out there who shouldn't.

So, changing the subject just very slightly, I worked in the finance sector most of my career and I was used to an almost 50:50 mix of men and women, I think. I wasn't really aware of ... perhaps because I'm a man ... some of the issues [laughs] but I know that AIT is quite interested in whether there's more that could be done and whether AIT can help in some ways to encourage more girls and women to get into STEM. Do you have a view about this as a topic?

Yes, I think, I mean, it's obvious and it's a big conversation right now in the media, that there is a lack of women, not just in technology but in many professions, but if you think about technology the numbers are quite low. The number of women working in technology is low and those numbers go down the higher up you go, so the number of senior women working in technology is very low. So I think there is definitely work to be done to change that and it's twofold, I think. So, encouraging girls and women to consider taking technology courses and going into STEM fields is important, I think, so to work with schools, to work with parents, to show what these careers can do is a good thing, but also I think companies have to look at how they treat them and how they retain their women, such that they can advance to senior positions in organisations, because without both of those things, things aren't going to change.

It's also how technology's kind of gone backwards because I do a little bit of part-time lecturing and one of my slides about the history of technology shows a computer room in about 1975, I think, or 1980, and three-quarters of the people there are women.

Yep.

But it's flipped the other way now, so what's happened to our industry now, what is it?

It has flipped the other way and I can't say I know exactly – there's hypotheses about, computers became very marketed at boys in the eighties, I've heard that said, and therefore encouraged women out and into other subjects but I'm not sure if that's a cause or a symptom of what was happening.

It's a big complicated discussion, isn't it.

It is a big complicated discussion and I think you're right that the numbers have changed a lot and I think they've changed as the nature of technology has changed as well. So it was a very different job in the past.

00:47:29

Yes, that's true. Ok, let's change subjects slightly if we could. I think we've walked through your career so far. Let's talk a bit about AI, machine learning, robotics, that sort of thing, as I think it's not very clear to many people, me included really, it's a kind of taxonomy in this area. Let's talk a bit about how you define the boundaries of these topics.

Sure. So, AI is probably the one which most people hear in the media now, as it's become a very popular term in the past couple of years, but three, four, five years ago it wasn't anywhere near as popular in the media, but I like to think of AI as being the field of trying to make computers do things which seem smart. So whether that's a smart speaker, like Alexa or Siri, or whether it's a robot that can move around the world or whether it's a system that can diagnose from medical images, all of those kind of things I think, and AI encompasses the whole host of different things. So it's not just how you do it but also it might include the engineering around that, so how you build the systems and the design of the system and the user interface. So I think it covers more than just the core technology, the way you do some of these things. Then, if you think about machine learning, machine learning is one of the ... it's a group of algorithms really, which allow you to learn from data so they allow you to

learn things like identifying what's in a picture or identifying what's in speech. They're a set of algorithms where you have a set of data and you learn from that set of data how to do something. So machine learning has come on a long way in the past few years and it's one of the drivers that has made artificial intelligence systems more capable. Then speech recognition, which is speech and language which I've worked on, heavily rely on machine learning algorithms to do the core speech recognition and the core language understanding,and the same with medical diagnosis from X-rays, you can learn from X-rays, you can learn from a large database of X-rays how to diagnose particular diseases and, of course, there's still a lot of work involved in that and it's not something you can just stick a database of X-rays and a machine learning algorithm and plug them together and make it work: you have to understand something about the problem and the data you've got and understand it works to be able to build. So I think of AI as being this umbrella and machine learning being a key set of algorithms but underneath that umbrella.

Do you think we're still at the innovation stage in these core technologies or are we moving towards exploitation now?

So there's a few different approaches to machine learning which people talk about and there's one in particular called supervised learning which is where you take a database of images and those images might be labelled and from that label data you can predict. If you get a new image it gets categorised, which of those labels it falls into. Supervised learning is well proven in research and, I think, is fuelling a lot of the exploitation in industry right now. So people are taking supervised learning and applying it to lots of different tasks. And there are things that fall under supervised learning, like speech recognition, image categorisation, machine translation, language understanding, these sorts of things, so there is still research going on in them but we're definitely at the stage of exploiting them in industry. And then there are other forms of machine learning. So there's something called reinforcement learning, and there's unsupervised learning as well, that people talk about which haven't yet been exploited as much and applied as products so these ones are still very much in research stage.

00:51:40

So if someone's looking to start out in a career in technology now, do you think this is a really vibrant, thriving area to try to get into?

I think so, yes, there's a lot of interesting things going on here in machine learning and artificial intelligence right now, and I think it's a great area to look at because it spans a bunch of different things. It spans maths, it spans computing, it spans user interface design, it spans other adjacent – project management – a whole bunch of capabilities which means that you don't have to be somebody who is going to dive right into the algorithms to be able to work on the technology. You might have a preference for interface design, or for software engineering, and they are also things that you need to be able to build these and so there's a large choice of different career paths you can take now.

So someone who is doing A levels, or Btecs now perhaps, thinking of maybe going to uni, they wouldn't have to specialise necessarily in a machine learning topic, they could be generalised computing or engineering or physics degrees and still be appropriately skilled for it.

Yeah, and I think there's a lot more courses now, ... there was online education happening a lot, things like Coursera have a lot of machine learning courses and lots of videos on YouTube and so you can actually learn a lot about machine learning independently of university, but I think that the technical skills you learn from something like physics and maths and computing at university are very useful, but equally people come from different angles too, so somebody might come from a linguistics background or psychology background and gradually, over time, move in different ways into the field.

You said earlier, about the way Amazon's developing things, it means there's roles for people like social scientists too, aren't there, psychologists and ...

Yeah, and I think there's lots of interesting ways you can get into working on machine learning and artificial intelligence and some of the things that you bring from those

other fields are very useful and complementary to some of the other more technical physics and maths subjects.

Do you happen to know if there's a lot going on in the, kind of, defence world? I know a lot of students have talked about being quite attracted to that area of being involved with the military for all sorts of reasons. I'm guessing at this point, do you know?

I don't know much about the defence world.

Ok, so then moving on a little bit to voice technology, as you know I come from a voice infrastructure background, I remember when voice-over IP suddenly came on, actually it didn't suddenly come on, it seemed to take a long time to get on this, but it obviously revolutionised the way we built buildings and structured telephone systems. It seems to me that world has gone a bit quiet in the last ten years or so. Where do you think the future of voice technology is?

I think voice technology is becoming much more easy for people to access. So most people have a smartphone, I have for example, which has much more computing power than anybody had access to ten or twenty years ago. Certainly, when I was studying the idea of having this phone, this computer in your pocket was not really there. And so I think that building networks for connecting computers and connecting things to the Internet is changing the way people are accessing this technology, so you can, for example, use your phone and there are some voice technology you can run directly on your phone. The other thing that's happened, then, is that once you've got this infrastructure for building the internet you have things like the Amazon web services where you can build clusters of computers which are running from the Cloud and people will run speech recognition in the Cloud, so it means it's much more powerful and then everybody can access the same computing power from their phone which enables you to have one location where you are building and applying your speech recognition system but making it accessible to a lot more people. So these changes in connectivity and computing power are really changing the way that people can access voice and language technology.

I think people use the telephone less than they used to, actually, certainly in the corporate world, I think it's much more about messaging, it's much more about video and audio conferencing rather than just picking the phone up on a one-to-one conversation.

Yes, that's true.

So what are people like Microsoft and Apple doing with the voice space, do you know? Do you keep track of them?

I don't keep too close track, no.

[laughs] I was just curious. So, it seems that we have the capability to think about extracting value from the voice stream somehow, that's what we've been talking about, essentially that but applied to consumer products, if you like.

Yep.

In the commercial, corporate world, can you seen opportunities there for extracting value from the voice stream?

So I think there's a lot of ... a lot of things are still carried out by voice, even though people aren't phoning each other up anymore, you still do have plenty of phone calls and plenty of audio and video conferences where things are being discussed, and using, recording those in corporate settings and being able to identify, for example, in a meeting what actions have been decided on, what information was shared, that can be valuable for lots of people. And there are, of course, industries where you have to do that so where you're being audited for compliance with laws you may have to record your conversations and be able to look at what happened in them, so being able to do that much more efficiently, analysing voice rather than having to have somebody manually transcribe is one area.

Do you think there's any opportunity for multi-language type stuff? I'm thinking back to the days when you spent audio conferences with people in different continents and language was occasionally a barrier.

Yes. You are starting to see ... people are building machine translation systems which will take one language and ... you can speak in one language and it will translate it and then speak out in the other language so you can have a conversation like that with some people, and I have seen those systems start to be used in various places, and embedding it in a corporate setting where you actually have people who speak different languages. I had a meeting the other day with somebody who spoke French and not great English and so we had this slightly difficult conversation with somebody else trying to translate bits of it, and you can see the applications of voice technology there.

It's interesting, actually, because it would go some way towards reversing the kind of flattening out we're seeing at the moment, that every country is starting to look a bit like every other country sometimes, but being able to speak in your own language and to use your own terms is a way of reinforcing your culture identity, isn't it, and retaining that. That might be quite nice in a way. [laughs]

I think it would be nice, especially when you want to talk to people who have a different use of language to you and you want to be able to have a conversation with them because they're doing interesting things and you want to understand more about each other.

00:59:29

Ok, so if we look, get your crystal ball out, and look kind of five, ten years ahead, have you got any thoughts about what we might be using – phones or voice communication or ... how are we going to be talking to each other, do you think?

I think people will still talk to each other .. you mentioned that there are a lot of people using different ways of talking to each other now, so send messages, messaging systems as well as voice and whatever, not picking the phone up as much

to each other, and I think that will still continue, I think people like the convenience of sending text based messages to people and then replying at later dates. But most human relationships are formed on actually sitting down and talking to people and so there still will be a big place for sitting down and actually talking to each other.

Yes, I agree with you, absolutely agree with you, I mean, I suppose round about 2,000 say, plus or minus around five years, at that time when networking and the corporate environment became easy to do and everyone networked, there's the temptation to hide behind email and I was a project manager and I used to get up and walk around and talk to people and go out and look at things. Otherwise you're stuck in your own little world and speech is so rich, actually.

Yes, and you get more than just the words from speech.

You look in the eye, the body language...

Body language you pick up, you pick up intonation from what people are saying, you pick up nuances from the way that they say things and you understand how they're thinking things through when they're talking it through because .. and reading into something somebody has said on the screen is so easily misinterpreted and you can have ... I don't know if you've been in long email discussions where people are back and forth in emails because things are misunderstood that could have been cleared up with a five minute conversation.

That's right, you end up emailing about something completely different that's nothing to do with what the original topic was at all.

And so, yeah, there is still definitely a place for voice.

Good. Alright Catherine, thank you ever so much, it's been lovely talking to you. Is there anything else you would like to comment on about what's going on in your field or are we done?

No, I think we've covered loads of stuff and it's been great to talk with you.

Well, thank you ever so much again and on behalf of Archives of IT, thank you for giving us so much time. I'm sure that for students, or for younger people, thinking about making a career in this area, this is the sort of stuff they ought to know about, really, and understand a bit more about what they might be getting into. Thank you very much.

Thank you.

[recording ends at 1:02:09]