

Please note that this transcript has been lightly edited relative to the original audio, in order to improve readability.

# AICSCSHistory\_Interview004\_AndrewBlake

## SPEAKERS

Vassilis Galanos, Andrew Blake

### **Vassilis Galanos** 00:03

So this is the 24th of April 2024. My name is Vassilis Galanos and I'm here with Professor Andrew Blake. This is part of our history of AI, computer science and cognitive science at the University of Edinburgh project that I co-research with Professor Chris Williams and Ms Xiao Yang. So we are really pleased to have Andrew Blake with us today. So Andrew, before we delve into the specifics of your notable career and the evolution of AI and computer science in Edinburgh, could you please share an introduction about yourself highlighting your major contributions to the field, and your personal journey within AI and computer science with an emphasis on Edinburgh?

### **Andrew Blake** 00:58

Of course, well, I did my doctorate at Edinburgh. I had never really intended to do research, it was kind of an accident. So, after my undergrad, I went to MIT and did some electronics. And I came back to the UK and I was working in Edinburgh with a company. A company called Ferranti that only old people will remember. It became GEC and then it became Marconi and then it went bust. So originally, the project was all about electronics and distance measurement. And it had a military slant to it. But then we started using cameras for things, measuring speed, you put a camera on a moving helicopter, and it was looking down at the ground. And just the way the ground was kind of passing in front of the camera was a way of allowing the helicopter to measure speed. And for various reasons this was a good way to do things. But then the project got canned. And so I thought, oh well, actually, in that case, I might as well move to the university, and then I can do what I want, instead of having some boss running out of money in the middle of the project. So then I looked around to see where I could go. And at the time, there was the guy in Swansea, who did an inspection of manufactured objects, and I don't think I can quite remember his name now. But there was another guy called Professor Duff, who was working in UCL in London, and he was building a new kind of computer especially to do image processing. And I really wanted to go and work with him, I liked that idea. But he didn't have a place, and so I carried on looking and I went to the library and realized that Edinburgh was a big centre for artificial intelligence, or maybe had been a big centre, it wasn't quite clear. I didn't quite get the nuances at that time. But I could see Donald Michie has produced this whole series of machine intelligence books that are numbered, one through ten and then beyond. I can't remember even how far it got quite a long way I think. And so I thought, oh wow, artificial intelligence is happening right here in the city. So I went to see him. And as it happened, he was delighted to have me because he got a new project to do robotics, partly in collaboration with this group in London where I'd wanted to go. And he didn't have anyone to do vision. He was doing more learning. But he had this sort of nostalgia for the previous decade when he'd had a big lab, and, you know, all skills in the lab. And so now he had a little lab, that the University allowed him because he'd been a little bit in disgrace, you probably know all that history. He'd been in disgrace, and he was allowed this little unit called, at that point, it was called the Machine Intelligence Research Unit, in Hope Park Square, next to part of the AI Department, part of AI was down there and the other part was up at Forrest Hill. So he had his little unit. And he wanted to have someone doing vision

Please note that this transcript has been lightly edited relative to the original audio, in order to improve readability.

because he felt a robotics lab should have somebody doing vision. We didn't have a robot, but he wanted still to do vision. And from my point of view, I decided I wanted to do a PhD, but I'd had a salary. And I didn't want to go back on a grant. Because you know, by that time I had a house and a mortgage and all that. I couldn't really afford to just live off a student grant. And he gave me a research job. Then it turned out that the University by a quirk didn't charge me fees for the PhD, because they had this arrangement where if you were researcher, they called it 1A, if you were a 1A researcher, you didn't have to pay fees for your PhD. But how did you get to be a 1A researcher? Answer, you needed a PhD. So this was really a little bit of a catch 22. But Donald Michie managed to break the deadlock by saying, well, he has the equivalent experience of a PhD, because I'd had this extra three years in industry and it stuck. I did quite well out of that, and I also worked in his little consultancy company. So we did learning research for BP, for their plant up in Aberdeen. So they had, you know, the oil rigs in Aberdeen and lots and lots of data coming ashore from the oil rigs. They didn't really know what to do with the data, there was just so much of it. And we used logical induction. A chap called Ross Quinlan, an Australian researcher, I don't know if you've come across his name at all, but he did this quite famous bit of work that was called ID3. That was one of the very earliest pieces of code for doing learning. So we were using that learning code to kind of analyze all the data coming across from Aberdeen. And that was the consultancy. I was sort of set up for three years then to do a PhD. But at this point, I'll stop. In case this is not the kind of thing you want.

**Vassilis Galanos 06:28**

It is absolutely the thing we want, and you're more than welcome to share as many anecdotes if you like, but perhaps this is the right point to move to something like the next question. So you arrived at Edinburgh in 1977?

**Andrew Blake 06:50**

Well, it was 1980 at the time I arrived.

**Vassilis Galanos 06:53**

1980. So when you arrived, did the University of Edinburgh have a reputation for being significant hub for AI and computer science? You mentioned the old glories and so on that Donald Michie was probably referring to. So in your view, what factors made Edinburgh this kind of pivotal place for the development of AI during the early days? What kind of stories did you hear when you arrived? What did you find out? How did you experience that kind of momentum?

**Andrew Blake 07:33**

I think at the beginning, I was a bit naive, and I didn't quite know who was who. And I had this strange experience that my boss at Ferranti when he heard I was going to Edinburgh, he said, are you going to the AI department? And I said, no, it's not the AI department is the Machine Intelligence Research Unit. And he said, oh, if I were you, I would go to the AI department. That will be better. And I said, why? He said, oh, well, I don't really want to talk about it. He was this very old school guy, who was not very transparent. And then it was two years later that I bumped into his wife, who was a secretary in the AI department. It turned out, of course, she knew all about me, a lady called Jean Parker, who was a bit of a kind of friendly dragon, who presided over the AI department. I suppose she was Jim Howe's most senior administrator. And, you know, you did not cross Jean Parker, she was she was a force to be reckoned with, but you know, ultimately benevolent. So then I realized that this man, John Parker, who had been my boss in Ferranti knew all about the AI department and the Machine Intelligence Research Unit and the schism between the two, because it was a big schism there. So in the heyday of Donald

Please note that this transcript has been lightly edited relative to the original audio, in order to improve readability.

Michie and his setup in Edinburgh, I mean, he initiated AI in Edinburgh and he brought in Christopher Longuet-Higgins and, just lost his name, the Bristol professor of perception, AI and the Eye

**Vassilis Galanos** 09:25

Richard Gregory?

**Andrew Blake** 09:27

Yes, Richard Gregory yes. There's the three of them were setting up the AI unit and it all went very well and it was all you know, they did some very famous work. There's that wonderful poster, which probably is still up somewhere I don't know, of Pat Ambler and Barrow who I've met since, and Popplestone, of course, who I knew quite well. There was a whole band of them that became quite well known. See my memory is not quite as good as it was. You will know the guy Steven Salter, who continued to be a figure in the University, he's a mechanical engineer he went on to do Salter's duck, which was the machine for harvesting energy from the waves, something completely different. But that was what he went on to do, he had this enormous swimming pool built up at Kings Buildings to experiment with the duck. And so they've been very famous. 1975, I think, was about the peak. And there's this very famous paper, that begins Ambler, Barrow, Burstall, Popplestone. You know, these are all the authors on this multi authored paper and this was the Freddy paper. This was the paper that documented Freddy the robot. And this great big poster that used to be stuck up in Forrest Hill showed all of these people in their 20s, or possibly 30s, sitting around the Freddy robot. And nowadays, I think Edinburgh hosts a video of the Freddy robot somewhere on the departmental website, and it's speeded up 16 times, but otherwise, what it does assembling this wooden toy, it's just amazing for the time. I mean, if you showed the video running at that speed now to a conference, people wouldn't think it was out of place. So it was a fantastic piece of work, in its time. So that was the kind of highlight, but then there was this report. Well, then there was some scandal around Donald Michie, and I'm sure there are many versions of this story. But the version I heard is that Longuet-Higgins, who was a rather famous scientist, got pettish and jealous of Michie's success. Michie was really the more junior scientist, I mean, Longuet-Higgins was a Fellow of the Royal Society. He was even a Royal Society Research Professor, which is a very elevated kind of scientist. And as a chemist, which was his previous career, he become very famous. And there's a book by a man called Coulson, who really is the godfather of covalent bonding in chemistry, completely different subjects. And when you look in the back of his book, of course, there are many references to Coulson, because he wrote the book, and the only other person who has as many references as Coulson is Longuet-Higgins. So he was this brilliant scientist who got a professorship in his 20s in Cambridge, and carried on defining the field of covalent bonding in chemistry, a sort of combination of physics and chemistry, absolutely groundbreaking stuff, until he was 40. And then at 40, he decided, you know what, I've done as much as I can in chemistry, I'm going to do something new. And he wanted to look at the brain, and artificial intelligence and so on. So that was how he got into AI. But anyway, after a few years, so the story goes, this is what was handed down to me, by word of mouth. So it could be wrong. But this is the version I heard, is that he got very jealous of Michie, because Michie was getting all the attention. You know, Michie was a fantastic communicator, and publicist. And all the glory, the greater part of the glory for this triumvirate and their effort in AI, reflected on Michie. And Michie was also one of these guys with a lot of energy, not terribly careful. So in the end, Longuet-Higgins supposedly told tales on Michie that he'd been doing company work, consultancy using graduate students, and not paying them. That he'd been sort of transgressing those boundaries between public funding and private work, and using SRC, as it was called then, computers for private work. So SRC is a predecessor of EPSRC that we have now. So that was the story. And so then, it was said that he told the SRC about this, that the Michie was doing bad things. And the way they dealt with

Please note that this transcript has been lightly edited relative to the original audio, in order to improve readability.

that was to bring in a kind of very senior citizen, who was a mathematician whose name I will remember in a minute, or maybe just cheat and use the web.

**Vassilis Galanos 14:45**

Lighthill most likely.

**Andrew Blake 14:47**

Yes, thank you. Yes, it's Lighthill. So they brought Lighthill to kind of do a hatchet job on AI. I'm sure many people have told you about the Lighthill report and all of that, and there's this TV programme recorded debate between Lighthill and then Michie and all of this. Then we start the AI winter, which was somehow synchronized with the US AI winter. Even though the story there is very different, more to do with IBM that was big on AI for a while, and then they lost confidence in the whole AI programme. And so the US had their AI winter about the same time as we had ours. And then when I got to MIRU, shortly after 1980, maybe 81 or 82, they decided to reverse this AI winter by getting in a figure as big as Lighthill, which was Sir Peter Swinnerton-Dyer, magnificent name, and a mathematician, did you come across him? So he was the anti Lighthill. He was the antidote, I should say to Lighthill. So Lighthill started the AI winter, then they commissioned Sir Peter Swinnerton-Dyer to reverse it. And why did they do that? Because the Japanese had started something that they called the fourth generation of computing. And this was where the new computers would natively do logic. It never quite happened. Actually, the closest we got was, you know, there was the Lisp machine. Still not quite there, it wasn't Prolog. Then people talked about Prolog machines, but I don't think it ever quite happened. But that was the idea. And the Americans got very worried about it, that the Japanese are going to steal a lead. So they kind of geared up and started programmes. And then the British realized, oh well, we'd better do something too. But hang on, AI is in disgrace. How can we administer the antidote to the antidote? So when I was a student, Peter Swinnerton-Dyer was touring around the computer science departments of the nation, interviewing people. And he came to MIRU, and I remember sitting around in a circle with Peter Swinnerton-Dyer and telling him what we were doing. And in the end, he delivered the verdict that he had been commissioned to deliver, which is, after all, AI is important, we should invest in it. For me, this was fantastic timing, because then, by the time I'd finished my PhD in 1983, they'd started a new programme called the Alvey programme that you've probably heard about. And so just as I moved to computer science at that point, because I wanted lectureship, I wanted a proper job kind of thing. So I moved to Edinburgh Computer Science, not Edinburgh AI, because they wouldn't have had me, actually the person you should interview, if you can find him is a guy called John Hallam. Have you heard of him? Have you talked to him?

**Vassilis Galanos 17:50**

We just agreed on an interview end of May. So very good.

**Andrew Blake 17:55**

He's a sort of very old friend. I mean, we've kind of lost touch nowadays, we don't really communicate very much. But you know, he was a very old friend. And he was doing the same kind of research that I was, but he was doing it in what turned out to be the right place. And I was doing it in what turned out to be the wrong place, which was because computer science wasn't supposed to do computer vision. I mean, they hired me, because they thought my thesis was interesting, I guess, and I gave good the interview. And, you know, I beat one of their own to this post, a guy called Gordon Brebner, who eventually got a post himself, but not that one. And then John Hallam got a post in the AI department, and so when it came to sort of expanding this activity, the university was very clear. They said, no,

Please note that this transcript has been lightly edited relative to the original audio, in order to improve readability.

we're going to do this in the AI department, not in computer science. So there was a competition for a new lectureship in computer vision. And there were various bids for it. And the AI department came at the top of this bid, and computer science came at the very bottom. So that was a kind of very powerful signal for me, I thought, oh, okay, it's going to be difficult to make progress doing computer vision in the Computer Science department, even though it was a lovely department, very nice people. Lots of them that you would have come across, Robin Milner was one of them. I'm sure you've had lots to do with him. I liked him. He was very friendly. He even wanted to get me into the LFCS because he could see I was doing good work. And he was looking for a rationale for me to join LFCS. But we both concluded that what I did really didn't fit in LFCS. So that was a pity, because they were good people. And similarly, Plotkin was always extremely friendly. But we did such different things. There was quite a few people around, like Alistair Sinclair and Mark Jerrum, who were more my generation. And people like Hamish [added post interview: Dewar]. These guys who did computer systems. And in those days, you know, they built their own computers, there was something called the APM, was the computer the computer scientists built. And they built their own language and operating system, the language was called IMP. And everybody was learning to programme in IMP. It was completely a local product. But somehow you could do local operating systems and local languages back then, you know, you can't really do that anymore. Because it's just not mainstream enough. So I had a fun time there for three years. Towards the beginning of the three years, I got a thing called the Royal Society Research fellowship. So then I didn't have to teach anymore. So I think I taught this course on graphics twice. And that was it. After that I didn't do any more teaching. And I happily did my research in the computer science department, I had my grant, I bought some of the latest computers, the Sun computer was very new then. I bought some of those, they were very advanced with not quite touchscreens, but sort of the wysiwyg. That was all great. And I got a research assistant. Now this is very interesting, because this guy, this research assistant, who came to me in computer science, I hired him on my first grant, has become by far and away the most successful and famous computer scientist in the UK, by a huge margin. And Edinburgh people don't seem to know about it, which is extraordinary. So this is a guy called Andrew Zisserman. And if you look him up, he's in Oxford. He's been in Oxford since I moved to Oxford, which was 1987. We moved together. And you know, at the time, he was very little known, but he has become, by a huge margin, the most successful computer scientist in the country. He's a Fellow of the Royal Society, he has a Royal Society Research Professorship, he works halftime for Google DeepMind. This guy is just an absolute giant of the field. And honestly, Edinburgh University really ought to get him back up there and get an honorary degree. I mean, they really should recognize that this guy is a huge success. I think lots of them, they didn't even know because he was very junior when he was in the University. I mean, he was just working for me as a research assistant, he's a very quiet guy, he didn't put himself about. I don't know if he even gave so much as a seminar when he was there. But he just became enormously successful. So if there's anything you can do to help Edinburgh University embrace their unrealized success. So anyway, we were quite a good team at the time, and we wrote a rather beautiful book together. Actually mathematicians like this book, because it is quite a mathematical book about 3d computer vision. And then we moved off to Oxford because as I say, I'd had this signal from the University that computer vision in the Computer Science department was going nowhere, that they were not going to be any opportunities for me to grow my activity, they were not going to promote me, you know, nothing was going to happen. And at the same time, a certain Michael Brady was a professor in MIT, who came to Oxford and started this robotics lab in Oxford. So I was one of his first hires, as a young Lecturer in Oxford. And so that was the beginning of the rest of my career, kind of thing. And I didn't do that much with Edinburgh after that, although I still had friends up there. People like Andrew Fitzgibbon, I don't know if you've come across him? So he did his PhD in Edinburgh with Bob Fisher, who you probably have met. And, he's also been very successful and has come down

Please note that this transcript has been lightly edited relative to the original audio, in order to improve readability.

[added post interview: to Oxford, then Cambridge]. I eventually moved to Microsoft in Cambridge after 12 years in Oxford, I moved to Cambridge. Andrew Fitzgibbon he came to Oxford first, and then I hired him to Microsoft in Cambridge.

**Vassilis Galanos 24:43**

Very interesting. So all these notable figures, they all started from Edinburgh. I guess a bit of a follow up there is your everyday interactions with these people. What was it that made those minds flourish, like what some people say it's the sort of geography that the topology even of the place that, you know, everything was like very near each other. So, conversation sparked more easily. But in your view, what was it that made all these amazing people sort of experiment with these ideas and put them to place?

**Andrew Blake 25:27**

Yes, interesting question. I mean, the place did certainly attract some brilliant people. You know, Alan Bundy. I mean, I know more about the AI side than the computer science side, of course. And I guess your remit covers computer science as well.

**Vassilis Galanos 25:45**

Yeah, yeah. Although, I mean, my personal expertise is rather in the history of AI. So I'm also a bit biased towards it.

**Andrew Blake 26:01**

It's quite interesting, there was Burstall and Plotkin were both right in there doing the AI thing with the Freddy robot and all that, they were all core part of that team. And they completely kind of flipped to computer science. So I think they left AI behind, decided that. I think they both liked doing logic. And so they decided to do logic in the context of computer science, for things like, in the case of Plotkin, program semantics. So both disciplines talk, for example, about grammars. So if you wanted to understand language in those days, in the wake of Noam Chomsky from MIT, then you would have to understand about grammars. So those sorts of AI people were into grammars, but then so were computer programming language people, they were also into grammars, and then you know, grammars could explain the syntax of a language. But Plotkin wanted to understand the semantics of the language, what did it mean? You know, how could you translate precisely a computer statement in a programming language into its meaning in terms of, you know, how it would take effect inside the computer. And this is not an area that I understand very well. But he became this kind of pillar of the establishment doing that. And then at quite a young age, Milner and Plotkin both got elected to the Royal Society, which was very unusual then. There were really not very many computer scientists, and no AI people in the Royal Society at that time. Other than Longuet-Higgins, but then that was because he was a chemist. So he was in the Royal Society by virtue of having been a chemist, Michie never got elected to the Royal Society. Despite that, because he was sort of a very controversial figure. One book that was written about AI called him the Rasputin of artificial intelligence. I don't know if you've come across that quote?

**Vassilis Galanos 28:14**

No. Do you remember what book was that?

**Andrew Blake 28:18**

Yeah, that's a good question, if I can remember the name of the author of the book. [Added post interview: Maybe by Ed Feigenbaum?]. But this would have been a book written around the early

Please note that this transcript has been lightly edited relative to the original audio, in order to improve readability.

1980s, of the history of AI. So it wasn't easy for AI people and computer scientists to be accepted into the mainstream, of academic distinction. So nowadays, I'm a Fellow of the Royal Society. I've been a Fellow since 2005. And I worked very hard to get computer science recognized by the Society, and I've been quite successful. So nowadays, we have 70 or more computer scientists in the Royal Society as Fellows, and probably 10 or 15 of those do AI. So we've been much more successful, but back in the day, for Milner and Plotkin to get recognized in that way was an amazing breakthrough. The only other person who was already in the Royal Society was Tony Hoare. You might have heard of him. I don't know. He was a computer scientist in Oxford, who also won the Turing Award like Milner did. So you know, those are probably our two most famous computer scientists sort of still living in the 1990s I suppose. They'd been a few before that. So I don't know. Yes. Edinburgh did have pulling power. I'm not sure I completely can explain it. I guess, Michie and Longuet-Higgins and Gregory were very charismatic people, if you ever met any of them, each in their own way very charismatic. I mean, Michie was just an amazing storyteller. He could be talking about anything, the subject matter could be dry as sawdust, and he would somehow manage to weave a story where you were just gagging for the next page to be turned. He was just an amazing raconteur. And Gregory as well, he was much more kind of a lecture demonstrator. So he used to like giving psychological demonstrations of visual illusions, this kind of thing. And he would do it always, not just with a drawing on the page, but often with a piece of apparatus specially built for the purpose, with kind of a box with a coloured glass on the front and lights inside and so on. So used to give fantastic lectures. And then Longuet-Higgins was just an incredibly clever guy, he used to give lecture demonstrations with music. So one of his side interests was building computer programmes that would do interpretations of music. And he liked, particularly the work of J. S. Bach, because it was a little bit more systematic and susceptible to musical interpretation. So I got him to give lectures to one of my classes once and he came in with this kind of computer connected to a keyboard. And it did an interpretation of J. S. Bach's Italian concerto. I don't know if you happen to know that, but it did it in a sort of semi human way, with rubato, you know, where you kind of accelerate and decelerate in the right places and had a bit of phrasing and it was not entirely mechanical, it was something a little bit better than just a machine. You know, better than a Victrola, those player pianos they have where the paper tape is running in the piano, those performances are very mechanical. But this had an interpretation. And every time it played, the interpretation would be slightly different. So they were each fascinating people in their own way. And so I'm not surprised that they attracted very interesting people, although I wasn't around to sort of actually see how they did it. Because the other person that went through that mill was Geoff Hinton. And I mean, if Andrew Zisserman is the most famous computer scientist in the UK, Geoff Hinton is the most famous in the world. He really is the top dog, I would say. There are some very famous ones, and not everyone would agree, but for somebody who loves artificial intelligence, then it's a no brainer, Geoff Hinton is the top. You know, some other computer scientists might say, oh well, there's this guy here who does algorithms like Don Knuth or something. They will say, well, he's the most famous and, yeah, for sure, he's very famous, too. But I mean, you could defend the position, especially at the moment, that Hinton is the most famous computer scientist in the world. I don't know if you get to talk to him. He's probably a busy man.

### **Vassilis Galanos** 33:23

He seems to be quite busy. But we hope to get him in May as well. Yeah. Did you overlap at all? I'm not sure, based on what I know, from history, you shouldn't have overlapped, but maybe you did.

### **Andrew Blake** 33:39

No, he left just before I arrived. On my desk, in the MIRU, was a copy of his thesis sort of left behind, which actually was quite relevant and quite interesting, because in those days, the work he did was a

Please note that this transcript has been lightly edited relative to the original audio, in order to improve readability.

bit different, a bit more like what I was doing. So I enjoyed his thesis. He's a fascinating guy, and I guess, Edinburgh can take some credit. I think his inspiration was mostly Longuet-Higgins, who was his supervisor, and got him interested in neural networks, more or less. And, you know, at the time, neural networks were quite primitive, but Hebbian learning, I guess, comes from the 50s. And then there was the perceptron in the 60s. And then there was Minsky and Papert's famous book that killed off the perceptron. And then what interested Geoff was how easy it was to get past the proofs in the perceptrons book that neural networks didn't work. So the perceptron book was a beautifully written book, but its proofs were all about two layer networks, an input layer and an output layer. And Geoff, maybe along with others, realized that if you put an extra layer in, then none of the proofs in the Perceptrons book applied anymore. So the kind of depressing effect of those results, the spell was broken if you like. So after the Perceptrons book, everybody said, oh, neural networks are dead. You know, there's just no point in studying them because Minsky and Papert proved that they're no good. There's some very obvious things that you would want a network to do that they will never do. Well, those sorts of results are all very wonderful. But Geoff's contribution to begin with was to say, well put another layer in the network, and then the proofs are all irrelevant. And we're free, we've broken the bonds of Minsky and Papert, and with self-respect we can do research again.

**Vassilis Galanos 35:50**

That's the most accessible version of Hinton's contributions to the perception controversy, I think very, very interesting. Thank you very much for this. Rumour has it that Longuet-Higgins had just read the perceptions book before Hinton became his student. And that kind of resulted in their debates about the topic. Very interesting.

**Andrew Blake 36:16**

Yeah. And of course, David Wilshaw was in there, too. I guess he must be about the same age as Geoff, probably did his PhD about the same time, I'm a bit hazy on exactly the chronology. Were they research students with Longuet-Higgins at the same time?

**Vassilis Galanos 36:35**

Yes. Or with a few years difference? Maybe a couple of years? Don't quote me on that. Very interesting. So you've mentioned many, many important names, many important figures. Part of this project, we're also looking for less noted individuals whose work might have significantly contributed to the advancements of AI or computer science, but might not have received the widespread recognition they should for one reason or other. Do you think you can recall any person, an individual from that time?

**Andrew Blake 37:26**

Zisserman is the biggest one by a country mile, that as I said, it amazes me that people in Edinburgh don't realize that he comes from Edinburgh. But you know, it's probably not quite what you're looking for, because his big innovations have been done much more recently. And in that, he left in 87, and he got really famous in this millennium. So they can claim him as a famous son, but it wasn't really work done at Edinburgh. But if you think about people at Edinburgh, I guess there was Pat Ambler from that original paper Ambler, Barrow, Popplestone, that one. Yeah, she was very quiet, but an important part of that team, I think. And the language, POP-2 and POP-11, they were quite important. I enjoyed them. I used them for my PhD. And that was a sort of team effort. There was also Chris Brown. Yeah, I think he's on that paper too actually, Ambler, Barrow, Brown. C M Brown. And he's been a Professor in Rochester, New York for most of his career. He'd be an interesting person to talk, to he's a bit older than me. So he might be a little bit better on the older history. And he's sort of retired. I guess he is retired,



Please note that this transcript has been lightly edited relative to the original audio, in order to improve readability.

but I can probably put you in touch with him, if you wanted. He's a very chatty and fun individual. So that might be helpful. And he wrote a book, Ballard and Brown on computer vision, which was the standard book for quite a long time. When I was doing my PhD, it was already doing well. I didn't know who else is there from back then? I guess you might have caught up with Lincoln Wallen? He's an interesting individual. He was one of Alan Bundy's research students. And he's had a very interesting career since, turned out to be an important person. He moved to Oxford as well, maybe a bit after me, I think, and was a Lecturer there. But then since then, he he's now kind of CTO of a company called Improbable in London, which is kind of attracting as much attention almost as Deepmind, it's like, you know, almost in the same breath as talking about Deepmind. He's turned out to be a big figure. And there were others like Robert Kowalski, did you come across him?

**Vassilis Galanos 40:13**

I've come across the name but not. Yeah.

**Andrew Blake 40:18**

So he's a sort of contemporary of Alan Bundy's. Again, you know, I think what happened to some extent is when the AI Winter came in the late 70s, there was a bit of a sort of explosion and talent fled all around the Earth. So Harry Barrow was the another one that I mentioned to you just now. And he went to America as well and joined Fairchild labs, and did a very famous piece of research on something called intrinsic images. Barrow and Tenenbaum were the authors of intrinsic images. That was a very interesting idea in computer vision. And now Tenenbaum's son, actually is a very well known professor in MIT in artificial intelligence. Who else was there? Early on, then there are kind of a lot of language people who I didn't know so well. But Edinburgh was very well known for language. There was Henry Thompson, is he still around? I'm not sure. I've a feeling he may still be around Edinburgh, and another guy called Graeme Ritchie. And, you know, there was a very strong tradition in language. Another one called Mervyn Jack, I think. Just one more, I'm trying to remember the name. So now they have a kind of quite a strong group in dialog systems. What's the name of the professor there? [Added post interview: Steve Renals.] Who is still, you know, very active, but I just can't..., but the reason that there are such a lot of people now in Edinburgh working on language is because back to the Alvey programme, this research program in the early 80s. The language people were very successful, raising funds from the Alvey programme, and they built quite an impressive activity. And there was that whole department building on South Bridge. I guess it's the one that eventually burned down, I think it's the same one, isn't it? You know, that's where all the language people hung out. That's been a big success of Edinburgh. Whereas the robotics kind of died off a bit. Edinburgh was not good at robotics until the last decade and a bit. So people like Sethu Vijayakumar, he's very well-known now for his robotics. But that's all relatively recent, that's well into the current millennium, I'd say. That that's that stuff has emerged again. But there was a long time in between when nobody in Edinburgh was really particularly good at robotics.

**Vassilis Galanos 43:15**

Yeah, very interesting. So you've mentioned the Lighthill report before, and you brought up a very interesting observation that it has been captured in some of the AI history books. But I haven't found any authoritative account of the one that you just mentioned, that after the Lighthill report, people started traveling, people started looking for different funds. Would you like to expand a bit on that? Like, where did people go? Another thing that I have picked up from people I was in conversation with, is

Please note that this transcript has been lightly edited relative to the original audio, in order to improve readability.

that, because of the winter during the Alvey times, although what was funded could fall under the umbrella term of AI, the nomenclature wasn't really around AI, there were different names to describe it. And I don't know to what extent this was true or not, or how did you perceive it during that time? If you talk about AI, did you talk in secrecy about AI, but didn't?

**Andrew Blake 44:31**

Yeah, it was very unfashionable to talk about AI for a long time. When I moved to Microsoft, I think even then we didn't start talking. We didn't think it was safe to talk about AI until about 2010, it was really a long time before people were doing it, I guess. But I used to say I did computer vision because that sounded more sensible. Whereas nowadays, I would say I do AI vision because it's more sexy, attracts more. If you're giving a talk about something, and you want the audience to actually come to your talk, then it's a much better idea to say AI vision in terms of publicity. So yes, it was quite unfashionable for a while, I'd say. Nobody thought it was working, you know, when I was doing computer vision, nobody thought this was actually working, in a way that you could say, this is a piece of technology that you can use. It was all very futuristic, and, you did experiments to see what might one day be done. So, my experiments with computer vision that I did in my thesis would take an hour or more to process an image. And there was this special purpose computer down at UCL which would do it instantly. That was what was so interesting about that machine. But for most people, the computer would be grinding away for an hour just processing one image. So it was very far in the future that we were expecting AI to go live, if you like, it didn't change until about the turn of the millennium. And that was when I moved to Cambridge, joined Microsoft, because I thought, oh, actually, we can really do this now. We can do something commercial with this technology. So yes, I mean, my memory of all that is a little bit hazy. I expect people in the AI department in Edinburgh were probably more willing to use the term AI than I would have been. Because I was thinking more that the orthodoxy was computer science, and AI was slightly wacky, you know. So you kind of took cover with the term computer science. And for me, that lasted quite a long time. There was another thing about that schism between Donald Michie and his Machine Intelligence Research Unit and the AI Department. And the story, as he told it, I'm not sure that this is really a balanced view, but the story as he told it was, the University asked me to become a teaching department of AI, after this sort of scandal with the computers, and the misuse of the computers. And his view was, AI is not ready to be a teaching subject, it's still too experimental, this would have been in the late 70s. And so he wasn't willing to stop being a [researcher?]. So you know, because there was no funding, the University said, well, the only way these people can continue in the University with their salaries paid is if they teach students. So that was how most of them justified their salary. Donald Michie had a very secure position as a Professor. In those days, if you were a Professor with a personal chair, you could do just about anything. The contract was incredible, you could only be sacked for gross moral turpitude, so he could just do anything he wanted. So he wasn't going to move, and neither was he going to teach. So they solved the problem by putting him in this little research unit with minimal resources. He had this tiny grant from the University for his research unit, which he used to take great pleasure in once a year, spending most of the grant on a single lunch, that, he would invite some visitors and this was him sort of turning his nose up at the University, because even though he was sidelined, he was perfectly capable of raising funds from industrial outfits like BP. So you know, he was always perfectly okay. And even SRC gave him some grants. So he was fine. And ridiculously self confident. You should try and find some recordings of him if only the Lighthill report. But there's one or two other lectures by him on the web, and you'll see what I mean about this, incredible ability to do storytelling. And, of course, he was one of Turing's men. And one of the hilarious things is that this year with the 60th anniversary in the University, and the AI safety Summit at Bletchley Park he's been totally resurrected, having been ignored and spoken about only in hushed tones for decades. The government

Please note that this transcript has been lightly edited relative to the original audio, in order to improve readability.

has said, of course, there was this great pioneer of Bletchley Park, Donald Michie. Well, nobody has uttered the name of Donald Michie in establishment circles for decades and decades. So I found that quite remarkable.

**Vassilis Galanos** 49:43

Yeah, this is extremely interesting. And I think his archive is down in London.

**Andrew Blake** 49:52

Where is it then?

**Vassilis Galanos** 49:54

In the National Library, I believe. In the British Library? Yeah. Which is extremely interesting. Why isn't Michie's in Edinburgh? So this is extremely interesting. And I know that time is pressing. So nearing the end of this of this conversation, drawing from your extensive experience and all the things you have just shared with us, what message would you like to share with future researchers in AI? And especially those who are now here in Edinburgh, and they kind of stand on the shoulders of giants? If I may, what pitfalls should they try and avoid? Where should they focus in the future. I'm aware that you're still deeply involved in the field of AI and you monitor the landscape very closely. So I think you're definitely in position to orient younger generations around the legacy they are building on.

51:16

I don't know, if there was a sort of really infallible prescription, then, everyone would do it, I suppose. But I mean, you know, to go back to Geoff Hinton, he was sidelined by most people including me for decades and decades, we never thought that neural networks was going to be the winning technology, it really didn't look like it. It looked like this kind of crazy kind of obsession that he had for 30 or 40 years. And, he kind of stuck it out. He was also very difficult to understand because he was not very mathematical, he was a very intuitive thinker. And so for those of us who are more mathematical thinking, it was very hard to follow what he was saying. There were one or two geniuses around who could interpret. And actually, Chris Williams was one of them. So Chris was one of Geoff's students. And he was one of the very few people who had a mathematical mind, who could, nonetheless, interpret Geoff's ideas. The other one that I always enjoyed was a guy called David MacKay, who, who's no longer around, he died eight years ago or something. But I don't know if there's a really a general moral. But certainly in this case, success has come from an unlikely quarter, I would say. And that's certainly a message for funders, you know, people who choose where to send the money. Actually, Paul Nurse is right and John Bell is wrong. If you've been watching this current debate, there's John Bell, a senior physician, Sir John Bell, etc. in Oxford, very famous physician, and Paul Nurse is the Nobel Prize winner, who was president of the Royal Society and so on. And the UKRI is the big body that gives out all the money nowadays, as you probably know. And I don't know if you've been following this story, but the head of UKRI is a Cambridge professor who said, well, she's going to retire from this position next year. And then government sources have been saying, oh, we need a business person to run UKRI and not a scientist. Whereas, you know, traditionally, it's been a scientist, and certainly the current incumbent is a scientist. And John Bell, this physician, has chimed up and said, yes, there's not nearly enough strategic research direction from UKRI, meaning, defining the topics and deciding where to send the money. Whereas Paul Nurse would say, and after all, he's the one that got Nobel Prize, he would say exactly the opposite. He would say, no, let 1000 flowers bloom is what they often say, to represent the opposite pole, you know, just give good people money and let them get on with their research, and many good ideas will arise, and it's no good trying to guess which are going to

Please note that this transcript has been lightly edited relative to the original audio, in order to improve readability.

be the good ones, because that never works. You know, so let 1000 flowers bloom, then, one of those flowers will turn out to be, you know, cannabis or something really important. Or some cancer drug or something. I still believe in that to a great extent, you know, funding public research, the bulk of it should be let 1000 Flowers bloom, and that way you bring up the next generation, and the strategic has to be quite carefully drawn, quite selective. And governments have to be appeased, so you must do some strategic research. But getting that balance is kind of tricky. But for PhD students, the lesson is, well, if you've got a good idea, don't be put off, pursue it. And the other thing I tell any PhD students is, choose your supervisor very carefully. It makes a huge difference who is your supervisor. If they're too famous, actually, it doesn't work, because then they're gallivanting around the world, and you never get to talk to them. But choose ones that are not too old, so that they're not yet too famous, but really smart, and who inspire you. And I guess I didn't have that actually. Because Donald Michie, though he was very smart and very interesting, he didn't work in computer vision, didn't know anything about computer vision. He just wanted somebody in the lab that did it. So I was completely on my own to begin with, and only eventually found some muses, guiding lights, but that was outside the lab.

**Vassilis Galanos 56:12**

Yeah, very interesting. I've spent loads of time to try and navigate where AI stands in this sort of spectrum between curiosity-driven research or sponsorship-driven research and how funding allocation happens. So it's extremely interesting to hear that from you. And I think it's a wonderful way to reach a tentative end to this conversation. I guess, I want to ask you, if you have anything further to add, any other story, any anecdote you would like to? If you want to, I'm happy to spend hours.

**Andrew Blake 56:57**

It's been a wonderful ride, I think. At one time, Edinburgh, MIT and Stanford, were the only people really doing AI, and maybe Carnegie Mellon as well. You know, it was those four places that were building robots and experimenting with AI. And I always used to like to say, it's not just about building interesting machines, it's also about understanding the nature of intelligence. And I have dabbled a bit in the psychology of vision, not just computer science. Because I found some interesting people that were doing the psychology of vision. And in the end, the race to understand intelligence has been completely won by artificial intelligence. We have beaten the neuroscientists into a pulp, you know neuroscience is still way behind. I just met one outside the primary school where I help to teach coding. I met a dad of one of the kids there who's working in probably the most famous laboratory in the country, the MRC Medical Research Council. LMB, it's called the Laboratory for Molecular Biology, they have 16 Nobel Prizes in one lab, and it's just an incredible laboratory. But they've been studying a little worm called the C elegans for 20 or 30 years, they've sliced it up very finely. They know everything about this worm, its genetics, its wiring diagram, everything about it. They still can't explain how it works. And you know, even if they have a complete blueprint, they still don't know how it works. So if you can't do that with a little worm with 306 neurons in it, how are you going to do it with the human brain that has 10 to the power 11 ( $10^{11}$ ) neurons in it? And of course, every C elegans worm has the same wiring diagram. That's not true of humans, we have about  $10^{11}$  neurons, but I may have a few more than you, or a few less than you. You know, we don't have the same wiring diagram. So the idea that what they nowadays call the Connectome, if you've come across this idea, the Connectome Project where you can sort of slice up the brain, trace the pathways, it's amazing that you can do it, but it won't shed any light at all, I don't think, on the nature of intelligence. And I think neuroscience has completely failed to explain what intelligence is, and AI has completely succeeded. You know, the last year has just been the most incredible success. I mean, it's just unbelievable, I think, the extent to which.... We were already doing speech-to-text transcription, that's already amazing. We were already doing language-to-language

Please note that this transcript has been lightly edited relative to the original audio, in order to improve readability.

transcription, amazing. But this year, building machines that can do question answering over a broad range of topics and compose essays, it's just unbelievable. I think, nobody thought, even Isaac Asimov didn't completely anticipate that. He anticipated the voice driven dictation machine. But he didn't, I don't think. The Hitchhiker's Guide to the Galaxy is the other wonderful piece of science fiction, this is it, you know, this is actually the Hitchhiker's Guide, plus now we also have the Babel Fish.

**Vassilis Galanos** 1:00:30

The Babel fish.

**Andrew Blake** 1:00:38

You know, we've got it all, it's all actually happened. And it was supposed to be crazy science fiction. The only thing we don't have is the infinite Improbability drive, it would be nice to have one of those, but I'm not sure, there maybe some physics problems there. And so, I just think the story of AI over just 60 years, or actually 68 years, AI is 68 years old, and I am also 68 years old, we were born in the same year, 1956. It's just been an incredible story, the subject going from nothing to explaining, going a long way towards explaining one of the greatest scientific mysteries that we have, which is, what is intelligence? I just think it's incredible. The intellectual kind of splendour of that story is just amazing.

**Vassilis Galanos** 1:01:28

It's an amazing way to end the conversation. I have millions of questions, but I think that I will make myself available at your disposal if you ever feel you want to share more. So I want to thank you again, Andrew.

**Andrew Blake** 1:01:44

You're welcome to send me if you're writing something, I don't know if you're going to write a book or what?

**Vassilis Galanos** 1:01:51

So, shall I stop the recording, and I can do let you know on this?