AICSCSHistory_Interview009_BobFisher

SPEAKERS
Bob Fisher, Vassilis Galanos

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

Vassilis Galanos 0:03
Okay, this is Vassilis Galanos. This is the 13th of August 2025 and I'm here with Professor Bob Fisher. This is part of the Edinburgh history of artificial intelligence, cognitive science and computer science project. So Bob Fisher, before we delve into more specifics about the history of AI at Edinburgh, would you like to tell us a bit about your trajectory, your own research contributions, and like we said, Before, you were pointing towards your thesis that is just next to us here.

Bob Fisher 0:39
All right, so well, the trajectory started in high school, and there was, I had access to a programmable computer at the local junior college that nobody was using, and my high school math teacher said, you might want to go, do you know, see something about it? And I went, and I had a lot of fun writing the hello world sort of program. And then, and then I had a little mind puzzle that I was having trouble assembling. I thought, Oh, maybe I could write a computer programme to solve it. Okay? And actually, that was an AI program, 1968. I didn't know, I'd never heard the term AI, but it was essentially an AI search program---backtracking and generate and test and all the kinds of things that you would find in the early AI. So in retrospect, I felt very proud that I had invented all that myself, and it did solve the puzzle. And what it led, is when I ended up in university, at that time, there were no computer science degrees, so I was a math major, but one of my colleagues in the dorm that I was living in was a PhD student in essentially an AI group at that time. And I had sort of asked, I guess, a bit about the possibility of summer jobs. And so he introduced me to the professor, and I got an internship, summer internship after my first year and second year, and kind of after the third year. In the third year, actually, I was an intern in IBM translating the project that we had been working on there at Caltech to the labs to demonstrate how it might work. And anyway, it was an AI project that I was working on. It was a natural language processing project where you could type in questions in natural language, and it would essentially search through databases of tables. It was a relational database kind of system and extract the answer and then generate natural language text to give the answer. And that was all very exciting for a young, impressionable person who had no idea of the whole rest of the world of AI, but it was very exciting, and it kind of essentially led me to doing a PhD here. I did work for a couple of years after a bachelor's degree, and found it sort of interesting, but leading me in a direction, you know, in the commercial world that I wasn't really quite so interested in. I remember I really had loved the research at the University, which is what led me back to the PhD. I was interested

in a PhD in AI, and applied to the big four universities, and was accepted at Edinburgh. And I don't know where this is entirely relevant to the history, but because I had not been a student for four or five years, officially, it was my supervisor, my advisor back at Caltech said, Well, maybe, maybe potential PhD supervisors might think you're a little bit old and not quite as malleable as they would like, or whatever. So that might be part of it. Also, during that time, I had a worked, part time master's degree at Stanford in computer science. So now computer science degrees were starting to appear. I think when I was an undergraduate, maybe MIT might have had the only was sort of computer and Computer Engineering Bachelor's degree or something like that. But anyways, very early in kind of the computer sort of scheme of things,

Vassilis Galanos 4:55
Who was the Caltech supervisor and were they the same person as the supervisor you had for your internship that you mentioned?

Bob Fisher 5:08
So yes, my academic advisor was also my Caltech internship supervisor. It was Frederick B Thompson, okay, was his name. He was kind of big in natural language back then. Right anyway, so that, so that kind of led me here. At that point, I was kind of interested in knowledge representation as a research topic area. And when I got here, my PhD supervisor, who was Jim Howe, said no, knowledge representation is not a topic, it is a tool. Go find a problem to solve. And I did a little bit of stuff in AI and education for a while, and I tried a little bit with behaviour simulation for a while, and none of it really sort of stuck. But one of my buddies, a fellow PhD student, was doing computer vision, and he was kind of doing low level computer vision, and I'd been doing some reading. And eventually I settled on that sort of high level computer vision as my research topic, and got happily tucked into that. I think, well, things are quite, you know, different now, I think I really kind of settled into a topic after being here for at least a year and a half, maybe almost two years, or something like that. And whereas you know your topic more or less before you even start, and you start working on your thesis proposal pretty soon after you arrive. So there was a lot more flexibility to explore things back then. It was also much more informal. When I arrived, there were only about 30 people in the department, including the secretaries and the servitor. And we all would meet for, you know, tea twice a day. So you kind of knew everybody that was in the Department of AI at that at that time, and we were in the Forest Hill building.

Vassilis Galanos 7:26
And what year was that?

Bob Fisher 7:31
So I came in 1980. Yeah, so spring of 1980. I had gotten the acceptance in the fall of 79 but at that point they didn't have a studentship for me, and I didn't have enough money to pay for it all. I started working as a self-employed contractor, and perhaps did make enough money over the six months to pay for the three years here as a sort of student wages. But, but in the end, I did get a scholarship. I was sort of figuring I

would work for a year and then come, and that was the agreement. But they had the scholarship after six months, and I came. I think it might have been the forerunner to EPSRC, was it SERC, or maybe it was a university scholarship, I can't remember exactly. And it was really pretty minimum wage kind of stuff, but it was fine. And anyway, so it was about a year and a half, two years into the thesis that I kind of discovered the computer vision connection and started working on that, and worked on that, off and on. It took me six years to finish the PhD. After four years, I got my Lectureship here, and at that time, I'd estimated I only had, like, three or four months worth of work to finish it off. But, you know, with a new lectureship, there just wasn't that much time, and I was in a relationship, so that, you know, also was a very enjoyable time. So anyway, so it took me about a year and a half, almost two years after I got the lectureship, before I finished the PhD. There was another factor back in that, and that is supervision. My supervisor was great, but I think we only had a supervision meeting once a year. And this is Jim Howe and, you know now, I mean, we see our students, you know, like I see my students every week, you know. And, you know, I work with them and guide them, and all that kind of thing. I mean, he applied very high level kind of guidance, and just left me to myself very much. And as a consequence of that, well, as you can see, the PhD thesis is rather thick, and that was because he never said, you can stop now. I had no idea, you know, what one should do. I mean, I had a project in my head, and I was doing the project, okay. Nowadays, things are pretty different about that.

Vassilis Galanos 10:32
There are word limits too.

Bob Fisher 10:33
Well, I hit the word limit. I was under the word limit. And back then it was 100,000 words, not 70,000 words, as it is now, but it was pretty close. And I don't think that's the longest. I think somebody had a two volume version, but it's much more manageable now. But the good thing is, I'll come back to the PhD in a minute, but because of that content, this is a time when there were hardly any books on AI, and certainly almost none the computer vision area around. And so publishers were keen to get AI books, and so John Wiley agreed to publish, essentially my thesis. Their editorial advice was cut out a third of it. And that was it. That was, that was good advice.

Vassilis Galanos 11:31
Was it like a textbook or monograph?

Bob Fisher 11:35
Research monograph. Anyway, so all that effort that I put into the thesis, you know, and because it was computer typeset in latex at the end, it was easy to convert that to their style file, and and it also kind of led on to the Russian translation that I mentioned as well, too. So we were talking a little bit before about the latex. So there had been another typesetting system called nroff. I think it was maybe a DOS kind of package or something like that. And laser printers were just, you know, appearing, and latex was

just appearing. And one of my fellow PhD students had a lot of heavy maths to typeset, and nroff definitely was not up to it, whereas latex was. And I saw his thesis, I said, oh, your thesis looks so much nicer than mine. So I spent a couple days and re-typeset the whole thing in Latex. Latex's pretty easy, and with a little bit of help from him, it was Lincoln Wallen and so I had the second latex thesis in the department of AI. There might have been people in computer science who also used latex, but Lincoln was the first and I was the second. Lincoln went on. He did a PhD, did a post doc. PhD with Alan Bundy in automated theorem proving and reasoning, and went on to do some research at Oxford, and later moved off to the commercial world. Probably the high point of that, he was the chief technical officer for DreamWorks. And he stepped back from that a year or two ago, and I'm not quite sure what his new company is.

Vassilis Galanos 13:46
Is Lincoln his first or his last name?

Bob Fisher 13:49
Lincoln is his first name, and Wallen is his surname.

Vassilis Galanos 13:57
And who was your computer vision buddy that you mentioned from the first year

Bob Fisher 14:01
Robert Beattie. So that comes to another story. So he got his PhD. He was a student of Robin Popplestone's, got his PhD, and went off and went into the commercial world, and is still involved in that. He has a company that does essentially computer vision guided robot welding and things like that. He builds the vision components for the welding systems. And anyway, so on to the lectureship. So, as I said, when I arrived at the department, there were only about 30 people, and it was sort of kind of going, it was struggling a bit. It was surviving, largely because Jim Howe realised that the research money was not there, but the School could do teaching. So there was a first year and a second year, kind of optional course. I'm not sure about other universities, but here, our first year students do three courses each semester or each term back then, and one of them is like their degree course, maybe two of our degree courses, and the third is a free option. And so AI1 and AI2 were options, and we got students from all over doing those courses, and that brought in essentially student money, which allowed the department to survive. There were, there was also some research funding as well SERC, but because of the Lighthill report, it was a real struggle to get money. And you certainly couldn't get money for, let's say, abstract AI stuff, it had to be very practical. So there was research money for the AI and education, applying AI methods to help improve, you know, teaching, working with students with various forms of mental handicaps, neurodivergent students. And there was also money for robotics, not for robotics per se, but because of links with manufacturing and advanced manufacturing. So there was, there was money for that. And so Jim, at the time I came, was getting most of his money for the education type stuff, and Robin Popplestone was getting most of his money for the robotics kind of thing.

Vassilis Galanos 16:55
That must be between 1980 when you arrived, and throughout the first years of the Lectureship. So late 80s, mid late 80s was that sort of Alvey programme funding or different sources.

Bob Fisher 17:09
So when 1980 was around, it was just SERC. Alvey came along later. So I think it was about 1982 or thereabouts, the Japanese introduced their fifth generation computing project based around AI. And the UK said we're strong in AI, and then they looked around and discovered it was gone. I think I heard a statistic that said in 1972 there were 500 people in Edinburgh working on AI related topics, you know. And that was probably the reputation that led me here. When I got here, there it was about 30, so that was definitely the low point. But anyway, the the UK introduced the Alvey programme, and through that Alvey programme, there was research money for projects. And Popplestone had one of the projects. I don't remember if Jim had, because the Alvey programme was kind of a bit more about advanced manufacturing and things like that, I think. Anyway Popplestone had one of the projects. There were also a number of academic posts that were created. The goal was essentially to create AI conversion courses from other subjects. So you'd have a bachelor's degree in math or computing or something, and the conversion course would give you AI skills. So there was Masters in AI, actually, I think was called, actually a Masters in Knowledge Based Systems, I think that is maybe what it was called.

Vassilis Galanos 19:00
I was about to say that what I've heard from other interviews is that that time you said there was this rumour, of 500 people working. You arrived and there were 30 people. So there is some disagreement between people and whether these 500 people would be here, they would just refuse to call it AI, so they would adopt different terms in order to continue to get funding, and present it as something different. Did you call it AI then, like, how did you experience the sort of jargon-related question?

Bob Fisher 19:40
Well, I was just, you know, a little ignorant guy at that time, so I probably called it sort of AI and computer vision and stuff like that. But, you know, I wasn't at a point of applying for research grants or anything like that. I was just working, you know, happily embedded in my [PhD] project. Yeah, but you're right. I think because of the bad name that Lighthill had given AI, I think the strategy to getting money was to call it something else. And so what was called knowledge based systems, which were founded on sort of two principles of AI, one was sort of probabilistic reasoning, and the other one was sort of rule based reasoning. Anyway, that looked like there might be a lot of money to be made from that with expert systems, or knowledge based systems that could up the game of companies to improve medical diagnosis. There's legal reasoning. One, I remember there was one on analysing ground sonar signals to try and find oil deposits and things like that. I think that came out of Stanford anyway. But the naming had become knowledge based systems, or expert systems.

Vassilis Galanos 21:07
I want to connect with the Masters programme you mentioned. So the conversion course, was that speaking to the knowledge based systems? Was it like promising that you could take your domain and then use an expert system to apply it? Or was it more general?

Bob Fisher 21:32
I think it was a little bit more theoretical. We were, I guess we were training students with all sorts of skills. So I taught on that. When I got recruited as a Lecturer, I taught a computer vision course. It was called Machine Vision, Computer Vision. At various times, it's been called different things on that programme, and there was certainly a knowledge based systems component. I think there was probably a natural language component, an automated reasoning component, probably building on the skills that the staff had. And we were still such a small organisation that we didn't have a large amount of resources for a variety of topics. I mean now our Master's course, and back then, our Master's course would have like 20 students, or something like that, and our Master's course has like 350 and there's like 70 or 80 different modules that will contribute to it, based on the different expertise of all of the staff. Whereas back then it was, you know, there was no flexibility, you know, everybody had to, you know, pitch in. So computer vision was my bit and and I also did some teaching on the first and second year courses as well, too. So, so another historical bit is, okay, so when the Alvey programme came along there were these lectureships, and Edinburgh got two or three lectureships, maybe four. I think there were maybe two in natural language, which I think Graeme Richie and Henry Thompson got. Have you interviewed Henry Thompson?

Vassilis Galanos 23:23
No, but I intend to.

Bob Fisher 23:25
Right, okay, so I was here a few months [as a PhD student] before he started as a Lecturer, but he came shortly after in 1980 as a staff member. Anyway, [in 1984] there were two lectureships, one was robotics and one was computer vision. I got the computer vision post. Alvey had created two computer vision, or the funding from Alvey led to two computer vision posts. One was at Sussex and one was at Edinburgh. And at the time, there were really only three real candidates for those two lectureships. Okay, one was myself, the other one was Robert Beattie, and the third was David Hogg, who was also a PhD student at Sussex, later went to Leeds University and was he might still be there. At some point he was assistant principal or vice principal for something or other. Anyway, there were the three of us, and we were all buddies, because we were the only three doing this kind of stuff. And you know, we'd always meet up and have a beer or two at, you know,
conferences and that sort of thing. So it was a little bit of a of a challenge when there were three of us for two jobs or whatever. But we, you know, we managed to cope with that. And in fact, I think it probably, probably worked out for the best, because I think Beattie was much more entrepreneurial than either David or myself, so the

direction he kind of got nudged actually was probably the right direction for him. But you know, who knows how your life would go according to which door opens and which door closes? Anyway, so it was at that point then the Department started growing. I mean, we were bringing in the Masters students. And it kind of went from 20 to 40 to 60 to 90 Masters students over the years, and that brought in, and also a lot of these students were overseas students who started the kind of the overseas income kind of thread, which also supports the ability to recruit additional academic staff. And the summer projects sort of feed into people's research proposals and things like that. So things started to go like that. Popplestone left shortly after I became a lecturer, and so I inherited his Alvey project, along with Pat Ambler, who kind of did the robotic side, and I kind of did the computer vision side. We sort of collaborated on that for a few years. So fast forward a little bit, so sometime around that, I don't know the exact date, there was the REF institution, the REF assessment process [added post-interview: actually the RAE at that time] came along, and I was not personally involved in that, but the University submitted, AI had a submission, Computer Science had a submission. And Cognitive Science, which may have been a department by then, maybe not, anyway, made independent submissions, and they were incoherent and overlapping. They were criticised for all of that, and that kind of feedback came back to the University, and the university essentially said, right, you three guys are going to go together. Gonna put you into room, you can walk out as friends, or just one bloodied person can walk out, leaving two corpses behind, or whatever, but you're gonna merge. Which is what did happen eventually. Alan Bundy was the first Head of School. It wasn't called the School back then, wasn't called the department. Forget what exactly was called. [Added post-interview: Division of Informatics.] For about a year or so it was called something, then the University decided everything's going to be a School, and so we were renamed.

Vassilis Galanos 28:12
Was it late 90s? You think? Or early 90s? Immediately after Alvey? [Added post-interview: 1998]

Bob Fisher 28:19
Yeah, I could probably look that up. You probably got it in your paperwork somewhere already.

Vassilis Galanos 28:26
The REF exercise is interesting, how this played a role. Because this I've touched, well, interviews have touched on the general impact of the REF, but not this particular one, which I find interesting. What you said about the merge is also very interesting. Yeah, we can look it up later.

Bob Fisher 28:45
Yeah, I would guess that it's sort of, do you remember when the fire was?

Vassilis Galanos 28:54
2002

Bob Fisher 28:55
2002 so the merger would have been before then.

Vassilis Galanos 28:56
So the Informatics School, I think, was established on year 2000 or 2001 and so this must have been, I guess, late 90s.

Bob Fisher 29:14
Yeah somehow like 96 sticks in my head. But I don't know. I don't know. It's just a number. I don't know why that number pops into my head.

Vassilis Galanos 29:23
And I remember from other sources that at the time, around 2000, if you got a PhD from this sort of recently merged School, you could decide on whether the title would be a PhD in Informatics, in Artificial Intelligence or in Computer Science, and you could select based on your own career preferences.

Bob Fisher 29:49
Well, there was, it took a while for the School to really gel as a School. I mean, you know, there were the three Heads of Schools [added post-interview: Heads of Department] you know, all carried on. And they all, you know, had strong personalities and didn't always agree with everything, and I think there was maybe a little bit of resentment. I mean, for example, I think Jim Howe kind of really didn't, you know, he had worked hard to build the AI brand, and I think he kind of resented losing that brand when it became Informatics. He was really, in my mind, he was the saviour of AI here at Edinburgh, because with the, you know, Michie and Meltzer and Longuet-Higgins and the fourth person [Added post-interview: Richard Gregory], you know, the fallout between them, it all could have, you know, just evaporated. And Jim, you know, picked up the pieces. He might have been even just a post doc at the time, or something like that, or maybe just a junior lecturer, picked up the pieces when, you know, these big guys were battling it out a little bit, and put a lot of effort into, you know, essentially saving what was the Department of AI at that time. There were a couple of other units around, but everybody had left, except for Michie. And of course, Michie, being such a strong personality with quite a reputation, there was always competition between the AI Department and Machine Intelligence Research Unit. And Michie went down the route, I think, largely of commercialization and research, whereas the AI Department did the education type stuff. We did research, but our bread and butter was the education, the AI1 and AI2, and then later the Masters course and that, you know, put us on a firmer foundation, because, you know, what is a university for? You know, you don't get paid to do research. You get paid to teach, you know. You know, we're here to create, disseminate knowledge, including teaching. So Jim put a lot of effort into that. Also, you know, whenever there was an opportunity, with all the rise of the knowledge based systems and expert systems, you know, he was right in there. There

was a spin out, quasi commercial organisation called the AI Applications Institute. Have you run across this?Yeah, that was sort of sat parallel to the School they were doing the commercial stuff, and Jim was like the Director of that, as well as the Head of Department. And this is, I think this is part of why I, you know, saw him formally once a year for my PhD. He was just too busy with other things, and I was not a problem, so he let me get on with it. In later years, I did ask him, you know, I think, you know, if we'd met more often, I probably would have finished my PhD sooner and stuff like that. And he said, well, you know, when I got my PhD, I saw my supervisor twice. On the first day he handed me a piece of paper. Said, here's your project, come back when it's finished. That was the second time. So anyway, things have been improved, I think, a bit since then.

Vassilis Galanos 33:53
If I ask you to time travel back to 1980 and so on, you have just arrived here. How did you navigate your understanding of the different departments, the Department of Artificial Intelligence, the Machine Intelligence [Research] Unit, the different people, the personalities that you have already alluded to. How did you encounter all these characters? What did you make of these stories?

Bob Fisher 34:29
I was not a particularly social person, so I didn't have a great deal of contact. I was aware that Computer Science existed. I'm not sure if I remember talking to anybody from Computer Science when I was a PhD student. And they were also at King's Buildings, so that was quite aways away. And off the top of my head, I think I probably knew Michaelson. Sidney Michaelson was the Head of Department. I had heard stories that Michie had gotten the first computer at the University the year before Computer Science got their first computer at the University. That's a statement that probably needs some paperwork to document that whether that's actually true or not. I could also believe Michie would say that kind of thing, even if it wasn't necessarily true. With the Machine Intelligence Research Unit I only went once there. There was some materials that they had in their library, and I was interested in looking at them. And I went, in order to get access to them I needed to have a conversation with Michie. And it was sort of like a job interview, you know, I was just looking to get some papers. And you know, he really wanted to know what I was good for, I think. And at that time, he was doing a lot of work on chess, automated chess playing as a vehicle for automated reasoning and memory and learning and things like that. But anyway, it left me feeling uncomfortable. I felt like he was looking to see how he could use me, or whatever. And maybe that's just because I was a little guy and he was a big guy and, you know, something like that. I don't know, maybe I misinterpreted things or something like that. But anyway, I knew that there was a lot of difficulties between Jim Howe and Donald Michie, I guess partly because of the competing interests of the two departments, and also Jim was the, you know, kind of protege of one of the other three from the big four. And there was a lot of damage done all around in the breakup, and I think Jim suffered from that.

Vassilis Galanos 37:29
And who was the protege of whom?

Bob Fisher 37:31
I think Jim Howe was the protege of Meltzer. All right, yeah, I think, I think it was Meltzer. Again, like, I don't think it was Longuet-Higgins. I'm trying to remember who the fourth one was.

Vassilis Galanos 37:52
Was it Michaelson?

Bob Fisher 37:53
No. not Michaelson, he was a computer scientist too. There were four of them, Christopher Longuet-Higgins, Bernard Meltzer, Donald Michie,
[Added post-interview: and Richard Gregory.]

Bob Fisher 38:17
I can't think of anything right off. It's probably in some of the paperwork of Jim's. Another little anecdote. Robin Popplestone was a Lecturer here when I came, but he did not finish his PhD. In theory he wrote his PhD, typed it all up, and he sailed the boat from Ireland to here, and the boat sank with the PhD manuscript on it. I don't know if there's any truth to that, because only Robin knew, but Robin was really smart guy.

Vassilis Galanos 39:09
That's what the name POP1, POP2 derives from.

Bob Fisher 39:13
So that was a language development with him and someone in Computer Science, was it Burstall? There was a collaboration between him and a computer scientist to develop the language. The language was used very heavily in the robotics work that was done here. The robotics work was that was done here, I always associated with Popplestone, but in fact, actually it was Michie who started it all, and developed the first Freddy robot, and, you know, led for the funding that led to the development of the second Freddy robot, which is the one that's in the museum

Vassilis Galanos 39:51
Rod Burstall played a role.

Bob Fisher 39:59
Yeah, a little bit, not, not too much, but a little bit yeah. And who else, it was Steven Salter and Gregan Crawford. And somebody came from SRI, settled in the UK. I think he might have been British, might still be alive, but he will be in his late 90s, and my understanding is that he is not really communicative anymore. [Added post interview: Harry Barrow was actually at Edinburgh before going to SRI.]

Vassilis Galanos 40:38
Do you remember the name ?

Bob Fisher 40:42
It's coming, Harry Barrow.

Bob Fisher 40:52
Who was the fourth big [name]? That name may come anyway

Vassilis Galanos 41:06
When you arrived here was Longuet-Higgins still around?

Bob Fisher 41:10
No, he'd left for Sussex, I think. I met him a time or two at Sussex. And, okay, so when I arrived, it was Bundy, Howe and Popplestone I think were the three academics. I think that Epistemics, Department of Cognitive Science wasn't a department yet. It was like a School of Epistemics, yeah.

Vassilis Galanos 41:44
Did you participate in that at all?

Bob Fisher **4**1:47
No, I knew a couple of PhD students. Because there were so few of us PhD students, we hung out a little bit. And also it was, it was in Buccleuch Street, no Buccleuch Place, and so it was close enough that, you know, we would have events. You know, we would see each other at those events, but I didn't have a great deal of contact with them. I think I did write a report on sort of AI and education, sort of a survey report, literature survey, and I think I might have talked to a couple of the people from there as part of that. They were focusing a lot on natural language. Oh, another person, maybe it was Chris Mellish, was maybe here. Peter Ross, Chris Mellish, Peter Ross and Graeme Richie were here, they might have been post docs at that time.

Vassilis Galanos 42:53
I've met Peter Ross ages ago, based at Napier University

Bob Fisher 43:01
He's still around today. He and I went to go visit Nan Howe, the widow of Jim Howe, sometime in the last year for social visit. Catch up. You might be, have you talked to her?

Vassilis Galanos 43:18
No

Bob Fisher 43:19

You might find that interesting as well, too. She, you know she was, she lived through a lot of it along with Jim. So anyway, so what more? Okay, so there was the AI Applications Institute. And I think when the School formed, Jim kind of felt like, you know, as I was saying, he'd felt like he'd sort of lost his baby a little bit. I think he was a bit discouraged about that, and retired, you know, he was of age to retire, and I think back then you had to retire at 65 anyway. And so he retired. But I think he'd kind of lost interest a little bit as well too, but there was still the AI Applications Institute, and he came back to be [Added post interview: the Director again]. After he'd retired it kind of had started to suffer a little bit, and he came back as the Director of that for another year or two. But I don't know all the details of that, maybe Nan would know more about it. But have you talked to Austin Tate? Yeah, right. So, he'll have an opinion on it. I think there was, I think there was a kind of conflict between Austin and Jim.

Vassilis Galanos 44:46
From what you're saying, it seems that there was a bit of difference of ideas as to whether AI should be commercialised or be more like research oriented. Was that something that you were encountering?

Bob Fisher 45:07
Not really, I think there was, there was room for both there. And, you know, I think Jim, you know, enjoyed research, and enjoyed the research that his students did, and he had his own kind of AI and education research projects that I think he enjoyed. But I think he felt that, you know, an organisation runs on money, and research is all fine and dandy, but if, you know, if nobody's bringing in the money, you're not going to get any research done, you know. And so he really pursued that money, and that's sort of what made the difference, I think, yeah. And that's still a little bit the case, I think. You know, we're, you know, there's, there's better [support], we've got the business development executives, and they do a really good job here. But you know, we're very academic, you know, as an example of that, I do computer vision, which has, you know, hundreds, you know, of possible applications. In 40 years, I've had about maybe three consultancy contracts related to computer vision stuff. And, you know, because, you know, I guess my reward system is not money. My reward system is papers published and promotion and things like that, and so that's kind of where I put my time in. Whereas there's other universities that are very heavily commercialization focused, and they do quite well with it. Let's see, another thought. Okay, so we were talking about the Freddy robot briefly. When I came the robot was still in operation, it was still in use. Largely for demos, but it became a kind of test bed for several of the PhD students to use the robot for different assembly tasks. And so there was still a bit of that, but it was kind of, so this is the Freddy II that's in the museum. And it was kind of sat around a lot. I think that's also the nature of a lot of research equipment. You know, you have a research idea, you have some success in raising some funding for it, you buy some equipment for it, you do your project. The end of three years, that project is finished, you move on. And what do you do with the equipment? And for a while we had a little museum of robots around here, down in the basement. Here, there's the cafe or not the basement, but the cafe area, where there were some glass cases which have been repurposed. And all of the little museum, all the little robots that we built over the years,

they're sitting in a box somewhere. I was the curator of that little exhibit for a while, but they repurposed the area. And I don't know there was supposed to be a place to put the robots, but, you know, somehow it never happened, I guess.

Vassilis Galanos 48:25
What happened to these robots now?

Bob Fisher 48:29
The little ones, they're all just sort of sitting in a box somewhere in one of the labs. And the big ones, I guess, are sitting in the corner of somebody else's lab. Some of the big ones you could see in the in the glass case to wall, or the glass wall down there. Most of the ones that are sitting in the front aren't used anymore. Same kind of problem. Okay, with the formation of the School, there were the three departments that kind of came together to form the School, and Computer Science really had two groups. It was the theoretical computer science, Laboratory for the Foundations of Computer Science, and kind of a systems group, and they became two of the research institutes here. The cognitive science department really became the foundation of the Institute for Communicating and Collaborating Systems [added post-interview: it later became ILCC, the Institute for Language, Cognition and Communication], which is one of the research institutes here. Sort of natural language processing, the education workers, AI and education was in that group, speech generation, speech understanding. Those are kind of the main themes. There are large language models now. And then there were three more new little institutes that were created. One was the AI and its Applications Institute [added post-interview: initially CISA the Centre for Intelligent Systems and their Application; renamed AIAI in 2019], okay, which had very similar name to the spin out, but that's what it was. And they did a lot of the, so Alan Bundy was in that group, and that was sort of knowledge based systems and reasoning, planning and that sort of thing. There was a Adaptive and Neural Computation Institute, and that was sort of machine learning, and David Willshaw kind of led to the foundation of that. And I think Chris was maybe the second director, or maybe he was the first real director, or something like that [added post-interview: David Willshaw was the founding Director of IANC 1998-2004]. And then the final one was Institute for Perception, Action and Behaviour, which is kind of robotics, computer vision, and graphics is what we did. And I was the first director, the creator, and the first director of that. And it was kind of a way of helping amalgamate some of the different pieces of the original departments, and group them together, to get people to talk to each other, rather than stay in their own silos quite so much. And so the robot and vision people kind of came from the AI department, but the graphics came from, there was a guy, a Lecturer in Computer Science who kind of joined there. And eventually we were all together in the Forum after the fire. So that kind of helped as well, too.

Vassilis Galanos 51:39
Who was the Lecturer from Computer Science?

Bob Fisher 51:43

Eric, Eric, I want say Eric Grimson, but he's an MIT guy, Eric, Eric McKenzie. Eric McKenzie, yeah, so he was another one of these people, academics, who never had a PhD. So were there others around that I'm not aware of? I mean, Robin didn't, Eric didn't. There were probably a few others, you know, they came from industry or something like that, you know, because, you know, there weren't Bachelor's degrees in, you know, computer science and, you know, like Jim, Jim was a psychologist. I don't know, there were mathematicians and things like that, electrical engineers, But I think as the as the heads of the original departments kind of retired and left, that created an open space for people to integrate a little bit more, and also moving into the same building really helped as well too. Right, more about AI and the early days.

Vassilis Galanos 53:13
I usually ask people, you know, if there are any figures that have been lesser known, you know, their names don't really feature on most available histories of AI. Other people that you feel that you know, if they had like a chance to shine for different reasons, you know, who would these be? You already mentioned a few names. That's why I'm always asking, who is that person. But are there other people you might have in your mind that you'd like to highlight?

Bob Fisher 53:42
Well, I mentioned Pat Ambler before, and she remarried and became Pat Fothergill. I mean, she was a woman, you know, in robotics. I mean, she may have been like the second woman internationally in robotics. There were a few others, people in computer vision, robotics, these sorts of people from those early days, but there were very few women, and she was good. Oh yeah, there was, there was a little exodus to Aberdeen, Derek Sleeman, who was a, I'm not quite sure he might have been a student here or something as well, or maybe, I also saw him at Sussex, but eventually he landed in Aberdeen. And Chris Mellish, Graeme Richie and Pat Ambler all went to Aberdeen. I think when, when the School was formed, the atmosphere changed a bit. I'm not saying it was bad or anything like that, but before, you know, we'd been in a, you know, a small department, you know, everybody knew everybody, you know, and you didn't have an HR person and all that kind of stuff. You just, you know, talked to people and resolved things, and things happened. Whereas, once we became a School, you know, you want you, we went from, you know, just to something like 20-25, academics or whatever. And, you know, with little siloed institutes. And people, some people who liked the sort of family atmosphere originally felt a bit lost, I think. And Aberdeen was attractive in that sense,

Vassilis Galanos 55.47
Because it was smaller, and preserved a sort of family atmosphere, yeah. Were these the people you mentioned?,

Bob Fisher 55:57
So another person who left was John Hallam. He was a PhD student the same time as I was. He was in the kind of robotics and vision sort of area. He was a student of Jim Howe's as well. And he was, he was doing undersea localization you know. You know,

if you're a robot on the land, you can see, you know, though there's a tree, there's a building, nothing moves. In the ocean, everything moves, so how do you know where you are? Because where is a kind of relative concept if everything is moving, you know. So anyway, that was his PhD research, and later carried on with that, but he moved to Denmark, South Denmark University.

Bob Fisher 56:46
Okay, and again, I think it was because it was a smaller place. And also, I think Denmark, socially is a more progressive nation than the UK. So and John had come from Jersey, which, again, is a smallish sort of place, and I think he just liked it better there. And he's still there, we're still in touch. Other people from that time. So I mentioned Steve Salter, so he was involved with the robotics project, later became a professor in engineering here, passed away a couple years ago. Well, you know, we're all reaching that age, and these are the people who were in the generation above me, so there's very few of them left, if any, that's why I've got these [archives]. Their partner said, here, I don't know what to do with these, can you take them?

Vassilis Galanos 57:59
So the family atmosphere you mentioned, that's what many people have told me, that this is what made Edinburgh AI really unique at the time, because everyone could speak to each other and and so on. Do you have anything else to add, like, what made Edinburgh like significant at the time, maybe in terms of research and reputation, or do you think that Edinburgh has influenced? So John Hallam is a case who moved from Edinburgh to a different place. I've heard stories of people getting their degrees from Edinburgh and then moving to a different country and starting an AI sort of faculty project there. Have you heard of such stories like how the Edinburgh legacy travelled around?

Bob Fisher 58:49
Yeah, there's certainly a few of our PhD students who've ended up as professors at other departments. So Sam Steele was one of those. Who else, oh Roy Featherstone, he was a PhD student at the time, did a lot of sort of theoretical robotics related to kinematics, and went off to Australia. There was a lot to do with Prolog in the early days here, the language is still, still sort of kicking around. I adapted a version of it for a micro computer. Fernando Pereira. There was this spin out company, I guess, related to that. He went off to Google. I think, eventually, I think he's still there. He was here. He gave a talk a couple years ago. Prolog was going to be part of the foundation of the Japanese fifth generation computing project, as a kind of competition to MIT's Lisp machine. And neither of them really went very far, as far as I can tell. Had a lot of influences, but were unsuccessful commercially.

Vassilis Galanos 1:00:15
A side question on that. So I've heard many different interpretations of the impact of the fifth generation computer programme. Spanning from it's over- exaggerated, these projects would happen anyway, up to it has had massive influence. If it weren't for that,

we wouldn't have the American strategic computing initiative, we wouldn't have Alvey and so on. But I was recently, I recently realised that Prolog, Edinburgh based, was the basis for the fifth generation computer programme. So how did it feel to be in Edinburgh? Were there like collaborations with Japan at this stage? And what did these mean politically, when Alvey was a response to the to the fifth generation computer brand? And what if we situate that even at a global scale, because Europe, the European Union, well, the European Economic Union have also had its own project, as in response to the fifth generation and the United States. So how, how was that perceived at the time?

Bob Fisher 1:01:25
I don't think I can answer any of that. I think I was just a little bit too early career, you know, just a PhD student. And also, I've not been a hugely political person, you know, I tend to focus a little bit on my own thing, a bit too much, maybe. Were there connections, there probably were, but I'm Peter would Peter Ross might know more about that than I, because he was a couple years more senior. When I was the PhD student we did have a Japanese visitor, but he was more of a robotics sort of guy. He was doing, like a post doc year or something here, and later went back to be a Professor at the University of Tokyo. I was, as I was getting near the end of my PhD, I was thinking, you know, what am I going to do now? And I had a thought to go and be a post doc in his lab in Japan, and there that would be doing vision and robotics see things, nothing to do with Prolog. But as it turns out, I got my Lectureship here and I gave up the post doc opportunity. I have a vague feeling that there were Japanese people who did come around, but I wasn't really part of that connection. Let's see Prolog itself, I think really took off because a guy named David Warren, I think really developed it here. It was, I think maybe it came from a French researcher before that, but I think it was David Warren who really made it big, and also went off to start a company involving it. So that was one part of what made it big. The other part that made it big was that Clocksin and Mellish wrote a book, you know, a programming language book on Prolog.

Bob Fisher 1:03:36
I'm trying to get rid of all my stuff.I did have a copy of it for a while, but maybe I don't have it. Maybe I gave it away for the library or something.

Vassilis Galanos 1:03:57
I see some Japanese there, or maybe it's Chinese.

Bob Fisher 1:04:00
It's Chinese. Yeah, I had, it was an EC funded project. It was a big data project in the kind of the big wave of first big wave of big data. And most big data was text database scanning, you know, like Twitter sort of thing, or Wikipedia. And we looked at images, a big database of images. And it turns out the source of the images were undersea cameras. There were about 10 undersea cameras with our partner in Taiwan, and we were looking at tropical coral reef fish and doing detection, tracking and identification of, ultimately, what's about, approximately 1 billion fish instances. They may have been

the same fish seen a billion times. Well, there was a certain amount. There's a joke. I mean, some of the cameras were fixed, so they would look at a certain place. And a number of the fish that we were looking at are territorial, so they would always be in the same place. So you would see them lots. Anyway, that was an EC funded project. I was the consortium coordinator of it, and there was a partner in Italy, the partner in Taiwan, a partner in the Netherlands, and maybe another. It turned into a book.

Vassilis Galanos 1:05:28
Fish for knowledge: Collecting and Analyzing Massive Coral Reef Fish Video Data

Bob Fisher 1:05:42
So it was, you know, essentially, we were doing big data, but with images as the source of the data, rather than text as the source of the data.

Vassilis Galanos 1:05:56
Published 2016.

Bob Fisher 1:06:05
Yeah, so the co-authors are all the PIs of the different groups, and every chapter is, you know, is sort of written by a different either student or group or something. I wrote one of them in the introduction and coordinated the construction of the book, editor-in-chief or something like that. So back to the early days.

Vassilis Galanos 1:06:32
You mentioned earlier on, I think before we started the interview, you have some opinions on certain things, or interpretations you want to highlight.

Bob Fisher 1:06:44
Yeah, well, I mean, I mentioned a bit about my interview by Michie, and feeling uncomfortable about that. And also, as the before the School became the School, we were kind of an amalgamation of the three, and we were trying to do various things to kind of get us together. And one of the things that we thought we might do was try and establish a Graduate School for the PhD students to kind of amalgamate and associate and stuff like that. And I was charged with trying to do that, and I went around to the three schools and talked it up and made presentations, and it sounded like all, of the schools were on board. But then when I came back, we had to get it sort of approved by the Department of AI, and Jim closed it down. He'd commissioned me to go do it, and then he killed it, and this, and he never explained to me, you know, why he did that. That made me quite angry, because I'd put a lot of effort into it and that kind of thing. And, you know, and I was his student, you know. So why, you know? And anyway, so it, and I think

Vassilis Galanos 1:08:05
Killed it in what way?

Bob Fisher 1:08:06
Well, it just all three departments would have to agree to form this Graduate School, and he, you know, convinced our school to our department, to not go into this. And I mean, I've never been a great debater, and I think he pulled a few debate tricks on me. And, you know, I mean, I think he had decided, you know, what the outcome was that he had wanted. And I think he was quite good at chairing committees and getting outcomes that he wanted. Which is a good skill to have, but I but anyway, that I think that sort of set back the Informatics project a bit. And my kind of hypothesis about that was that he really didn't want the three departments to be amalgamated. As for reasons, you know, because he invested so much in it already, per se, and he felt maybe that was the future. So you kind of kind of see a little bit of that. I mean, there's this deep, deep learning and large language models kind of stuff. I mean, that's like at the core of AI, but at the moment, it's just about at the core of everything Informatics does.

Vassilis Galanos 1:09:32
Yeah, so in retrospect, you know, would he be like the would you see him as the wise person who could foresee, you know, the eventual victory of AI over the term Informatics?

Bob Fisher 1:09:44
Yeah it's a different it's a different way of of looking at things. You know, to some extent, you know, you got sort of computer engineering, and computers and those are just tools, like screwdrivers that you use to do things with. Whereas AI is, I think, is, in my personal opinion, is an intellectual revolution on the par with like the Industrial Revolution. And AI leverages brain power, you know, both it gives you tools to do more and deeper and faster and larger, but also because of the nature of AI agents you can now, you know, deputise a thing to do something to do work for you. Both intellectual work and, you know, practical work. And that, you know, I think you know, is going to change the world. It is changing the world and will continue to change the world for quite some time. You know, if it is the will of Allah, of course, and nothing else bad happens, like energy collapse, economic collapse, environment collapse, all that sort of thing.

Vassilis Galanos 1:10:55
External factors, external factors

Bob Fisher 1:10:59
And AI is everywhere

Vassilis Galanos 1:11:02
Have you collected all these thoughts somewhere, like your own understanding of what AI is and what it does?

Bob Fisher 1:11:09

Yeah. So I have a, I have a paper which is in review for the AI Magazine, sort of the US AAAI kind of vehicle, this AI Magazine, and it's called "What do we want from AI?" Peter Ross had been very helpful with some comments on that, as a number of other people, I think it is about to be published. It's had a first round of reviews, and they came back pretty favourable with, you know, they wanted some revisions, which I have done and sent back. So I, my feeling is they're probably not going to object, and will probably accept it. But I guess I'm happy to to share a draft version with you. It's kind of a response, I mean, there's, there are a large number of international organisations which basically have documents that say what we don't want from AI, and I, I felt we should turn it around and say what we do want. You know, rather than closing off directions, we should open up, you know, directions more. I mean, okay, so I take a kind of Western liberal approach to that, but, you know, essentially seeing that, you know, AI is going to be a great tool for the future or whatever, but you know, we should use it to create jobs, not to take away jobs, kind of things like that. It

Vassilis Galanos 1:12:54
It is interesting that when these first, let's say, policy or regulatory voices sort of emerged post IBM in 2009 and ImageNet, sort of successes. My PhD started looking at the gap that very few AI researchers were keen in proposing what AI has to do. So it was left to other commentators, vendors, public figures, to even politicians, to determine what AI is. But there was very little expertise from the domain involved. So it's very interesting to see, you know, initiatives like this.

Bob Fisher 1:13:38
I mean, there's nothing deep about it. I've got, you know, sort of 18 kind of principles or whatever, but the most of them, I think, are perfectly obvious, the only one I can think of that maybe has not been articulated before, is that, you know, what I want for AI is to be a proper engineering discipline. You know, in the same way that you know, if you build a bridge, you've got components, you've got theory, you've got tools for determining whether this bridge is going to collapse or not. And at the moment, AI, I think, is still largely an art, rather than an engineering discipline or a science. And, you know, it'd be nice to start to have, you know, standard building blocks that if you want to build an AI system, that you can say, well, if I need it to do this I need to use that building block. It's sort of heading that way a little bit, you know, you can download components from, you know, Python, you know, the pytorch kind of libraries. There's the MATLAB libraries, there's Git, lots of Git resources and stuff like that. And probably the, you know, the biggies, you know, IBM, Microsoft and Google are probably producing, you know, starting to produce tools, or whatever that you can just, you know, package together. I'm not really aware of it, but, you know, but part of what's engineering is not only having a package, but also knowing how that package performs. And that's again, with, like, the all the deep net stuff, you know, that's just black box magic, you know. It, you know, works sort of sometimes

Vassilis Galanos 1:15:26
It works fantastic, but we have no clue what happens.

Bob Fisher 1:15:30
Yeah, and when it's going to fail, whereas you can tell when your bridge is going to fail,

Vassilis Galanos 1:15:35
Well, at some point you will tell when the machine learning is going to fail as well, when somebody doesn't get a loan, I guess. But I wanted to ask you, like, based on your experience to computer vision and your trajectory and your knowledge. So building a bit on a conversation we had with Geoffrey Hinton on the terminologies. So when we talk about AI, what's the purpose of AI? For many years, the current version of AI that is deeply connectionist, you know, wasn't even called AI. So Hinton went for, you know, neural networks, right? And this sort of the AI effect thing. So different conferences dedicated to neural networks, and all of a sudden, now, after these successes machine learning becomes AI. What is your opinion on that? Like and when you think about AI becoming a discipline similar to engineering. What is it, in terms of, you know, the foundation, foundational principles behind it, going back to earlier terminologies and symbolic AI or good old fashioned, and from the computer vision perspective, what is your take on this?

Bob Fisher 1:16:54
Well, this is projecting ahead into the future. I think the in all the machine learning stuff and the deep net stuff, connectionist stuff, has been wonderful, but there's a lot of obvious problems with it. And my gut feeling is that's where the good old fashioned AI is going to come back in. The connectionism will be part of it, but, you know, the connectionist stuff isn't very good at, you know, sort of planning, you know, it just generates tokens, you know, probabilistically. But there doesn't really have, like, a big picture. And so, like, planning stuff, you know, maybe we'll put a framework for using that as a good engineering component in it, and planning itself as a technology. And maybe it might have some use there. I think there's a lot of work, particularly in the natural language area, where they're trying to kind of merge the kind of the large language models with, you know, grammar and corpus understanding. And I'm getting a little bit out of my domain here, but essentially traditional natural language kind of theory,

Vassilis Galanos 1:18:18
Dendrites and Chomskian theory, yeah,

Bob Fisher 1:18:21
Yeah, linguistic theory and pragmatics and all that kind of thing. So I think there's, there's probably a place for using, you know, some of the things that we know already to help shape the machine learning based products that we produce. And whether we call it AI or anything else, I don't know. I think that's a job for your academic grandchildren to decide, or something like that, and once, when the dust has started to settle a little bit more. So, for example, I'm the President of an association called the International Association for Pattern Recognition, and we're, we're 50 years old, and we kind of existed a little bit about five, even five years before that. And so our name

encompasses pretty much everything that AI would encompass now. Maybe we don't literally do natural language, but you can use pattern recognition technology to do natural language applications and stuff like that. Mostly it's machine learning and computer vision is what people who call themselves, people who submit papers to the pattern recognition conferences, really do. And we're an umbrella organisation of about 50 different national organisations. In the UK it's the British Machine Vision Association that does the kind of computer vision work, and they are affiliated with the international association, along with in the US, it's the IEEE PAMI. And there's a unit of the Russian Academy of Sciences, unit from the Chinese Academy of Sciences, they all affiliate into this organisation. And so we have, where I'm going with this little background is we call ourselves the International Association for Pattern Recognition, and that's the name that we started with because there wasn't really an AI association . You know, I guess, people had sort of started to use the term AI. I'm not quite sure the politics of why we, I think, I think it was probably the person, people involved is who set up the institution.

Vassilis Galanos 1:20:41
So you said the institution is 50 years old. So that goes back to the 70s. Were you part of this at the beginning?

Bob Fisher 1:20:51
No, there was a forerunner to the British Machine Vision Association that I was part of, the British Pattern Recognition Association, which was affiliated to the big thing. Then later we changed their name, but as a PhD student, I went along to their meetings and later became a member of that, but have gradually, you know, been involved with the organisation. But anyway, the point I wanted to make was, you know, what we do and what we call ourselves are kind of two different things. And we have talked about changing our name, but then we kind of lose our brand. But on the other hand, our brand is a little bit stale. Okay, so from a practical purpose, you know, it's like we have we run, every two years, we run an international conference on pattern recognition. We get about 1000 delegates. It's been growing. I would say over the last 10 years, it's gone from maybe 500 delegates to 1000 delegates. On the other hand, I was the one of the general chairs for the European Conference on Computer Vision, and when we proposed for it, typically they were getting about 600 delegates, 600-800 delegates. Two years later, they were getting 1600 delegates. And then two more years was the year I was going to be running it. Turned out it was affected by COVID so we did it online, but actually we had 4000 delegates. And, you know, and if you look at the NeurIPS, you know, they get like 10,000 people, or something like that, you know. And so all of these other and, you know, CVPR gets about 10,000, the International Conference on Computer Vision probably gets about 6000. I mean, all these other things have kind of grown whereas, you know, the pattern recognition brand has grown very slowly. And well anyway, so you know, do we change the name, okay, and lose our brand, but maybe rebuild it or something, get a little bit more of the people in or something?

Vassilis Galanos 1:23:13

Time travel back to. It's interesting, you said your initial experiment with the puzzle solving was in 1968. I believe that was roughly the year. So that was the period when Minsky and Papert's Perceptions book came out, that is meant to have left a sort of negative footprint on the pattern recognition sort of approaches. Do you think that these are sort of valid claims? How did you experience the role of that book?

Bob Fisher 1:23:45
Well, I didn't really run into that book till I was an early PhD student. And you know, David Willshaw, who was one of the early machine learning [pioneers] based on, you know, neurophysiology, you know, primate and other animal neurophysiology, you know, had a kind of neural network called a Willshaw net. yeah,

Vassilis Galanos 1:24:11
Associative nets.

Bob Fisher **1**:24:12
Yeah, associative nets. And the, you know, I'd, I think, okay, so what's the difference between, you know? Why are we not suffering from the same problem now, as Minsky and Papert showed, and I think they, they, they were dealing with very small, very small networks, you know, like half a dozen neurons or something like that.

Bob Fisher 1:24:41
Maybe it's easier to draw a picture or something like that, I'll find a space, but anyway, you know, if this is the problem you're trying to solve, the Minsky and Papert sorts of neural networks that they were thinking about sort of carves off little bits of the space and say, okay, here's three different solutions or whatever. And because their neural networks were so limited, that was about all they could do. And they found some problems that they showed that this is completely inadequate. Whereas modern neural networks, really the boundaries are like this, you know, they pick up all the subtleties by virtue of the scale of the number of components, etc. And also part of the reason why is that the neural networks that we work with now, these, you know, here I draw a two dimensional picture, and that's pretty close to what Minsky and Papert were thinking about, but modern neural networks really are working in 1000 dimensional space. And nobody understands the geometry of that very well. But it sort of kind of works and draws these boundaries or whatever that seem to capture the nuance of whatever the problem is.

Vassilis Galanos 1:26:06
I was just asking mostly on the basis of the early years, you know, and how this was perceived, you know, in the 80s and so on.

Bob Fisher 1:26:14
There was a lot, there was a lot of scepticism about neural networks, I think partly because of Minsky and Papert. And I think there was like, a lot of enthusiasm by the,

you know, the younger people saying, oh, this looks really cool. And in fact, in my PhD thesis, there's even, well I didn't call it a neural network, but it was a data network where information flowed back and forth. And so there was kind of like a pressure in that direction to use those as memory devices and or computation devices. But at the same time, there was the Minsky and Papert. And what that meant, I think, was the next generation, you know, our seniors kind of really didn't encourage us very much. Okay. And what changed that was the knowledge-based systems. Now there was a kind of, there was a kind of probabilistic sort of theme that emerged out of it. There was a rule-based theme, there was a kind of probabilistic theme, there's like Bayes networks. And Chris Williams knows all about that kind of stuff, and probably the history of it as well, too. And I think that's what made the difference. We didn't call it AI, and we didn't call it neural nets. So it was, we dissociated it from, you know, Lighthill and also the Minsk and Papert, but it was now a tool that we could use to solve some problems. And by framing it in this probabilistic way, a kind of probabilistic network, it became a valid tool to use. And so the knowledge-based systems and quite a variety of different architectures evolved out of that, that sort of direction. And I guess it kind of laid a little bit of the foundation for the neural, the large-scale neural networks. I don't know, it was, I think it was pretty much Hinton's group that kind of, it was the group in Toronto that really opened the door for that.

Vassilis Galanos 1:28:41
And that's an interesting story with Edinburgh, so Hinton sort of getting his PhD from Edinburgh but never having a postdoc here, so eventually moving with Alan Newell, and then, yeah, starting his career. There was also work in parallel with, from Yann LeCun and Yoshua Bengio.

Bob Fisher 1:29:00
Yeah, I don't know so much of that. I mean, that all happened before I was here. There was another postdoc, Tim O'Shea, before I was here. He ended up being the Principal here, you know, 30 years later or something like that. But I don't really know what he did, you know, he was here and gone, you know, several years before I arrived.

Vassilis Galanos 1:29:20
Yeah. Very interesting. I know time is running and you have a meeting at 2.

Bob Fisher 1:29:29
At 2. I need a little bit of food.

Vassilis Galanos 1:29:31
And some food, yeah. So any final thoughts or any advice for new researchers who want to pursue a career in AI? You already alluded to a few things, but, you know, final words on this.

Bob Fisher 1:29:47

Yeah, well, it's a very popular area and there's a huge amount of competition. I see that with my PhD students. You know, where 50 years ago we were asking big questions and PhD students could even try and do big questions. That's like my PhD.
You know, Eric Grimson at MIT did a masterful job of going from the image, he has a book called From Image to Surfaces. Is it here? I've given that one away as well, too. I mean, as I'm nearing retirement, I'm trying to give away as much stuff because I'm not going to have a place to put all these things.

Vassilis Galanos 1:30:26
I might steal a few of those then.

Bob Fisher 1:30:28
You're very welcome. You're absolutely welcome to take anything you want. I mean, other people have already picked through them. Anyway, Grimson did a book From Images to Surfaces, and I thought, I can do From Surfaces to Objects, and that's the title of my thesis. Right, here's one of his books. Anyway. Advice. Yeah, well, there's kind of a trend of PhDs that I see a fair amount of, which is you take a problem that had been solved with some traditional approach, and then you apply a deep net to it.
And the students spend a lot of time deep net engineering, and they know Python, and they know PyTorch, and they know deep net engineering, and they produce some slightly better results which are publishable, and they get their PhD. My feeling is that they've not learned the fundamentals of computer vision. What is it about? It's about shape, it's about geometry, it's about optics, it's about physics, it's about light, and how all these things interact. And then there's also all the knowledge-based type stuff, sort of like the ImageNet and WordNet. Our world is defined by the language that we use, and the objects that we have in it are linked to language. There's kind of a relationship between all of that. They don't learn any of these things. They just learn how to tweak a few parameters in a deep net to get an effect, or whatever. And my worry is that this is not giving them the foundational skills to kind of carry on with their career. I guess some become academics, and maybe those are the ones who did learn the foundational stuff. A lot of them become engineers, and in a couple of years they use a different engineering methodology, and the world changes very quickly. So my feeling is that something has been lost a little bit. The advice I would give is learn your fundamentals to underpin, why is my deep net working? I see an example of it. They're trying to do something with motion, or whatever, and their deep net isn't working. I say, well, have you tried using optical flow as an input? What's optical flow? Or something like that. Then they find a deep net program which computes the optical flow, and they plug it in, and suddenly their thing works better. I say, well, that's because it's related to motion, or it's related to shape, or something like that. That's what has been missing. Here's a way to add that in. I think some deep net research is kind of heading in that direction.

Vassilis Galanos 1:34:04
Do you have your teaching materials from the times you were teaching?

Bob Fisher 1:34:10

Most of them I don't think I have anymore. I had Jim Howe's lecture slides from 40 years ago.

Vassilis Galanos 1:34:15
Oh, nice.

Bob Fisher 1:34:17
I tended to sort of hand them on. In the later years I flipped my teaching. I didn't ever write any monographs or anything like that, I just had slides, later powerpoint slides. But in my later years I changed the way I was teaching so that all of the lectures were pre-recorded, sort of 10-20 minute topic blocks. The students had a pathway through all of those, and we used class time to discuss questions that may have arisen from the videos. Those videos are still around somewhere in the university system, and some people still use them. I think some people elsewhere in the world use them as well, too. They're probably a little bit stale, because it was all pretty much pre-deep net stuff. I think I had maybe four 15-minute blocks on deep nets, but it's probably 10 years old now.

Vassilis Galanos 1:35:33
I don't want to steal more of your time, although I have millions of questions coming up. We can continue another time, maybe.

Bob Fisher 1:35:40
Yeah, if you want. If you have more questions.

Vassilis Galanos 1:35:42
Yeah, sure. Okay, I will finish the recordings now.