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

AICSCSHistory_Interview007_JohnHallam

Vassilis Galanos 00:09
Well, this is the third of July 2024, and I'm pleased to be virtually together with Professor John Hallam. This is part of the Edinburgh AI and computer science history project, and I'm Dr. Vassilis Galanos. And so, John, thank you for being with us today. My first question is, before we delve into the specifics of Edinburgh's influence in AI and computer science, I would like you to describe your own contributions and your own research trajectory. What brought you to Edinburgh? And what did you do during your stay there, but also after?

John Hallam 01:07
Okay, so I came to Edinburgh from Oxford, where I was doing a maths degree. And the reason that I came to Edinburgh, really it was because one of my predecessors in the Oxford maths group that I was in had come to Edinburgh to study AI, that was Chris Mellish. In fact, he was, I think, two years before me, anyway, and he had sent a poster, advertising PhD places Edinburgh which was put up in the college. And that sounded interesting. And I was, at the time thinking, what should I do after my maths degree, maths being a very broad discipline. And I decided I was going to do a PhD. So I applied to Edinburgh and to Cambridge, and was interviewed at both and chose to go to Edinburgh. And so I joined the department in 1979, in October, as a PhD student. What I'd originally planned to do was to look at theories of colour vision. I don't know if you know the work of Edwin Land, he was the guy who invented Polaroid cameras. And he discovered that colour is computed in the brain rather than represented in the physical world. What he showed was that you could recognise the surface colour of objects, more or less independently of the illumination. And this is a curious thing, because the colour of the light that goes into your eye, which is what you naively think is the colour you're seeing, is not dependent only on the colour of the object, it depends on the colour of the light. If you put a green apple in very red light, it tends to look black. What he showed interestingly in his experiments was that with objects of different colour, if you can see the whole scene, then it's relatively insensitive to illumination. If the illumination is coloured in some way, you nevertheless get the correct surface colours, the objects, but if you only see a tiny piece of the object, you get the colour of the light instead. And what this tells you is that colour is a global property that's computed in the brain, essentially. We're evolved to try and recover the surface colour of objects independent of how we view them. And it's obvious why, you want to know the colour of things to know whether they're ready to eat or dangerous or whatever. So I was interested in that, but it turned out when I arrived that a guy from the previous year, called Wayne Caplinger, was already studying that. So I had to find another project. And Jim Howe, who was my supervisor, because I'd come wanting to do vision, had a friend at Heriot Watt, who ran the electrical engineering department, and they had an underwater robot called Angus. And Jim was talking to him, I think, in the staff club one day. And he mentioned that they had a need for acoustic-based communication between the robot and the surface vessel. And that was suggested as a project. And I thought that was really boring. That's just communication engineering. But in thinking about the underwater robot, I realised that there was a fundamental problem in the robotics of the day to do with navigation. So state of the art in 79-80 was Hans Moravec's CART robot at CMU. And basically what it did, it was powered by vision, navigation, as things are today. So it would take a collection of pictures very slowly, because technology was slow, it would sit and think for about 20 minutes, calculate where it was, then it would move about a metre, and

do the whole thing again. And it was obvious that for a robot underwater, this would be completely useless. Because you're not in control of your own motion. And in 20 minutes, you could have moved a long way. So my PhD, in the end, became solving that problem. And the result is SLAM. So I spent the next ten years basically working on navigating underwater robots using sonar. Having done a PhD, we then got a grant from Marconi, which was a British company that made sonars at the time. And they were interested in whether their sonar could be used to navigate an AGV. So we built a demonstration of this. And that went on, in fact, until 1990, that work, at which point, I stopped it, because it became clear that we couldn't go much further without a real underwater robot to try it out on. And those were very expensive. There was basically no chance. But a side effect of that was that a number of people had read my PhD thesis, including Hugh Durrant-Whyte at Oxford. And they, I think, got excited by the idea that one could navigate and build maps at the same time based on what you could see or observe, and the field just grew from there. So that that was largely what I did in the first ten years, I guess, in Edinburgh.

Vassilis Galanos 07:01
It's very interesting. So the University of Edinburgh has been a significant hub for research in AI. And I experienced that myself, being here now, but also slightly more different when I arrived in 2016. And I want you to talk about the research culture you experienced once you arrived from Oxford. What kind of research culture did you encounter when you were here? What was your average week looking like as a PhD student?

John Hallam 07:44
It was fairly informal. There was a set of general kind of introductory lectures on AI. Because there weren't that many textbooks at that time, and basically, no one knew what it was, except at AI departments. So I think we had one or two mornings a week we spent just being brought up to speed on the different subfields of the discipline. And that was very interesting, because that covered the whole breadth of AI. I mean, now when people talk about AI, when they talk about what they think is AI, they are generally talking about connectionist large language models. And that didn't feature at all in those days, because connectionism was a non-issue then, there wasn't the computing power to do it. There wasn't really computing power to do robotics either. But that didn't bother people quite as much. I think, because they thought robotics was easy. Like they thought vision was easy. I mean, all these things that we do so naturally ought to be easy, right? But unfortunately, it turns out that's not the case. And the same is true of language actually. So there was a lot of work that went on language, but we learned the basic computational linguistics and interaction via parsers, and so forth. The kind of technology that drives the precursors of today's LLM, which was ELIZA, Winograd 's work. So we did natural language, we did theorem proving, we did knowledge representation, we did some robotics and vision. And we had four faculty in those days and they took turns to introduce their fields. So there was Jim Howe, who was Head of Department, because I arrived following the explosion 10 years earlier. And things were fairly stable when I arrived, although there were still a lot of people not speaking to other people. There was Jim, there was Alan Bundy, there was Robin Popplestone. And there was Henry Thompson, who'd just arrived. And Gordon and. I think. Bernard Meltzer [post interview correction: Rod Burstall] had just gone off to, or were in the process of going to CS. When I came for interview, they interviewed me as well, and wanted to drag me up to King's Buildings. But I wanted to do AI, not computer science, if I wanted to do computer science I could have gone to Cambridge. But I wanted to do AI. So I stayed in town with the AI group. And in parallel, that we didn't have much to do with because of politics, there was Donald Michie's Machine Intelligence Research Unit, hidden somewhere out there in the university, which we only heard about by accident.

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

Vassilis Galanos 07:54
Okay.

John Hallam 07:54
I think that was around then, it may have come later, certainly within the period of your of your study. But that was that was another piece of the fallout from the breakup of the department before I arrived. So we had that. And then we had an amount of just studying by ourselves, basically. The whole of the PhD programme was quite loosely structured. So there was, at the end of the first year, you had to submit a thesis proposal document, which outlined what your planned PhD research was going to be. And then at the end of your work, you handed in a dissertation and got examined, and those were the only two formal fixed points. And you had meetings with a supervisor, obviously, on a mutually agreed basis. And sometimes not mutually agreed, some supervisors were very diligent and met regularly with their students, others didn't. Depends. Jim was about right. I mean, we met regularly, but not frequently. And that worked well for me. And one of the things that I had learned doing a maths degree is that a lot of work is invisible. I don't know if you've ever studied maths seriously. But basically, the way you learn to be a mathematician, is you tackle problems that you can't do. And when you're tackling a problem that you can't do, largely, you can't do anything. So you have to wait for inspiration. And when you get the inspiration, you do a bit more. And it takes you further, sometimes in the wrong direction, but hopefully not. And then you get stuck again. So actually doing mathematics is largely a question of being stuck a lot of the time, and therefore doing other things while your subconscious mind works on the problem and eventually generates new ideas to pursue. And so I treated AI research much the same. I did things like rowing and joining clubs and stuff. So it was relatively informal. And in fact, it was so informal, that when Bob Fisher joined, he came about four months after the rest of us, I think in January. He eventually felt moved to have a PhD students meeting once a week where we could discuss how things were going, just to share experiences. That was quite fun. The chief thing I remember from that is that Bob's program used to get a couple of orders of magnitude faster each week. And this went on for weeks and weeks and weeks and weeks.

Vassilis Galanos 14:05
Interesting. So you were interviewed in order to get admitted to the PhD?

John Hallam 14:12
Yes.

Vassilis Galanos 14:13
That's interesting. I mean, the rest I think the structure is pretty much the same. Maybe they try and offer some structured relationship with supervisors. I guess it never changes, some people are more active than,,.

John Hallam 14:28
I think that nowadays people try and do a bit more quality control. But yeah, in my day, PhDs were very much more like an apprenticeship than a course. Therefore, it was much less structured.

Vassilis Galanos 14:46
Yeah, yeah.

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

John Hallam 14:47
I think that the thing I appreciated most I think about those years was the breadth of coverage that we had. And one of the things I've noticed all through my career is that the breadth of AI that I'm comfortable in is considerably wider than many people's. And that goes back to us having a solid grounding in each of the different sub-disciplines.

Vassilis Galanos 15:20
Yeah, who are delivering those introductory lectures you mentioned?

John Hallam 15:25
Well, I this is partly inference and partly memory. Jim did I think vision, Robin Popplestone did robotics, Alan Bundy did theorem proving, I think Henry did the language, Alan would have done the knowledge representation. That was most of it. And some of the postdocs as well did bits and pieces. And if you were doing AI in education, then I think you got extra, Jim did that, because that was his original area of interest and expertise.

Vassilis Galanos 16:02
Yeah, I'm talking now to a researcher who's looking at the Logo turtle. Maybe I should put you in contact, maybe you know more about it.

John Hallam 16:14
It was still in use when I started there. They'd been doing experiments with the Logo turtle for some time. And Peter Ross, I think, started the same year as me. And his job was to implement Logo for the turtle, from what I recall. Okay, as postdoc for Jim. And they had a big picture in Forrest Hill on the wall of Ilona Bellos. She was one of the robotics RAs, her children playing with the turtle in one of the experiments they were doing.

Vassilis Galanos 16:49
What was her name?

John Hallam 16:52
Ilona Bellos.

Vassilis Galanos 16:54
Okay. Interesting. The person who told me about episodic memory was Peter Ross. Full cycle. Great. Interesting. Wow, fantastic. I was going to ask something about, about the research culture, but think was pretty much covered. Oh, you mentioned the explosion. I really enjoyed the term of the department. So you weren't there when this happened?

John Hallam 17:24
But no, that was 73-74. If you want to know the history of that, the person to talk to is Nan Howe, Jim's wife.

Vassilis Galanos 17:34
Okay, yes, we considered sending her a message. But we were a bit uncertain as to whether that kind of breaches, you know, private life or not. But we may try.

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

John Hallam 17:48
It's something to consider. And she may or may not be willing to talk about it, but she certainly knows what went on, where all the bodies are buried. I mean, all of the protagonists now I think are dead. Jim died five, six years ago? Richard Gregory? I can't remember. But Longuet-Higgins died earlier. Michie was killed in a car crash, I think? So that no, there's none of them around to get annoyed.

Vassilis Galanos 18:21
Exactly. Yeah.

John Hallam 18:22
But there was I remember Nan being very exasperated about it. And I think she viewed a lot of it as childish tantrums. But it was before my time. Helen [Pain] knows more, I think.

Vassilis Galanos 18:37
From your experience, was it something that impacted research?

John Hallam 18:43
No, I don't think so. I think the AI Department got on with what it did. It wasn't speaking to the Machine Intelligence Research Unit. But we didn't experience that as deprivation. That was just the way things were.

Vassilis Galanos 19:04
So you arrived in 1979?

John Hallam 19:06
Yes.

Vassilis Galanos 19:07
And I suspect you completed your PhD around 1982?

John Hallam 19:12
No, 84-85, because there wasn't so much pressure on finishing early from funding agencies. And because I'm from Jersey, not Britain, it was easy for me to get an extra year's grant. So I had four years of maintenance grant. Which took me to 83. And then I submitted my thesis in, it must have been end of September 84. Because I think that was when I became a Lecturer, because you had the Alvey Programme.

Vassilis Galanos 19:13
I will ask that you came on this transition point..

John Hallam 19:20
Because of the British government's panicky response to the Japanese announcement about solving all the world's AI problems. And then it turned out Edinburgh was the only place that had people qualified to do anything about it. So yeah, there were lots of jobs then and I was one of the lucky ones.

Vassilis Galanos 20:14

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

Do you want to talk more to the Alvey Programme and your own research? Did you interact with the Alvey Programme?

John Hallam 20:24
No, not really. No, I don't think we did. We did, at least I was involved in something called Prometheus which was a project, but I don't know what the provenance of the project was, I can't remember now. It was about self driving cars. That was in the 80s. Mike Brady from Oxford was one of the leaders of that. But otherwise, no, I mean, the Alvey Programme was largely about knowledge representation and AI computing. And I was interested in robotics and navigation. That's what I was working on then.

Vassilis Galanos 21:16
Would you like to expand on the way you approached the relationship between AI and robotics in that decade?

John Hallam 21:26
I've always thought that robotics is a subfield of AI. Because that's how I learned. The problem with robotics is that people think it ought to be easy. And I heard a story that back in the day, when they were doing AI in the States, one of the things they wanted to do was to assemble Heathkit televisions. I don't know if you remember Heathkit, you're maybe too young, it was this American company, where you could get a technological thing sent to you as a collection of components, and a big thick book of instructions, and you assembled it yourself. And they had a television and they had computers, all sorts of stuff. But the idea was that you would have a system that could read the instruction book, and then assemble the kits. And the immediate reaction of the AI people was that, oh, the thinking involved in figuring out how to assemble it is the complicated bit, we'll do that. And they thought that the vision required to read and understand the instruction book, and the manipulation required in order to build the thing were trivial. Well, or at least if not trivial, they were for the engineers, and the engineers would sort that out. And it became apparent in the course of the 80s, that that just wasn't true. That you could probably, and experience has proved this right, you could build reasoning programs that were as good as or better than human in a niche domain. And so it was easy enough to build an AI chess player that could beat the world chess champion, as long as it didn't have to move the pieces on the board. That was fine. But moving the pieces on the board was in some respects harder than playing the game. And while you get world class chess and Go champions, you don't get world class robot footballers. Because as soon as you start to interact with a dynamic world, all of the assumptions of 80s and 90s knowledge representation fall apart, the world is suddenly dynamic, you can't build a fixed model of it, you have to keep rebuilding the model of it, and it's really expensive. And then it turns out that if you want to understand just how the world behaves, you need another huge store of knowledge. Just to know like a three-year-old does, that if you drop a cup on the floor, and it breaks, the "it" means the cup, and the cup fell. And whatever was in the cup came out and splashed around. All of that is kind of common sense physics that we learn as children. But representing it in the classical knowledge representation mode turned out to be horribly, horribly difficult. I remember going to talk at one point where it was demonstrated that if you were using standard logic, and you wanted to describe the action of opening a car door, you had to do all sorts of gymnastics, because it was very easy to have a model in which the car door was stuck shut or stuck open. And if you're a little bit clever you can have a model of the car where you can open the door. But in between it might flap open and closed arbitrarily many times. But getting a model where you open the door, and it went from closed to open, and only do that once was actually remarkably technically complicated. And that was just a simple example. And it became clear that you have to do all of this to build systems that could interact with the dynamics of the world.

And that was largely what inspired Rod Brooks, because he was working, he did his PhD at CMU, I think, and he was working on vision there. And then he went to MIT to work for Tomas Lozano-Perez on robotic assembly. And he realised that they were stuck in this reasoning paradigm. And he'd stopped believing it was going to work. And he took the view that trying to model human reasoning was over ambitious, and that you should do something simpler. And insects were a good paradigm because no one would accuse an insect of being wildly intelligent. I mean, ants basically move at random, but they do it in rather clever, random ways. And yet, insects are one of the most successful genera on the planet. So why not? If we could build an insect, we'd be doing really well. I used to start talks about robotics sometimes by saying, imagine you want to build a robot that can work outdoors, whatever the weather and can come in and find its way around, it doesn't get stuck, and keep the grass down. And I put a picture of a sheep, I'd say, well, if you could build one as good as this, you'd be very pleased with yourself. But think how dumb that thing is! There are few animals dumber than sheep. But we couldn't build anything like that, and with the reliability. So that was one of the challenges that became clear over the 80s in robotics, that the classical methods of AI, of treating thinking as the manipulation of explicitly represented knowledge, just wouldn't cut it. At least if you could do that, before you could do it, you would need something that converted your sensory experience of the world into the explicitly represented knowledge, and something that converted explicitly representing knowledge into the action. And that those were equally, or maybe more difficult than the problem of dealing with the knowledge itself. So I think all of the robotics in Edinburgh more or less moved to that more minimalistic paradigm, less knowledge-based paradigm through the course of the 80s.

Vassilis Galanos 28:11
Did you play an active role in that?

John Hallam 28:14
Yeah, there were three of us. There's me, There's Chris Malcolm and Tim Smithers. And we were very involved in the direction of robotics at that time. That's not to say that there wasn't a lot going on in the more conventional robotics. In the early 80s there was Robin Popplestone and Pat Ambler, who were doing the knowledge-based robotics. But what they were interested in was assembly robotics, where the environment can be much more constrained and simplified. And they were looking at ways of describing how to put things together based on the way that you mate different geometrical features, like you put the peg in the hole, they wanted a language where you could say, put peg A in hole B, and the compiler would figure out what the robot had to do. And then towards the end of the 80s, maybe was the middle, anyway, Chris Malcolm took the position that the problem with that was that you were talking geometry. So what you were doing, your language was describing geometrical transformations, and the assembly kind of happened by accident. The robot didn't care whether it was carrying the peg or not. It would do the movement prescribed whether or not it had the peg. A human would never do that. If the human failed to pick the peg up, then they'd try again, the robot might not notice. So he became interested in doing a behaviour-based way of assembly robotics. And his Master's project in fact consisted of a demonstration that this made sense. Very impressive piece of work that was, it was called SOMASS.

Vassilis Galanos 30:00
Interesting.

John Hallam 30:02

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

Yeah, so there was that. And then there was also the design-to-product project, which was when Tim Smithers came, he was employed as its manager. Robin, I think, was the Research Leader. And the idea of that was to, that was a knowledge-based project, the idea was to capture all of the knowledge generated in the process of design of a possibly robotic system, a mechanical object anyway. And the motivational idea was that when you're making, for instance, a pump, the concept designer will design the overall function and the way it's realised, then there'll be a more detailed designer that will put together a mechanical structure that does the job, then that will get passed to a production engineer, who will decide that it's much too complicated to build that, and we'll adjust it to make it easier to build. And then maybe it'll get passed to a maintenance engineer who will decide that certain parts of it are going to be horrendous to fix, and we'll tweak it. But the way these guys communicate is through blueprints. And the blueprints contain a picture of what's been decided, but not why. So if the maintenance engineer wants to move a hole from here to there, on the grounds that it'd be much easier if the hole was over there, he doesn't know whether the hole is where it is now by accident, or there's some deep reason. And he has to figure it out from looking at the blueprints and thinking about the job. Whereas the designer who put the hole there in the first place knew why he put it there. He may have said, I've got to get rid of this hole somewhere, I'll stuff it there. Or he may have said, well, the hole has to go there for these and these reasons. If that were recorded, then the maintenance engineer or the production engineer would have a much clearer idea of what changes would be easy and what wouldn't be. And then one could take the document with all of this information in, one could provide assistance, like calculating clearances or checking tolerances or etc. So you can imagine the kind of thing. So the design-to-product system was intended to be this coherent design document all the way from conception to maintenance, together with a collection of assistants that help the designers along the way. It was ridiculously ambitious, 30 years ahead of its time, and people still haven't done anything like this. It was a really good idea, and that was something Alvey funded. And Tim Smithers came as a result of that. And I think his experience with that was what probably put him off this strongly knowledge-based approach to robotics as well, because he became more or less an evangelistic crusader for the behaviour-based robotics approach. And Chris was interested in it for assembly, and I was interested in it for practical reasons, because I'd come to the conclusion that robotics designers hadn't a clue what they were doing. You couldn't, and you still can't, I think, design a robotic system that's going to work. You can design a robotic system, and if you're lucky, it will work. But it is luck. Tim Smithers has had this nice metaphor, where he said, that if you look around Europe, there are these huge cathedrals standing. And they were built by people who had no idea of the theory of structures. Because the formal theory of structures was only developed in the 1920s, the mathematics wasn't really available before then. But nevertheless, they built these amazing things. And they're still there today. And you look at them, you think they obviously knew what they were doing. But actually they also built a lot that fell down. And it was very difficult for them to say in advance whether their building was going to fall down or not. Until it either did or didn't. And that was the situation that he felt we were in with robotics, I agreed with him, that a robotics designer who's experienced will design a robot for some task, and it'll probably work. But unfortunately, that's not good enough for what you might call commercial robotics. Because if you're a company that's delivering robotic solutions, you won't be able to guarantee them. This was before the Microsoft EULA, which includes this clause that says no, this software is not guaranteed to do anything, and if you believe otherwise you're an idiot. And it's amazing that they could get away with that. But no, that was software. In the real world you can't so easily eliminate liability for your products. So It was clear that as robotics designers, we couldn't do that, we couldn't guarantee any behaviour. We knew we could make robots. And if we were lucky, they worked well. And it seemed to me that if you took animals and studied how they worked, you might learn something about what kinds of robots were likely to succeed. Because with animals, you know that they at least work well enough. If they don't work well enough, they're extinct. If they work well enough, they survive. So you had a working example, and the question was what you could infer from that. So that led on to what I spent the next 30 years doing, which is doing more biologically inspired or bio-relevant robotics.

Vassilis Galanos 35:52
Very interesting. I have spoken to Tim Smithers a few times. I've also visited him and his wife in San Sebastian, not for the purposes of the interviews, but just because I happened to be there. So I find him an extremely great conversationalist.

John Hallam 36:11
He's a very interesting man. And he also had this much more philosophical outlook, I think. He was one of the drivers of this kind of philosophical questioning in robotics. And Chris also was more philosophical. I maintain I'm a mathematician, but people mistake me for an engineer. Not actually true. But I get on very well with engineers. I have very practical view of how things should be. But Chris, and I and Tim did a Royal Society paper on behaviour-based robotics back in the day, you can find it in Philosophical Transactions A.

Vassilis Galanos 37:04
Yeah, I think I have read that paper in order to get familiarised with technologies at the time. This is extremely interesting. Although I've spoken to Tim, he's never mentioned many of the things you mentioned about him, probably because he's also very humble man.

John Hallam 37:24
Yes, he may disagree with them of course.

Vassilis Galanos 37:27
He may, or he may not. But it's great to have your perspective as well. From what I know, I recognised some of the things you said about Rod Brooks, because I've read his book on robots, the popular Penguin book on the topic. And I remember he said that when he took the decision to look at the simpler structures and so on, he was considered very much a heretic, or even an outsider immediately, I wonder whether you, Chris and Tim felt the same during your stay in Edinburgh at that point?

John Hallam 38:03
I suppose to some extent, yes. I mean, in terms of the kind of industrial robotics world, to some extent, we were and in terms of mobile robots, I suppose also some extent. And we got interested in using learning to program robots, and I once went to a workshop. I was invited to go to workshop by Mike Brady, I think, on the strength of a presentation I'd done at his group on the feature-based mapping. But I talked instead about the work we were doing with a robot, where the robot learned how to interact with the world in a way that was successful. And they weren't very happy about it. And I thought, well, you know, this is where robotics is gonna go in my opinion, if you don't like it, tough. But that was quite interesting. So I think yes, in a way we were on the outside, but I've never felt the need particularly to be popular in my field. I've more felt the need to be interesting, to do interesting things. And I've just done interesting things. My somewhat in jest definition of the job description of a professor is do whatever you think's important, don’t waste our money or embarrass us too much. And that's kind of how I've done it.

Vassilis Galanos 38:27

Is it your definition, or did you pick it up?

John Hallam 39:50
No, no, it's mine.

Vassilis Galanos 39:52
Amazing. Yeah, I'll keep that in mind.

John Hallam 39:55
And when you get there, that's basically what they want you to do, at least that used to be what they wanted you to do, I'm not sure what it is now. Earn lots of money for the university now is probably what the definition is. But back in the day it was discover interesting things. Try and guess the future.

Vassilis Galanos 40:17
You mentioned a key term earlier on, the commercial sector, the commercial understanding of robotics and so on. Do you think, historically, that shaped AI's history from AI, robotics, bio-robotics? To what extent you think your academic trajectory in Edinburgh, but also beyond, was also shaped by commercial trends, even as a reaction to them or as a position?

John Hallam 40:52
I think my trajectory wasn't shaped by commercial trends at all. Deliberately, either for or against. I have absolutely no interest in applications. I'm much more interested in finding out how things work and why. I'm not against applications, but I don't go seeking them, they come find me sometimes. And there's times I felt that I've managed to spend 40 years doing the kind of AI I think is interesting, despite the increasing pressure from politicians and universities to do work which has some relevance to something. I never bought the argument of relevance. And it's, yes, you need people who do applied work, because you need something to take the work close enough to industry that the technical transfer people, which is one of the things my wife does, can actually deploy it, and have it earn money. But you also need the basic research. And the thing about basic research is that it's a long-term gamble. It may or may not pay off, but generally when it pays off, it pays off big. And it's almost impossible to predict which things will pay off and which things won't. The classic example, I think, was the work of Hardy, back in the 30s, he was a pure mathematician. And he despised military application, I think it's fair to say, he wanted to do work that had no military application, no possible military application. And he worked in pure mathematics, so he was off to a good start. You know, mathematics is this game where the pure mathematicians try to invent something which has no useful function at all. And the applied mathematicians try to demonstrate they were wrong about that. So Hardy worked in number theory, which is one of the more abstract bits of pure mathematics. But his work is the basis of all public key cryptography. So not only is it militarily relevant, it underlies pretty much all of the modern economics and e-commerce that we do. And he probably would have been horrified by that. But equally, at the time he was doing it, that was completely non-predictable, and yet his work has paid off hugely. And I think that the difficulty with letting politicians dictate research agendas, is that they will always be short term. Same with companies. The problem in Britain in the 80s-90s, probably still the problem in Britain, was that companies were run by accountants, I don't think it's different anywhere else in the world. But an accountant's time horizon is typically maximum three years. There's almost nothing you can do interesting in basic research in three years. You need a longer time window. There's not much you can do actually in applied developmental research in three years either. Brooks had this view that it might take you a year to make a lab prototype of something interesting. It would take you ten times as much effort

to make something that was robust enough to be worth pursuing. It will take you ten times more effort than that to make a viable product out of it. And it will take you ten times more effort than that to actually have a company selling things in bulk. And his kind of prototypical case for that was the vacuum cleaner, the Roomba. When Bridget and I went to visit him in 1992 in his lab and he had the prototype vacuum cleaner then, and it was a bit like a turtle, it drove around and it sucked stuff off the floor. And I think he knew then that it might be a robot vacuum cleaner. But that was in ‘92. And he says himself that they were so naive then, that they sucked all the dust past the electronics. And based on that, you know, he's then got this factor of pretty much 10,000 to get from that to the Roomba. So having accountants whose interest is your profit and loss at the end of the year, is not conducive to long term research strategy. And generally, I felt very happy not to be involved in any of that. I've explicitly declared, I'm not interested in applications, whenever anyone asked me. If applications came along, I would happily pursue them. And later on, I've done several applications that arose from the biological work that we were doing. But I had no particular drive to do that sort of research.

Vassilis Galanos 46:06
Very interesting. A follow up on something you mentioned earlier. During your stay in Edinburgh, did you have the sort of conversations around political influences or misuses of AI? Could you envision where AI and robotics might end up in two or three decades?

John Hallam 46:29
I think no. I don't mean no, we didn't have the conversation, I think no, we realised it was practically impossible to envision where what we were doing might end up. You only have to look at computing to see that's impossible. If you think back to when computers were invented, which was during World War Two in secret. In the 70s, people still thought that the one maybe two computers would be enough for a country like Britain. Why would anyone need more? There was this story that when the AI department wanted to get its first computer, they were told no, they had to use the DEC 10 computer that everybody else used. And Robin Popplestone and a couple of others, changed the application to ask for a machine to transcribe from punch cards to tape. Because tape was much faster to read over the phone to send to the computer, and they got an LSI 11 which was of course a computer, and they wrote POP-2. They kind of got their first computer of their own by cheating a bit. And nowadays, you've got more computing power on your wrist in a watch than typically banks had in their mainframe room back in the day. And how would you foresee that? But we did have conversations about what AI might not do. And everybody got very cross with Reagan and his Star Wars initiative and refused to work on it. Pretty much the whole of the AI community said no, we're having nothing to do with this.

Vassilis Galanos 48:17
Interesting.

John Hallam 48:18
Because there was this famous story. I think it's a book about it called Always Another Moonrise. The American long-range defence system had almost launched a nuclear attack on Russia because it mistook the rising moon for a missile attack. The fact was that there was a person sitting there who had five minutes to decide, is this a real attack or not? And decided, no, it's not, the moon’s coming up. And, you know, imagine putting an AI agent in there that says, yeah, I'm fully convinced. Think about chatGPT here. Are you sure? Definitely, certainly an attack, they are really attacking, probably now. And then some poor human's got to push the button based on what this thing's telling him, he has about five seconds to decide. Now, you're going to put the humans in the position which we're now trying to put driv-

ers in, where your autonomous car says "panic!" and you have to wake up and assess the situation and decide what to do in about a quarter of a second because that's all it's left you, or else you have to have sat awake the whole boring trip, ready to do that. It's just not feasible from a technological viewpoint, you can't make the technology reliable enough that you could trust it. And from the human factors viewpoint, you can't give the people enough time and space to make good decisions. You've rushed them. So basically, everybody said no, this is an immoral use of technology, the technology is not ready yet and may never be ready and we're not going to do it. I kind of wish they decided that with autonomous driving as well.

Vassilis Galanos 50:14
It's hard. And there's lots of debate in the UK now around autonomous vehicles and cars and so on. It's interesting. You mentioned the Star Wars project. I wonder, I guess the equivalent to the UK is response to the fifth generation in the US was the Strategic Computing Initiative. I was wondering whether you had conversations around that as well. Did you see that as part of Reagan's politics?

John Hallam 50:39
Yes. No, it was the Strategic Defence Initiative, which was the Star Wars programme, that particularly worried people. I think having lots of money to develop interesting computing technology, that wasn't an issue. Because people viewed technology, I think much more neutrally. People worry much more now about what technology will do. But it's a very difficult question. There was a text in a newsgroup once about, somebody said, when you asked me, what do I think of mobile phones? He said, well, do you mean the mobile phone that enables me to communicate with anybody at any time, wherever I am, whatever I'm doing, to call for emergency help if I need it? And he went on for about 100 words. Or do you mean the mobile phone that interrupts me whenever I'm engaged in something important, and means my employer can get hold of me in the middle of the night? And the problem with technology is that people are in charge of it. And I don't think anyone's found a successful way of changing people.

Vassilis Galanos 52:18
Unfortunately. There's this more contemporary debate on the design principles and the design values we embed in technology. But maybe it's a bit late. I guess the work you did back in the 80s deciding on a structural change, so instead of following a biblical version of artificial intelligence, where you start with a human, and instead of that you get an evolutionary process, where you have a, say, design principle that might have played a huge role in that. And there's a question of what if, what if that started earlier? Do you have any thoughts on that?

John Hallam 53:01
I think we would probably have got a lot further if people had taken seriously just how much of human activity and experience is sensorimotor. And humans like to view themselves as the thinking creature. Because we're so good at sensing and acting, because it's so easy, we think it must be easy. And it's not. And I think one of the things we learned in AI is that introspection is a really bad idea for figuring out how things work. And the other thing that we've learned is that symbolic thought is actually the uppermost layer of the icing on the cake, and that the fundamental problem is surviving in the world. And if that had become apparent earlier. There's actually a long history of behaviour-based approach to technology, going back hundreds of years in fact, if you look for it. That was never pulled together really until, I guess, Rolf Pfeiffer and Christian Scheier wrote the Embodied Cognition book, which was also something of a manifesto. And I think if it started with that earlier, but then I think if we'd started by taking Descartes with a bigger pinch of salt, the whole of Western science might have been more produc-

tive. But one of the things you learn in studying animals and evolution is that things are contingent, you can only start from where you are. But that was the kind of thing we discussed in the 90s actually. It had lots and lots of robotics activity, activities in the robotics community. But nevertheless, with this kind of flavour. There were people like Rolf Pfeiffer at the AI Lab in Zurich, Luc Steels from VUB in Brussels, us from Edinburgh, the people from Sussex. Quite a lot of others as well in Europe, and then Rod Brooks, Dan Dennett. And we'd all get together and have great time kind of bashing around some of these ideas and trying to think about how to do AI properly. What were the real problems?

Vassilis Galanos 55:51
Was Luc Steels in Edinburgh when you were here?

John Hallam 55:56
I don't know. I don't think he was there, when I was there. At least if he was, I didn't know him then. I met him in Brussels in 1990 I think, we were writing an EU proposal with Rolf and Tim Smithers and me and Luc, and a couple of others. And we spent a week in VUB, actually doing the work, writing the proposal. Can't remember if we got the project, but it was very interesting.

Vassilis Galanos 56:27
Interesting. You mentioned embodied cognition. And that sparked the question in my mind, were you also in contact with the cognitive science people in Edinburgh?

John Hallam 56:39
To some extent, in that AI and Cognitive Science shared this School of Epistemics venture. And so we tended to be quite involved with each other. Henry Thompson, I think was the one who was explicitly on the boundary. But there was also Keith Stenning and others that we knew, because of a lot of shared activities. I particularly worked a lot with David Willshaw. We did a lot of MSc project supervisions together, and various bits of research we did together. We had a very similar view of things.

Vassilis Galanos 57:24
He was the first person I interviewed for this project. What are your memories with David Willshaw?

John Hallam 57:33
He's a lovely man. I really liked David.

Vassilis Galanos 57:36
Yeah.

John Hallam 57:38
We had a lot of fun. We had a whole range of Masters projects that were basically looking at Marr's theories of computational neuroscience. Because the nice thing about David Marr is that he was a neuroscientist at heart. And because of his ideas about computational theory, when he stated a theory about something, it was very easy to prove that he was wrong. Because he'd made explicit all the assumptions and all of the conclusions that he got from them. So they made really nice Masters projects, investigating his neural models of the cerebellum or the hippocampus, and seeing what they actually did, and whether they worked or not. And because we had more computing power than he did, we could stress the models more. And yes, you discover that he was genuinely wrong, but wrong in interesting ways. That's what you want from a good scientific theory. You want to be able to demonstrate

that it's wrong, ideally, and learn from doing so. David and I had a lot of fun doing it. We had some very good students, Geoff Goodhill was one of them. Tim [Taylor] was another that I shared with David at the Master's level, and then Tim did his PhD on artificial life. But yeah, a long and very, very interesting collaboration, I felt.

Vassilis Galanos 59:14
Did your research overlap at all the Willshaw Net and your work?

John Hallam 59:19
Not really, we used some of some of the connectionist ideas in robotics sometimes, but when they made sense.

Vassilis Galanos 59:26
When they made sense, yeah.

John Hallam 59:29
For instance, when we were doing the robot learning, we used connectionist learning models for that. But generally, no, we didn't do a lot of connectionist things. And I've never been completely sold on what you might call strong connectionist AI, which is that you can build real intelligence out of connected computing units. Maybe you, can maybe you can't, I don't know.

Vassilis Galanos 59:57
It's a problem that puzzles philosophers since many centuries

John Hallam 1:00:02
And one of the things that's clear from robotics is that thinking doesn't happen in your head. Some of it happens in your head, but a lot of it happens in your body. And the whole central nervous system's involved in thinking, and also there's a lot of your gut, there's a substantial amount of neural activity goes on just in your gut. And then the other thing that humans do is embed their minds in the world, like we stick up street signs, and we organise our spaces in ways that are convenient. One of the problems that I think dementia patients have, when they go from their home to a care place, is that part of their mind has been left behind, that was embedded. And typically, the mainstream AI people don't think that way. You're not the symbolic functionalist or the connectionist, or even the behaviour-based roboticist. The embodied cognition people come closest to this kind of view. So I've never been completely sold on connectionism, because I don't think thinking is a function of the nervous system, primarily thinking is a function of the organism.

Vassilis Galanos 1:00:02
And the social, I guess?

John Hallam 1:00:26
Yeah, and with humans, and also, to some extent with animals, it is the interaction with the environment, and that includes other agents, that drives what we might call thought. But that led to a whole lot of interesting work in the 90s on all sorts of random stuff. So I had students who looked at models of emotion in connection with reinforcement learning. That was interesting, because it turned out that I had a student who was very taken with Damasio's work. And he has this model of emotion that, that engages both the body and the brain as a coupled system. And we were interested in seeing whether

having such a system would interact well with a classic reinforcement learning approach to things, which again is another feature that you find in brains. And it turned out, curiously, that the emotion system didn't make the learning work any better in the sense that it learned just as well. But it made it learn much faster, because you could use it to focus on the episodes where you needed to focus to learn. That was quite an interesting outcome. And so there were a number of things like that.

Vassilis Galanos 1:02:53
Very interesting.

John Hallam 1:02:54
And I spent an amount of time being internal examiner for PhD students who were doing cognitive neuroscience work with David. And that was very educational too. I even went so far as to buy a human brain colouring-in book that I never coloured it in. I don't know if you know about those, if you're required to learn brain anatomy, then one of the ways that you can do it is by colouring in, and colouring in is a very good memorising aid. Because interestingly, because you physically engaged with the material. One of the pieces of research that I heard about which I liked the idea of very much, was a study that was done about whether it was better to have your documents on paper or on the computer. And the conclusion from the study was that if you just wanted to read the document, it made very little difference. But if you actually wanted to engage with the contents and think about it, it was much better to have held it in your hands. Humans think with their hands.

Vassilis Galanos 1:04:07
Yeah. I personally completely agree. But I've heard people claim the opposite, but I don't know.

John Hallam 1:04:19
The notion even turns up in Asimov's Foundation books, some of the later ones he wrote. The interface between the man and the spaceship is done with hands, and it's like telepathy with hands. The argument being that hands are the surface with which humans engage with the world. I would have to add mouths, because when you're tiny, everything goes in the mouth. But it's not all in here [the head], it's distributed throughout the body. And as I was saying, out in the world, as well.

Vassilis Galanos 1:04:50
A question I'm asking to all interviewees is, have you got any people in mind from the period you spent in Edinburgh that you think might have deserved more attention or more recognition? Lesser-known figures who, for example, stopped their PhDs for various reasons or left, or people, whose work you think is particularly important, but never received the citations in your view?

John Hallam 1:05:34
No, that's not a question I think I have any answer to.

Vassilis Galanos 1:05:39
No, it's fine. For me, I think most people tend to say, the same

John Hallam 1:05:45
It's not something I've really thought about.

Vassilis Galanos 1:05:50

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

I have a sort of thing about finding lesser-known figures. That sometimes history is kind of written by the winners, as they say. And all these factors like commercial success sometimes, tends to overshadow approaches. Yeah, interesting.

John Hallam 1:06:16
I suppose I think, Chris Malcolm's approach to assembly robotics was particularly interesting, and was fighting against a particularly heavy establishment. And I'm not sure what happened to it after the 90s. Maybe the obvious example is Bill Clocksin, who wrote the Prolog book with Chris Mellish. He was someone who left without a PhD. And then, I think he then went to work at Cambridge at the Computing Lab for many years. I certainly met him there two or three times. They kept wanting to invite me down to have a job, but I kept deciding not to move. In the end I moved to Denmark.

Vassilis Galanos 1:07:17
When did you move to Denmark?

John Hallam 1:07:20
Formally, I moved in 2003. I went on sabbatical in 2001. In fact, we travelled to Denmark two days after 9/11 which was interesting. I had a metal suitcase full of electronics and stuff. And we flew down to Stansted to fly to Copenhagen. And they said, Oh, you can't check the metal suitcase in, you have to take it to the special luggage place. So I took it the special luggage place. And they opened it and kind of stared at it in horror. And then one of them said, is that a laptop? So I said, Yes, because it was. Oh, that's okay then. Shut it and sent it off. And I thought, interesting that, so much for airline security. And I've never really believed in it since then. Strip searching grannies in the corridors and stuff.

Vassilis Galanos 1:08:22
Just a chilling effect. A very performative act, rather than security.

John Hallam 1:08:29
It's called Security Theatre.

Vassilis Galanos 1:08:31
Yes. I wasn't aware of the term.

John Hallam 1:08:35
It's one in the computer security literature. It's called Security Theatre. The classic example was an Australian aeroplane where they smelt petrol in the cabin. And aeroplanes don't run on petrol. They run on kerosene. So this was worrying. And the cabin crew went along and looked in all the overhead lockers. And in one, they found a chainsaw, which was leaking slightly. So they said, Whose is this chainsaw? And the guy says mine. And they said, Well, how did you get it through security? And he said, Well, you know, I showed it to them. And they looked at it, and they said, well, yeah, we probably should do something about that. But it's not on our list.

Vassilis Galanos 1:09:24
Yeah, we had extensive conversations with Tim Smithers on bureaucracy and formalities and how they impact AI research as well. What's to tick on various boxes, and I wonder whether these kinds of measures weren't really present in the past, like in the 80s.

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

John Hallam 1:09:49
No, not so much. It became necessary to go faster with PhDs towards the end of the 80s because the government decided to start deducting research funding from universities whose average completion rate exceeded four years, I think. And this was regarded as radical and extreme and unfriendly. But then for the last 20 years, I've been in Denmark, which has a completely stupid PhD programme. You have three years to do your PhD. If you don't complete it in three years, you fail. Of the three years, you spend six months doing random work for your department, and you spend six months doing largely pointless courses. So you have two years to do your research. And the worst part about it is that the structure of the programme is written in the law. So it's completely unchangeable.

Vassilis Galanos 1:10:54
Wow.

John Hallam 1:10:56
Yeah. Nevertheless, people do get PhDs in Denmark. But I think having anyone invent something radically new and world changing in a Danish PhD is effectively unlikely to happen, because anyone smart enough to do it isn't going to be stupid enough to be doing a PhD in the first place. Under those circumstances, they'll have gone off somewhere else with the proper programme. And a student who's smart enough to complete that PhD is nevertheless going to rely on having a reasonably well-defined programme of work to complete, because they've only got two years to do it.

Vassilis Galanos 1:11:41
I mean, I've done my MSc in Copenhagen. It's two years. So it's essentially effectively like an MSc. But with some more substance, some more research

John Hallam 1:11:54
There's two years of research with distractions, so you get three years to complete it. It turns out, you can get extensions for various circumstances, but nothing that changes the fundamental stupidity of the approach. And then the university that I was at generally took the least helpful approach to any ambiguity in the regulations.

Vassilis Galanos 1:12:25
Was the change you mentioned, towards, quality control in the UK, was it during Margaret Thatcher's?

John Hallam 1:12:37
Probably, probably part of her mythical economic miracle. I don't think it was about quality control, it was purely an accounting measure. They didn't want to pay. So people had to finish faster, so it cost less.

Vassilis Galanos 1:12:46
Interesting. So that I would think that is a type of challenge towards research, and AI as well. If you have this kind of benchmark measures, then your research, if you want to ask the big questions, like the ones you mentioned before, how and why. You know, you have this sort of barrier that says, you need to stop questioning this, and write up and so on. Are there any other similar challenges you think shaped your trajectory, or the field in general, especially in Edinburgh, like, political, financial?

John Hallam 1:13:40

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

I think not too much political, but certainly things that have shaped my trajectory in robotics. Around the end of the 90s, having done a lot of work in the kind of more biology/robotics intersection, it became pretty clear to me that, a lot of robotics publications were basically crap. In that they weren't adequately described. They weren't repeatable. There was no analysis of assumptions, failure modes, scope. People would tend to write a paper saying "Look, I did this, isn't it interesting?" And the correct response to such a paper is -- so what? One of the things that happened a few years earlier was that three of the "old men of robotics, Henrik Christensen, Ruediger Dillmann and Raja Chatila. Respectively from Denmark, or Sweden in fact, in those days, that was Henrik. Ruediger was from Germany. Raja was from France. They were drinking a beer in a bar in Japan after one of the International robotics conferences, and basically they said to each other, why is it that European robotics is so invisible? And they decided that part of the solution was to establish a European Robotics Research Network EURON. And Henrik called a meeting to discuss it, which was in Salford. And I think 100 people turned up. There was no funding for it. The people who turned up were the ones who really cared about it. I was one of the ones who went. And the result of that was that the European Robotics Research Network was founded. And then the European Commission got excited and paid for it, paid for the network for the next 10 years or so. But one of the side effects of this is that we then had a mailing list for European robotics researchers. And so I sent a mail to this list saying, has it occurred to you people that a lot of the work that's published in our field is very poor for these reasons. And my friend Herman Bruyninckx from Leuven also was involved and signed it too. And we got something like 100 responses saying, yes, what can we do about it? So I've spent a lot of time over the past 20 years trying to get people to do better methodology. And the reason that I'm saying this is because this is another casualty of funding pressure. And also commercial excitement. If you look at the average paper from Google, or Deep Mind, about whatever the latest piece of crap software they develop now is, it basically says, we implemented this horrendously complicated model and gee, look how well it works because the curve goes down. And that's the content of the paper in summary. And you think so what? One of my students who was doing his PhD here in Denmark on whether AI could be used in commercial robotics, actually had a set of slides, he took to one of these conferences, where he said, "a lot of people will say to you, well, you know, it's important to be pushing the field forward", and he put up a big road sign, like on the motorway, where it says "going fast, but nowhere, or going slow, but to a sensible place". And I think the problem that overtakes AI in this current commercial world, and perhaps the source of a lot of the hype that we've had in the past, and now more extremely so, is that people are not careful about their research methodology, for whatever reason. I think, largely, now it's a function of both commercial pressure and university political pressure that you need to have publications, because if you've got publications with citations, you're obviously doing a good job. One of the things that we talked about in methodology is we made a set of guidelines for writing methodologically-sound robotics papers, and one of the things we talked a lot about was metrics. So, the fundamental thing with metrics is, what is it you actually want to measure? And the base argument for this is that when someone reads a paper, they're not reading it to find out what you did. Or even whether it worked. The question that the reader of paper has in his mind is, will this technique work for me in my application? Is it worth my time investing in this to try it out? That's the question your paper should be answering. So what metrics do you need in order to help the reader decide that yes, this is a great technique, and it'll probably work. So I'm going to try it. And it became very clear that a lot of work suffers from the lamppost fallacy. I don't know if you know this one? It turns out it's a story from 14th century Sufi wisdom, where the guy Mullah Nasrudin, who's the hero of the Sufi stories, comes across his friend one night, and his friend's a bit drunk, and he's looking for something on the ground. And Mullah says, So what are you looking for? And he says, my car keys or whatever. I'll give you a hand he says, so they look a round for a while, they don't find it. And he says, but we haven't found it yet. So where did you lose it? And the guy says, oh, over there. And Mullah asked, so why are we looking here? Oh the light's better over here. And, you know this is precisely what we do. We make metrics that measure things that we can measure, because we can measure them, not because they're what we want to know, but because they're easy to measure. Citation index, citation count is precisely such a metric. And what we want to know is how much impact did your research have? That's very difficult to measure. So we measure citations instead, because we can do it. And you see this in research papers too. The loss function goes down. Well, yeah, it's easy to measure the loss function. But does it tell you whether the thing works or not? No.

Vassilis Galanos 1:20:55
And it also suffers from the accountant's symptom as well, in these things such as h-index, for example, sort of measuring the impact you had in the last three years? Exactly the amount of years you've mentioned earlier.

John Hallam 1:21:09
But accountants are classic driver of this, accountants want numerical indices that they can easily compute. And they largely don't care whether they reflect anything important or not. And the whole management science notion of KPIs, if done right, is a great idea. But typically, the KPIs that you want to measure are really hard to measure. And so rather than asking the question, how do we measure these really hard things? They ask instead, what can we measure? Let's measure that.

Vassilis Galanos 1:21:46
I read a very interesting book called The University in Ruins by Bill Readings from the mid-90s. He unfortunately died before the book was published, in a plane crash. And his best friend proofread and published the book. But he's situating that transition towards metrics in universities, and impact and research excellence in the mid-80s, essentially, so I think that was a crucial time for universities. Interesting. But you said you haven't really experienced that during your stay in Edinburgh?

John Hallam 1:22:28
No, I suspect it was much the same. I don't think the scientific community has changed dramatically. But exposure to biologists made me much more critical of robotics. One of the arguments a roboticist would advance is that these robots are very specific, and how can you generalise from the robot? Everybody has different robots. And eventually, I would say, well, you know, biologists manage. Every biologist has different rat, and yet they can draw general conclusions. Why can't we? Because we don't think about it enough. And in fact, all we need to do is borrow some of the techniques the biologists use, and we could go a lot further. And I don't think that the methodological problem is new, it's just laziness. But it's driven very much by these pressures to produce, and to meet criteria which are arbitrary, and are not measuring what you want measured. And the thing about the citation indices is that they produce people who optimise their citation index, they don't produce people who optimise the impact of their work, which is what you want. But it's not what you're going to get. And this must be well known in economics, but it seems to be very poorly understood in business and universities.

1:24:00
Sort of quantum effect of the instrument influencing the results.

John Hallam 1:24:06
Yeah, exactly. But, you know, tt's completely common sense. If somebody says, we're going to evaluate you by this metric, then you work to improve that metric. So it's very important that the metric is mea-

suring the thing that you really want as an outcome. And the difficulty is that the things that you really want are not straightforward to measure. Some of the most fun work I've done has been with computer games. Fun partly because I have no interest whatsoever in computer games, almost no interest. But one of my PhD students was interested in measuring whether people playing computer games were having fun or not. And that's, how do you define whether someone's having fun? The answer turned out to be not to define whether someone's having fun, but to present them with different versions of the game and say, okay, did you find this A or that B more fun, and collect an amount of data like that. Then you can build a model that relates the parameters of the game that you're changing to how much fun the person experienced. You've completely finessed the definition of fun, nobody cares what fun is. But you get a number out of this module that says how they would rank this in comparison to other versions. And then you can turn that back and use it to control the game. And it seems to me, that is a good example of how to measure something that you can't actually measure, dodge round it.

Vassilis Galanos 1:25:56
And again, the importance of the broader social aspect, maybe having fun or being entertained, is escaping and other social context[s] you want to escape [to]. Even if you describe it in precision, it doesn't sound like fun. You're getting killed in a video game, or you think about it a lot in a negative way. But still, it's an escapade from your work.

John Hallam 1:26:22
The thing that you really want to know is, is the player enjoying themselves? And I think the core insight that we came to in the work is, why don't we just ask them? But ask them in a way that we can then quantify their experience, to the extent that we need it to tweak the game parameters to make it more fun. That was a really nice set of research that we did, I really enjoyed that.

Vassilis Galanos 1:26:56
I hope you had fun.

John Hallam 1:27:00
We did.

Vassilis Galanos 1:27:03
Okay, so I do not want to abuse more of your time. I know we're nearing one half hours, maybe more.

John Hallam 1:27:10
Yes, it's twenty past three here.

Vassilis Galanos 1:27:14
Exactly. So I guess the last set of questions, or one big question is your vision towards the future, if you have any advice to give to younger generations who want to pursue a career in AI, computer science, robotics, bio robotics, all these sort of neighbouring fields. And something on last messages about Edinburgh, what was your reason for coming back every now and then? And any other thoughts you may have on this?

John Hallam 1:27:52
I think my advice to youngsters thinking of going into AI is not to be taken in by the hype. And not to be pushed into doing things that don't interest them. But pick some interesting area of human or animal

activity, and try and understand how it happens. If that's what you want to do. If you're more interested in applied work, then you pick an applied area and work on that. But try and work well, to produce work which is quantifiably sound. And don't let the bastards push you around. As far as Edinburgh goes, I was very happy there.

Vassilis Galanos 1:28:49
Do you follow up with contemporary achievements in Edinburgh?

John Hallam 1:28:59
Not really, no. I think my loyalty was always to the AI Department. And the politics and circumstances of the creation of Informatics were a little annoying. Also, it has to be said the Informatics came at a bad time, because I had then been in Edinburgh for 20 years. And it was the point when I had to decide whether I was going to be here for the rest of my career, or move somewhere else. And I chose to move somewhere else. The fact that I moved to Denmark was largely Bob Fisher's fault, or not that it was Denmark but that it wasn't in Britain, because I mentioned to him that I was thinking of moving and he said, but Edinburgh is clearly the best for AI and computing, where would you go? And I thought well, fuck that, I can go anywhere I like for the price of learning another language. And I never found languages particularly difficult, so I committed to going somewhere else in Europe. And at the time, of course, it was wonderful go somewhere else in Europe because Britain was a signatory of the Treaty of Rome. And we were full first-class citizens. Unfortunately, not so long ago they screwed all that up for us.

Vasilis Galanos 1:30:23
Yes, yes, interesting. And you chose a country where the language is only spoken within that country, so that's interesting.

John Hallam 1:30:37
Yes. Well originally I planned to go to Sweden because I was already fluent in Swedish. I had long interaction with Sweden over the 90s. Largely because of an accident: Alan Bundy was invited to a quite high-level workshop in Sweden on philosophy of science and he couldn't go, and he mailed and said does anyone else fancy going? Because I'll suggest you. And there was a silence, and I said, well, I could fancy going to Sweden. So we went to Sweden, and that was in 85 or so I think. Then in 90 I went there for a sabbatical and was nine months in Linköping, kind of building bridges between the computer science and the electrical engineering departments, and doing some behaviour-based AI. And when I was there I decided to learn Swedish, and got the university to find me a good Swedish teacher. And then I was back and forth doing examining and so forth over the next 10 years, and became fairly fluent in Swedish.

John Hallam 1:31:42
So the plan was to go to Sweden, and we'd organised the sabbatical in 2001 to go to KTH in Stockholm. And it turned out they couldn't get any money.

John Hallam 1:31:54
So I contacted a friend of mine, Henrik Hautop Lund, who had just become Professor here in Odense, and said, how would you like a visitor? And he said, yeah, sure, come, we'll sort something out. And we came to Odense, and we liked the place so much that we started to think seriously about moving. And

we've been thinking about moving, as I said, but then it became concrete. [...] I think that the thing that drove us towards Scandinavia was a comparison that became apparent to me, I think, when I was here on sabbatical. Imagine that you've got a really good idea for some programme of research or technology, and everyone agrees that this is a brilliant idea and it should certainly be done. In Scandinavia, they'll pay you to do it; in Britain, they'll give you 2/3 of the resources. And it's no harder to convince the Scandinavians than it is to convince the Brits of the quality of the idea. But the response is completely different. Scandinavia believes in education deeply. In Britain it's kind of an irritating cost centre that you can't avoid in society. And eventually I got fed up a bit. Twenty years of that was enough. And I thought, right. So there was lots of motivations, partly family, partly academic, but the result was we moved to Denmark.

Vasilis Galanos 1:35:21
This is extremely interesting. It's a very nice nonlinear way, I think to end our conversation. I think it's a great message.