

Sheffield, UK
[1:03:43]
m4v

Alan Winfield

An interview conducted by
Peter Asaro

March 25 2013

Q: Okay, so if you could just start by introducing yourself and tell us where you were born and where you grew up.

Alan Winfield: Sure, I'm Alan Winfield. I was born in Burton on Trent, which is in the midlands of England, in 1956.

Q: And where did you go to school?

Alan Winfield: Burton on Trent Grammar School, in fact, one of the last of the grammar schools in England.

Q: And university?

Alan Winfield: University of Hull in a – not far from where we are now. So, it's Yorkshire, East Yorkshire.

Q: What did you study there?

Alan Winfield: I studied electronic engineering. In fact, it was – I knew I wanted to do microelectronics. So, I looked for university courses that were specifically microelectronics. And there were very few at that time. So, I – at the time, I regarded anything as more than five volts and half an amp as pair electronics. So, I didn't want to do that.

Q: What made you decide to go into microelectronics?

Alan Winfield: Because ever since I was a kid, I was fascinated by electronics. I took things to pieces. So, from a young age, I was taking apart radios, and gramophones, record players, sometimes putting them back together again and was just utterly fascinated. In fact, I remember speaking to my careers teacher at school when it came to choosing what to do at university. And I was delighted when he said why don't you do electronics? I'd no idea that you could even do a university course in electronics.

Q: Did you go straight to graduate school after university?

Alan Winfield: I did, yes. So, I went straight from a Bachelors to a Ph.D..

Q: And where was that, in –?

Alan Winfield: It was also in Hull. I was very lucky in that I did a final year project in error correcting codes. In fact, my major in electronics was communications engineering. And it just

so happened that the school I was in had offered a prize of a Ph.D. scholarship to whoever came top of the class, if they wanted it. And because I'd loved final year project so much, then it felt very natural for me to continue in that, into a Ph.D.. So, I ended up doing a Ph.D. in digital communications – information theory and error correcting codes.

Q: Who was your advisor?

Alan Winfield: A guy called Rod Goodman, Rodney Goodman who's still – I'm still in touch with Rod, an outstanding Ph.D. supervisor and long-term kind of friend and a mentor. Yeah, I had a very, very happy experience as a graduate student. And in fact, I remember at an IEEE conference on neural networks, it must have been about I guess around 1979 or 1980, and the guest of honor was Claude Shannon. And I, being a rather brash Ph.D. student, I just marched up to him and introduced myself. We ended up having an interesting conversation. So, of course, it was only after that that all the professors said, "Oh, do you realize who that was? The guy's a legend. You can't just walk up to him and introduce yourself." Well, I did. So, that was a pleasure to – he's one of my scientific heroes. And I have to say that even now, as a roboticist, I'm always looking for the information theoretic kind of perspective.

Q: So, where did you go after your Ph.D.?

Alan Winfield: Well, in fact, I left academia and started a business. So, I became an entrepreneur. I, almost as a kind of side project during my Ph.D., I had invented a computer architecture, and patented it, and decided to see if we could commercialize this with – in fact, with my Ph.D. supervisor, and the third guy who was the financial guy. We started a spin-out company. So, I ran that for the best part of ten years.

Q: So, what was the company?

Alan Winfield: The company was initially called MetaForth Computer Systems. And we designed, at the time, the world's fastest single board processor. It was running just under seven million Forth instructions per second. It ran a computer language called Forth but at a machine level. Forth was the native instruction set for this computer. Commercially, it was not a success. But the company, in fact, continued and continues to this day not in computers, but in safety critical data communications. So, in a sense, the – my original specialism, which was communications was the thing that the company became successful in doing. And I'm very proud of the – I never – I didn't make any money from it, but I'm proud of the fact that it persisted and now employs about a hundred and twenty people. Yeah, that is nearly thirty years old.

Q: How did you get back into academia?

Alan Winfield: Yeah, basically after, as I say, running the business – well, in fact I stopped being managing director after about five years and then became technical director. But it was exhausting and I – by this time, I had a young family. And I really felt that I wanted to provide my family with a bit more security than was possible in this crazy, still young, startup company that was constantly having cash flow crises and so on. So, I decided to look for a job back in academia, and applied for and got the post that I essentially currently have at UWE Bristol.

Q: And what would you consider your first robotics project?

Alan Winfield: Yeah, I got into robotics really because of meeting two people. And between us, we founded what is now the Bristol Robotics Lab. Initially, it was called the Intelligent Autonomous Systems Lab. One was, or is Chris Melhuish who's now the director of the lab. But the third member of the three of us was Owen Holland. And Owen still is a great friend and really, I would say, provided the inspiration to both Chris and myself as – I mean I was not a roboticist twenty years ago, except I suppose in the sense of still harboring, if you like, science fiction inspired dreams of cybernetics and so on. So, I suppose there was some robotics in me but unrealized. And meeting Owen, in particular, really opened my eyes to the potential of this kind of new robotics, new wave robotics, which was bio inspired, biomimetic. You could almost say biological robotics. And I think that was true of Chris, as well.

There's another dimension, if you like, to the story, which is that I had been hired as head of research in the faculty of engineering, a faculty that really didn't have very much research. And my brief was to create a new research activity for the faculty. Now, I knew the communications research very well. In fact, just down the road at University of Bristol there was a big, still is a big coms lab. And I knew that there was no point in us creating a coms lab. So, I was looking for a research activity that was not represented in the region, in the southwest of England, in fact really not very well represented in the UK at all. So, it was a kind of happy meeting if you like. Chris and Owen, in particular, provided that inspiration from robotics. He was the only kind of one of the three of us who was already a kind of proper roboticist. And we needed – I needed strategically a new research direction. So, the common two things fitted.

Q: And when did that lab start?

Alan Winfield: 1993. Yep. So, this is its – this year is its twentieth anniversary. Yeah, and quite literally we started in a little kind of space that we managed to beg in the department. In fact, some of the first work I got involved in was designing and building little – our own mobile robots. I mean we – in those days, you couldn't buy research robots. You had to build your own. And so, I was really using – taking advantage, if you like, of my electronics background in embedded microprocessors and software engineering and all of that stuff. And because we have no space to do experimental work, we actually used the sort of space in the corridor outside the lab to do some swarm robotics experiments. And this was – in fact, the swarm work really came from Owen Holland who previously had already had been doing work with some swarm guys in

Belgium. So, and Chris Melhuish, in fact, ended up doing this for his Ph.D. work. And so, I got involved really in the – initially in the design of the electronics for the robots helping to design the electronics. And it was only really, as it were, during the next I suppose five years or so that I started to become deeply interested in the algorithms, the whole question of the nature of emergence and self-organization really fascinated me. And it still does. So, I – it was that that – I can't say there was a specific project. It was, in a sense, the work of the lab when there was only a handful of us found itself I think because of Owen's – Owen Holland's work. We found ourselves doing, if you like, some of the early work in swarm robotics. And I was hooked then.

Q: And what were some of the early questions and experiments that you were looking at?

Alan Winfield: Well, the early focus of that work was minimalist swarm robotics. So, we were really asking questions about what are the simplest robots in terms of their behaviors, both the physically simplest and behaviorally simplest, that could collectively do something interesting. And Owen did some of the pioneering work in clustering, puck pushing, where effectively a bunch of robots with only two or three rule – and these are identical robots. And they only have two or three rules, behavioral rules, would collectively sort and cluster little night lights, little candles and so on pushing them around. And so, the initial focus was very much on the minimalist kind of approach. So, I then ended up designing – co-designing with Owen the first Linux robots. We called them Linuxbots. And, in fact, we had collaborations with a number of other labs including the collaboration with Caltech, who – and we designed a number of – we supplied Caltech with a bunch of these Linux robots.

Q: When was that? And who are you collaborating with?

Alan Winfield: So, that was – in fact, by this time, Rod Goodman, my Ph.D. supervisor, was a professor at Caltech. And his post doc who I still work with to this day is – was Alcherio Martinoli. Alcherio is now at EPFL in Lausanne. So, we designed these robots, which are rather large. They're about the size of a dinner plate, robots. But I remember in about 1998, installing Linux on the robot. And we – and Owen Holland and I wrote a paper about these robots forming a wireless network. So, again, you see the communications kind of coming in. So, I think we were one of the first labs to actually have a wireless networked swarm of autonomous Linux robots. And it – but it was very clunky. I mean this was before Wi-Fi. So, we were using the kind of first generation of Wi-Fi. And the drivers hadn't been written for Linux so I ended up hacking – finding whatever kind of semi-working drivers I could find and hacking them and trying to make them work. So, it was fun. I mean it was pretty hardcore electronic engineering, as well. Yeah, yeah.

Q: Were there other collaborations during that period?

Alan Winfield: Sure, yes we – I mean, for instance, we collaborated with BAE systems, who also had a – well they still have, in fact, a lab in Bristol. And they were very interested in mobile

robot and we – again the same physical robots. They ended up with a bunch of those too. So, in fact, one of the – the Caltech collaboration was, at least part of that collaboration was, funded by a DARPA grant. And it was part of the – I think it was called the – or at least not – this wasn't the scientific name, but it was called the dog's nose project. So, it was effectively a project concerned with humanitarian de-mining. And I think the scientific question that we were trying to answer for DARPA was is there some advantage – imagine that you have a land mine. And it's unexploded, and it's leaking. Essentially, you have some vapor. So, it's essentially an odor tracking problem. So, the question is, is it better to have – or was, the question was is it better to have a single robot that tracks to the source of the chemical plume, rather like a dog would track by kind of going backwards and forwards and somehow honing in, or to have a bunch of robots that each sample the odor from wherever they are and then share the information wirelessly, and effectively, collectively compute the source of the odor. So, the scientific question was really very interesting. It was a kind of is it better to be collective or single robot. What's the – I ended up coming up with the notion of sort of collective gain, is there some kind of notion of collective gain.

Q: And what's the answer?

Alan Winfield: Well, it swings in roundabouts because of course the – you need to look carefully at the energy balance because obviously, collectively, a bunch of robots, although they might individually move less than a single robot that has to whiz all over the place, collectively of course you still might have five or ten robots instead of one. So, it was a qualified yes, I would say. But I recall a particularly interesting – it was a new experience for me, of course, going to a DARPA principle investigators meeting where you had to bring all of your robots. And they were – they had to do their stuff, and perform. But that was very, very interesting. Yeah, very interesting.

Q: When was that?

Alan Winfield: That was 1999, 2000, that period. Mm-hmm, yeah.

Q: So, what were some other of the more interesting swarm problems that you worked on?

Alan Winfield: More or less that time – excuse me, I – it's cold in here. More or less at that time, I became interested in mathematical modeling of swarm behavior. And I suppose – I suppose I – because this kind of biological robotics – so you have this extremely interesting kind of interface between the biologist, and in our case, the ant biologists who are kind of hypothesizing about the rules, the behaviors of the ants in order to do this amazing foraging and collective behaviors, and so on. And because there was the kind of biology and also artificial life thing going on – and I felt a little bit uncomfortable at the time. I remember feeling uncomfortable going to conferences – kind of simulation of adaptive behavior and the early a-life thing. And I was thinking well there isn't much engineering science going on here. There was far

too much of what I call proof by video, that someone would give a paper and show a really cool video of these robots doing stuff. And they'd say, "Well, that's it I've finished. Isn't it great?" But of course being a kind of old fashioned electronics engineer from having – and having worked in industry, particularly on kind of safety-critical production engineered systems, I really felt that there was a missing component, an engineering science component. And I remember writing a paper, which – I think it's pretty well cited now, where I was asking for a kind of new discipline of what I call swarm engineering. I was setting out the – a more disciplined methodology for engineering real world systems based on the principles of swarm intelligence. And part of that methodology – I was suggesting that we needed quite a number of new methods. And one of them was mathematical modeling.

And the guy that I mentioned, Alcherio Martinoli, he – his Ph.D. work was in mathematical modeling. I was very attracted by the methods that he'd been working on, and started to see if I could extend those approaches, and develop this kind of probabilistic mathematical modeling that he'd initially worked on. And not the only one, but there aren't very many people, there still are not very many people in the world who are interested in trying to mathematically model self-organizing systems. I think I was also attracted by the fact that this modeling almost abstracts away from the fact that these are robots. So, I was interested in the idea that you're not just doing robotics, you're doing science. And you're kind of doing – thinking about emergent and self-organizing systems in the abstract. So, I think that was another thing that attracted me to the mathematical modeling aspect.

Q: And when you got into that, did you find similarities to parallel search algorithms or other kinds of computational algorithms? Or were you really trying to look at say complexity analysis of these things or their optimization potentials?

Alan Winfield: Yeah we – there were several things we were trying to do. It – certainly there was an optimization aspect to this. We were – I mean essentially – but fundamentally, I was interested in trying to do, if you like, a validation exercise. Rather than just saying look here's a swarm, either a real swarm of robots – swarm of real robots or in simulation, look it works. What I was trying to say was yeah, but here's the math that actually validates these behaviors and shows mathematically how the macroscopic properties of the swarm are related to the microscopic behaviors of the individual robots. And, of course, as part of that, you're parameterizing. You're looking at the boundary conditions. And hence, you're able to optimize.

So, for instance, the most recent work I did with a Ph.D. student was looking at mathematically modeling an adaptive – a swarm doing adaptive foraging. So, in other words, the swarm was foraging with adaptive division of labor. So, this is a very interesting property of ants – not all, but some ants. That the – that ants – this surprised me when I first learned this. Some ants have a kind of labor pool. So, in the ant colony, there are some ants who are just milling around waiting for something to do. And it turns out that for these particular ants that if there is a lot of food in the environment, then more ants will be recruited from the labor pool to go

foraging, whereas if there's limited food in the environment, then more ants will stay at home in the nest and fewer will be foraging. And somehow the colony adapts the ratios of foragers to resters according to the amount of food forage, in the environment. And we weren't the first to develop foraging algorithms, but I think we were the first to mathematically model these algorithms. But what was, if you like, the icing on the cake was that we were able to put the mathematical model inside a genetic algorithm in order to do parameter optimization, which was something that's really hard to do in swarm systems.

Q: So, were you starting from the ants' algorithm? Or were you trying to prove a more efficient algorithm than what they used?

Alan Winfield: Yeah, that's a good question. I mean we – my observation of swarm intelligence research is that very rarely does one do swarm robotics researchers actually directly implement, if you like, the behaviors hypothesized by ant biologists. More often what you do is that you're inspired. So, you kind of – you have a rather abstract idea of how the ants may do this. I mean bearing in mind of course that we can't make robots like ants with pheromones and pheromone concentrations and so on. So, effectively, you're kind of inventing a robotics version, which doesn't really bear direct one to one comparison with the ant behaviors. And often, in fact, I've actually had this experience. The ant biologist will come along and say wow, you've come up with the same behavior, but with – the same collective behavior, but with fewer individual rules. So, in other words, we – the ant biologists may hypothesize a rather more complex ant than we end up demonstrating with the robots. Yeah.

Q: And did you go back to robots after mathematical modeling?

Alan Winfield: Yeah, I mean I've stayed in swarm robotics, but I've kind of diverged in – moved into several new directions. So, for one thing, I felt that swarm robotics has potential beyond, if you like, swarm intelligence. So, I became very interested in social intelligence. And I'm still convinced that if you want to build smart robots, really smart robots, it's crazy to just make one. The smartest animals on the planet as far as we know are human beings. And we're a profoundly social species. We wouldn't have all of this culture, and language, and art, and music, and science, and stuff if it were not for the fact that we are social animals. And it's always felt me as being – struck me as being very odd when people try and work on cognitive or learning robots, but they only ever make one. So, what I've been trying to go in the direction of is a higher level of cognition in swarm. So, in a sense, make the transition from swarm robotics to social robotics where – these are not social in the sense of human robot interaction but in the sense of robot robot interaction, which again is a little bit abstract but – so, for instance, I just finished last year a four and a half year project called the emergence of artificial culture in robot societies and what we did with that project was to implement social learning. In fact, learning by imitation, embodied imitation. So in other words, robots were able to watch each other's behaviors and then imitate them, but because it was embodied imitation, the robots imitated imperfectly. In fact for precisely the same reason that you and I imitate imperfectly. It's

because all we have is our own first-person perception. We have to infer each other's movement if we're trying to imitate a dance, for instance. I have to infer how you move your body if I want to copy your dance moves and inevitably I'll get it wrong.

So the interesting thing is that by implementing social learning, imitation learning in this way in real physical robots where the robots have to use their own senses and sensorium to do the imitation then you find that you have imperfect imitation, but because you have a swarm of robots and because they're imitating each other and of course you have robots imitating a behavior and then another robot imitates that behavior and then another one imitates that behavior then what we have is – well, firstly we have variation, which we get for free just through embodiment, we have heredity because we have multiple generation social learning, and in fact we also implement very simple selection. So a robot will imitate several behaviors of other robots, but will then decide which of those learned behaviors to enact. In fact, the simplest selection operator is simply to choose at random with equal probability which is surprisingly powerful, actually, is the selection operator. So what we did, of course when you have those three things you have the three Darwinian evolutionary operators and sure enough we demonstrated I think for the first time anywhere embodied memetic evolution and it was open-ended memetic evolution. So in other words, we couldn't predict where evolution would go and we certainly demonstrated some interesting, novel, new behaviors.

Q: Memes, huh?

Alan Winfield: Memes. Yeah, that's right.

Q: So what were they able to achieve. What kind of behaviors <inaudible>?

Alan Winfield: I mean we're talking about very simple behaviors. I mean the robots were raw, if you like, pathological imitators. That's all they do. So they just copy and then reenact and then copy and so on. So what we found was that we seeded behaviors and the behaviors in this case are the simplest case were just movement patterns. So we seeded like a little square move, these little wheeled robots. We're now onto E-pucks which are little miniature Swiss open source robots and so what we found is a number of very interesting things is we found that behaviors would emerge that were quite different to the original seeded behaviors, but they would become persistent in the population. So you could think of them as a kind of new behavioral fashion or a tradition. I think I prefer the word tradition and certainly we found that if for instance by happenstance you had a very poor quality imitation, so a behavior was very heavily distorted, there's a large degree of variation, but then if that highly distorted behavior was then copied with high fidelity and if this happened early in evolutionary history then you found that sometimes this very new and different dance became persistent. So what I'm saying is you saw the emergence of novelty just by virtue of a very poor fidelity imitation instance of social learning being copied by chance by a series of very high fidelity instances of social learning.

Another interesting thing, well firstly we found that the memetic evolution tended to drift towards sets of behaviors that somehow seemed to fit the sensorium and morphology of the robots. In other words, bearing in mind the robots have no sense of utility or value or ease with which the behaviors are learned, nevertheless those behaviors that seemed to be easier to learn and to pass on seemed to become dominant in the population. Perhaps the most interesting outcome or result was that we tested different memory strategies. So we tried these experiments with for instance no memory, with infinite memory, or with limited memory, and the interesting thing is that we found that with limited memory where a robot essentially has no choice but to forget the oldest of its learned behaviors, we found that the limited memory case gave rise to the smallest number of large clusters of related behaviors, which is very interesting. It seems to suggest that forgetting might play a role in the persistence of behavioral traditions.

Q: So traditions are more resilient with limited memory.

Alan Winfield: If you are to believe this. I mean this is a kind of very abstract but nevertheless embodied model and I've become very interested in the idea of robots as if you like and this is really referring to the Dennett, Dan Dennett's idea of a constructionist approach to science. So my journey in a sense has taken me from engineering, engineering science, and now to almost pure science in the sense that I'm now much more interested in using robots to try and understand the nature of intelligence, evolution. I haven't mentioned the work, but I'm also doing an evolutionary robotics and even culture by building working models. That's really the thing that's now grabbing my attention more than anything else, the idea of trying to understand what intelligence is by building very simple working models. I mean they're not actually simple. They're rather complicated with lots of robots doing interesting things, but nevertheless compared with animals, most of all humans, these are really very simple models.

Q: But you had to go back to the material robot, but not just the mathematical simulation.

Alan Winfield: Yes. I mean I haven't lost the engineering science aspect, but so analysis is still very important, but nevertheless, in a sense my journey in robotics is much more toward using robots as working models for – you could think of it as experimental philosophy.

Q: But you need the materiality to give you the right kind of variation.

Alan Winfield: Yes. I mean I'm perfectly convinced that doing all of this in simulation is not appropriate and I mean we are talking about complex systems. So with unexpected emergence. I've certainly seen lots of examples of unexpected emergence that firstly you would be really hard pressed to predict just theoretically or with pen and paper and you certainly wouldn't see in an agent based model or a simulation.

Q: So tell us a bit about your work in evolutionary robotics.

Alan Winfield: Sure. So again I've become fascinated by artificial evolution and evolutionary robotics, of course, has been around for a little while, but more recently the question of how can you evolve swarm behaviors, how can you evolve collective behaviors, and I became involved and still am involved in a European Union funded project so called future emerging technologies project called Symbrion which stands for Symbiotic Evolutionary Robot Organisms. So it's a different kind of direction. It's still arguably about modeling, but of a different sort. So here what we're trying to do in Symbrion is to build a swarm of robots that can act as a swarm independently. They can do interesting self-organizing collective behaviors of the kind that I've been discussing, but if the robots or if one of the robots, for instance, discovers that there's an obstacle in the environment that it cannot overcome on its own, then it can seed, if you like, the formation of a new organism that the robots therefore can autonomously physically self-assemble and then once they've self-assembled in two dimensions, they can then lift themselves up and form a kind of three-dimensional artificial organism that somehow can overcome. Imagine climbing over a wall that's too high for a single robot. Well, a bunch of the robots can self-assemble and stand up and then climb over the wall and then once they're over the wall, they can go back to being a swarm. So it's a sort of two-way process. So this is modeling of a different kind.

We're now modeling almost kind of the process of embryogenesis. So you could think of the individual robots of the swarm as stem cells. So when the organism forms and you have this kind of morphogenesis, the individual robots have to assume a particular function in the organism. Some become foot bots, some knee bots, and some sensory bots, and that's a deeply interesting question. How can you build a self-organizing system where the individual robots will differentiate, taking the analogy from cell differentiation, but again going back to evolutionary robotics, we're very interested in modeling the process of the formation of multicellularity. So I mean we haven't yet got this, but we're still working on the project, but my dream for this project, well our dream for this project – I'm not the author of all of these ideas. One or two of them, perhaps, but is that we can create an artificial environment where these robots are acting as a swarm. What environmental conditions would prompt them to evolve multicellularity? Under what conditions would it be better for these robots to form multicellular organisms than to remain as individual cells? That's a kind of pretty interesting question which really speaks to the origin of multicellularity. Yeah.

Q: So maybe you could just tell us about some of the Ph.D. students you've supervised and where they've gone off to.

Alan Winfield: So yeah, I mean in fact several of my early Ph.D. students are now themselves professors in the lab, which is really great. So I'm very, very proud of that fact. One of my Ph.D. students in fact was at the meeting. When I say former Ph.D. students was at the meeting on robot ethics that we were at today and he's now on the international standards organization committee defining the new standard in personal care robots. So I regard him as a –

Q: What's his name?

Alan Winfield: That's Chris Harper. So yes, Chris is, well, a very successful professional roboticist that I'm proud to have been his Ph.D. supervisor. But I have several Ph.D. students who one is now a post-doc and more or less immediately after graduating got a post-doc position in EPFL and he's still there. Another student who –

Q: Can you give us his name?

Alan Winfield: Oh sure. That's Julien Nembrini. So he was a Swiss mathematician who came to do swarm robotics with me and worked on those big Linux bots and very unruly robots, but did some really interesting work on kind of emergent morphology control. Sort of showed how you can get very interesting morphological variation in a swarm of robots just from single parameters that you vary a little bit. You get the difference between radial axial morphologies. Again almost hinting towards this kind of developmental biology kind of analogy. A much more recent student who worked on my artificial culture project that I mentioned earlier, he went back to Turkey and he now has his first junior faculty position as a roboticist in Turkey.

Q: And his name?

Alan Winfield: That's Mehmet Erbas, E-R-B-A-S.

Q: Okay, and what was it like working with Owen Holland?

Alan Winfield: Oh, an absolute pleasure and he and I still keep in touch and I regard Owen as one of the handful of most interesting people that I know on the planet and Owen is someone for whom it's impossible not to have an interesting conversation with.

Q: And who else influenced you theoretically, whether just in writing or through teaching or collaboration?

Alan Winfield: In robotics, well let me say in general, scientifically Arthur Koestler. One of the best science books I read years ago, his history of science, "The Sleepwalkers." I also read a number of his other books which still influence me a little bit. "The Act of Creation," for instance. I was inspired much more directly and almost personally by W. Grey Walter. Now there's some interesting history there because Grey Walter, as you may know, designed and built arguably the world's first electromechanical autonomous mobile robots and certainly again there were only two, but nevertheless, the first to do robot-robot interaction with Elmer and Elsie and his lab was very close to our lab in Bristol. It was literally just within a mile down the road. Now Owen Holland in I think about 1995 took it upon himself to track down the whereabouts of the six robot tortoises. Walter called them tortoises after Lewis Carroll because they taught us,

and he managed to discover the whereabouts of all of them and only – I can't remember. Two or three still existed. One in the Smithsonian, but he contacted Grey Walter's son Nicholas and discovered that Nicholas still has one of the tortoises in fact in a basement in Islington and persuaded Nicholas to lend the tortoise to the lab. So around 1995-'96 we had on loan this tortoise and Owen and some of the other guys in the lab, in particular the technician Ian Horsefield, got it working. So restored it to full working order, but still very fragile. I mean remember this robot was built in 1950. It was built basically for the festival of Britain 1951. So Grey Walters' technician, a guy called Bunny Warren who amazingly was still around in 1995-'96 for us to consult with and he gave us some original parts and so on.

So we not only were able to borrow this robot, but to make some replicas. So we built two as faithful as we could replicas and my brief involvement, apart from being inspired by these robots, was that I then became involved in – the lab was commissioned by – I kind of led the project to build some simulacra of these robots for the millennium dome exhibition, the millennium exhibition at the turn of the century and at that time I corresponded directly with Nicholas Walter who was still alive and had a very interesting correspondence and became fascinated, I think as anyone would be by that particular piece of robotics history. So Grey Walter definitely an inspiration. Remarkable that he more or less invented, I suppose you could say, in a minimalist robotics connectionist robotics behavior based you could argue approach and collective robotics. I mean all of that stuff which was not really rediscovered I think properly until Rodney Brooks in the 1980s.

Q: And did that shape how you structured your lab in Bristol?

Alan Winfield: It certainly influenced us. I mean for one thing it made us rather humble because wow, our marvelous robots with microprocessors really weren't much more capable or interesting initially than these 1949-1950 robots that have three vacuum tubes and two motors. Three vacuum tubes and two motors, but extraordinarily rich behaviors. I mean certainly I was very influenced – I think we all were by the richness of the emergent behaviors and by the again observation by Walter that the richness of the environment contributes massively to the complexity and richness of the behaviors of the robots. Yeah.

Q: So at the time you started the lab in Bristol, what other robotics labs were operating around England?

Alan Winfield: Well, perhaps the best known not so much robotics, but cognitive systems was the lab in Sussex, which is still of course the very well-known lab and also in Edinburgh and in Oxford. So I would say there were three very well established labs. I mean of course artificial intelligence lab started by Don Michie of course who worked with Alan Turing. I met Michie on a couple of occasions. Remarkable guy, really one of the founders of modern AI in the UK. So he was the founder of the Edinburgh AI lab which then became robotics and then of course the robotics lab of Michael Brady in Oxford which although Michael is retired, but that is still a very

strong lab, arguably responsible for inventing SLAM and a number of other groundbreaking developments in robotics. Yeah, so not many. I mean there are more robotics labs now in the UK, but not many big ones.

Q: What was Sussex known for?

Alan Winfield: Well, Sussex really known for cognitive systems in robotics and really for co-inventing evolutionary robotics. So I think Phil Husbands and Inman Harvey were kind of coming up with the ideas of evolutionary robotics more or less the same time as Dario Floreano and Stefano Nolfi. So there was obviously something in the air, I guess, but yes, really groundbreaking work in evolutionary robotics. I mean you're asking. I've kind of become a bit distracted by Grey Walter, but as an inspirational figure, but I've also been inspired by Susan Blackmore, the psychologist, for her work in both consciousness and memetics and certainly from a philosophical and methodological point of view strongly inspired by Daniel Dennett.

Q: Great. So the other question we kind of wind down with is for young people who are interested in a career in robotics, what kind of advice do you have for them?

Alan Winfield: Gosh. Well, I would say think about the science. Think about the big questions. Something that people sometimes ask me is "I want to do important work. What should I do?" Well, the answer is work on important problems. So even for someone who's starting out in robotics, I think it's a good thing to think about the longer term. Think about the big picture. Do you want to be the one who cracks artificial consciousness? Do you want to be the one who figures out how to put a swarm of robots on Mars that would build a human habitat? I would say don't limit your thinking to the problems of industry here and now. I think robotics is in a sense too important and too interesting, too fascinating, too far reaching, too inspirational to be, as it were, shackled by the economic needs of this or that company. I mean of course those are important things, but think longer term. And I would also say of course you need to do mathematics and programming and physics and you need to understand some principles I think to be a roboticist, but at the same time I would say think broadly.

Robotics is no longer a discipline of mechanical and electronic and software engineering. It stopped being just mechanical, electronics, and software engineering nearly 20 years ago. So a robotics lab like ours now has neuroscience. It has biochemistry. It has psychology and philosophy and material science and animal ethology and so don't think of robotics as a narrow traditional set of disciplines. It's now much broader than that and for that reason, read widely. Read evolutionary biology. Read philosophy. Read Dennett. Read Richard Dawkins. Read Stephen Jay Gould. Read about the theory of mind. Don't limit yourself to just becoming an engineer. Of course that's important, but think broadly because really the big questions in robotics and artificial intelligence are really very big questions indeed.

Q: Great. Is there anything you'd like to add or anything we missed?

Alan Winfield: Well, no. I think you've just about covered almost everything I can think of. So yeah.

Q: Great. Thank you very much.

Alan Winfield: You're very welcome.