

Flagstaff AZ
1:20:56
m4v

Ray Jarvis

An interview conducted by
Selma Šabanović
with
Matthew R. Francisco

August 30 2011

Q: The interview goes, as we start with where you were born, and your name, where you were born, and then go through the different places that you went to school, the kinds of things that you worked on, people that you worked with, institutions that you were a part of, that kind of thing, and then we end up with where you think robotics is going.

Ray Jarvis: Okay, now you said you had limitations, so –

Q: Yes, so our next one is at one-thirty.

Ray Jarvis: Okay. I just want to make sure, so I don't make it too –

Q: We'll keep an eye on the clock. I'm doing the times, and I'll try and make sure, yeah. So if we could start with your name, and where you – ?

Ray Jarvis: Okay, I'm Ray Jarvis, I was born in Rangoon, Burma, 9th of January, 1941. My parents had kind of British background, and some Asian background mixed. We came to Australia in 1947, landing in Freemantle on the west coast, and I went to a Christian Brothers school in a suburb near where we were, and finished up going to University of Western Australia, where I did my undergraduate degree in Electrical Engineering, and subsequently completed a PhD in Electrical Engineering. Following that in 1968, I spent two years at Purdue, in contact, somewhat with Professor Fu, who was one of my examiners. But I actually worked with Professor Ed Patrick, who was working on pattern recognition in those days. Went back to Australia to join the Australian National University as a senior lecturer, and was essentially given the task of setting up a computer science department, which began as an embryo group within the department of statistics, and a few years later, we became a separate department, but we were still, strangely, in the faculty of economics, which is a bit odd. So I worked at the Australian National University for 14 and a half years, went to a Chair of Electrical Engineering at Monash University, which is in Melbourne in 1985, changed the name of the department to include computer systems soon after, and established an intelligent robotics research center in 1987, and have been director of this group ever since. So the group, essentially included other academics from in the department with backgrounds in computer systems, with several different interests in robotics. So my colleagues had complimentary interests to mine, and we tended to work in a fairly democratic way, collecting a few students, roughly 10 to 12 between the four of us, expanding up to about 24 when the funding was good, and dying back to about ten when it wasn't so good.

My own robotics interests started off with combination of global search, as it was called in those days, and some pattern recognition background, drifting into image processing, and then computer vision and robotics. So the computer vision, robotics, about 1980 to '82, still while I

was at ANU, so when I established the center at Monash, my main interests were things like robotic hand-eye coordination, various range finding methods, stereo depth and then I switched between the hand-eye coordination to robotic navigation of vehicles, generally wheeled, sometimes tracked, and mixed indoors and outdoor experiments, so about halfway through, I suppose, my time at Monash, I got interested in larger outdoor vehicles that could operate in rough terrain. Those that sense the environment in sufficient way to carry out either tele-operated tasks or fully automated tasks. And I suppose the biggest project I worked on the last five years, we completed a year and a half ago, it was working with the country fire authority who were interested in robotic vehicles that could assist with putting out forest fires. And that work split up into several different approaches, one was vehicles that could go and do forward scouting, to check whether some area was safe, others that were capable of clearing the path for another vehicle, and then finally, the vehicle that, itself would carry water and extinguish flames. And the idea there was never to be too concerned about the degree of autonomous control, but just make sure that the human operator was at a safe distance. So the concentration there was on rich sense of feedback, and very simplistic control of existing mechanisms like steering, brake and accelerator, without modifying the vehicle more than absolutely necessary.

I guess another project which I had several goes at, was the – creating a semi-autonomous wheelchair, and the idea there was to allow for an adaptive shift between the wheelchair being fully under the control of the user, to taking over a certain amount of control when that user showed some signs of fragility, either from poor sight, or tremor, or tiredness, and I guess the key factor there was to be able to shift the responsibility between the user and the robotic equipment on board, so this vehicle was able to avoid collisions, but also give advice as to safe directions, and sometimes pull the – sort of the vector of intended direction toward the safer, and to change the parameters so you were switching between the fully autonomous to fully human mode. I had several goes at that. The first vehicle was simply one where you moved a joy stick and you had indications of where you could go safely and you were shifting around. The second version of that was where you had an eye-gaze tracker looking back at the user, so that a severely disabled person could go wherever they liked, simply by looking, so it was just where your eyeballs, not even your head, necessarily, you could keep your head still and look sideways, and that turned out to be very successful technically, but incredibly expensive, so nobody is going to work with that. Then I had a variant of that, which was a walking frame, that was motorized, that could convert into a wheelchair, again, with the same thing, and the last version is using connect sensor, which now brings the price down. One connect sensor feeding forward to work out the obstacle space, one feeding back to see what the user is doing for gestures.

So I guess I've mixed my interests between path planning, vision, robotic navigation and generally have been absolutely convinced that experimentation is the only way to go. I have developed a few theoretical ideas on clustering, on distance transform methodology, for navigation and convex hull, sort of geometric, some computational geometry. So I've done a few theoretical things that have solid founding, but most of my interests are making things work in physical space, and to show, I suppose, the relationship between concepts and reality and

adjustments that are required to take the existence in a cluttered space into account. My ongoing interests remain in the area of human machine interactions, and particularly in thinking about the different types of intelligences required for human-robot interaction.

And I had this idea that there were three types of intelligence, highly related, first, I call relational intelligence which is about how the robot can find its way around without bumping into things, in an efficient way, how it can recognize and handle objects. This is all about things in space, and what they are, and how to move. The second I call transactional. Now some of the computer scientists would call this dialogue. I call it transactional, because it's about the way the human and the robot communicated, but beyond that, how the robot came to understand what the human required, but within its constraints. So the negotiation took place to get the robot and the human to make a deal about what was going to be done, and there's a little bit more than just the gesture, or just the voice or just understanding of dialogue. Now the third level, I call social intelligence, and it's how to have robots mix with humans in such a way that the humans feel comfortable. So it's not just fulfilling the task efficiently, it's about having some notion of cultural limitations about proximity and approach and interruption and then timing, and also to include, perhaps, some aspects of, I suppose you call it entertainment and flexibility and a bit of softness in the whole approach, and I think that's the – one of the hardest areas, particularly, it's extremely hard to get funding in that area, because it seems too much like social science, but I think, in the end, we'll have to master that part as long with the other two, sort of harder layers.

So that's the work I'm doing now. I've sort of been semiretired and hopefully will get an emeritus status to stay linked with my equipment and colleagues at Monash, but my main interests are continuing this work about multiple robots in a human interactive environment, sensor rich, mainly interested in vision, and laser range finding, navigation, gesture recognition, voice and somehow protocols for human cultural limitations on behavior. So that's probably brings me up to date.

Q: That is a great story. So we're going to try to delve into some of those episodes of – into the different projects and some of the details of things you were interested and the kinds of technical and conceptual challenges that you were dealing with so if we could start with when you – well before we start with your PhD, I was curious, how did you decide you wanted to go into electrical engineering, in the first place?

Ray Jarvis: Oh, okay. I guess I was interested in making things. As a kid, I was always making things like projectors for film – for photography or models of things, and I always had this notion that I would like to be in a situation where I could make things, but I didn't really know whether it was bridges, electrical things or mechanical things. And I suppose I didn't decide that strongly until I was in my final year of high school. And then it seemed to me that electrical engineering appealed to me more, because it looked like it was going to explode into lots of different areas,

and as I got more and more into it, I think I was reasonably pleased with my choice, so I never really seriously considered moving sideways into mechanical, civil or chemical engineering. I went to a – the school I went to at university of Western Australia, had this common first year, which many engineering schools have, and then they kind of split up into specialties. But in fact, there were still some common courses that were run right up into my third year. It was actually a five year degree, so the specialization only came halfway through the fourth year, and then continued into the fifth. So I actually, looking back, enjoyed the mixtures of topics.

For example, in addition to electrical engineering, I learned a lot about mechanical design, about surveying, about strength of structures, material science, and a great variety of things, which I suppose at the time, we did not think were relevant, and in most areas of electrical engineering, would not have been relevant. But in robotics, happened to be incredibly relevant, so even knowing how surveying was done, was a huge benefit when you start to think about how to localize a robot in an environment, and the use of theodolites and lasers and all this kind of stuff. I also remember that we were taught to sketch things a lot, and I did technical drawing in high school as a special subject, which wasn't taught by my school. I actually went to night school and did technical drawing, because I like the idea of visualizing, so I think that's where I started to think as visualizing things as being the thing I really was interested in.

Q: Do you remember any things that you built when you were younger, that really stand out, like, any specific –?

Ray Jarvis: Well, I guess I did build a number of different versions of photographic enlargers and projections, and I built models of boats and things like that, and my brother went into chemistry, and he used to make explosive things, so I enjoyed making things that shot rocks out of pipes and things like that, slightly dangerous things, but I was generally interested in things that one could make with very few tools, and inexpensive bits and pieces. So I don't suppose there's any particular thing. I just made whatever I could with the things I could find around me, cheaply. So I didn't really know too much about how to go to a disposal store and buy things. I also had a sort of hobby interest in radios. I happened to have a next door neighbor at one stage who had an interest in this, and used to give me old radios and muck around with valves and make crystal sets and twiddle around trying to hear what was happening. So I had a general interest in gadgets of one kind or another, and then this interest in gadgetry has actually stood me in good stead throughout my career. So I think if any of my colleagues were asked, you know, what sort of captures what I've been interested in, I suppose you could say it's gadgets of one kind or another, which I've accumulated quite a few.

Q: So after you got to the university and during your EE degree, undergraduate, did you – were you working in the lab at all at the time, or was that in your –?

Ray Jarvis: No, there wasn't much opportunity to create things at that stage, but I do know that our engineering course was very, very lab based. We did many, many hours in the laboratory. In fact, for most of my classmates, the laboratory work was considered dull, but I actually enjoyed it, because I just liked to see how things worked, so if anything, I had a greater interest in the experimental laboratory work than I did in the more theoretical side. And I think that this is partly, of course, why I've always considered experimentation as the way to do things, although I perfectly happy to slot in theoreticians and people who have more philosophic views. I guess I'm just saying, yes, we need all that, but the thing I do best is experimentation.

Q: And when you got to your PhD, what kind of environment did you get into, and who were you working with?

Ray Jarvis: Okay, now when I started my PhD, I had recently developed an interest in computation. This is 1963, '64. The university at that stage, had a very small computer, it was an IBM 1620, and I, with the help of a friend, developed some programming to optimize the circuitry for switching circuits for binary switching circuits. That was my first introduction to actually using a computer, so as my PhD developed, I got interested in hybrid computing, which is combining analogue computation with digital computation. In those days, it was considered that, to build dynamic models of differential equations on a hybrid – well sorry, on an analogue computer, allowed it to simulate physical systems at, maybe thousands of times faster than you could with a digital computer in that day. But as digital computers became more powerful, then their need for the analogue side dwindles. And I was in that in between stage, where there was still benefit of combining two types, so I had an analogue computer with a link to the campus computer which upgraded from a poor old IBM 1620, to a Digital Equipment Corporation, DEC 10, which was apparently the second time shared machine made by this company, that was outside the US continent.

So the university of WA was the first outside the US to actually acquire one of these seemingly very advanced machines where multiple users could all clatter away on their teletypes and share memory and do their computations by either mixing with their colleagues during the day, and sometimes in blocking large chunks to be hands on, single user at night, to do the heavy computations. And I had a long cable between the physics school where the computer was housed, and the electrical engineering lab I was working in, about 200 yards apart. So part of the problem was how to get digital signals backwards and forwards without noise and this sort of stuff. So I built the interface between the two computers, I built models of different dynamic systems on the analogue, and I got interested in using the digital computer as the means of optimizing parameters, so that's where the optimization of parameters came from on a hybrid computer. So my thesis had a very long title, it was all about using a time shared computer and a hybrid system to do global optimization. In those days, there was always that – well people still today realize that if you have an optimization problem which is monotonic, that means there's only one valley if you're looking for the bottom of a valley, you could use steepest descents, slide

down the valley and so on, but as soon as you had multimodality, it was a different problem. So I was deliberately trying to work out what happens if you have a parameter space which has almost chaotic multimodality? And it turns out that you have to consider how to spread your testing as well as concentrate simultaneously.

And so that was called global optimization, and that's where I first got to know about King-sun Fu, was at Purdue, because one of his students had done a version of this, using a probabilistic method. So what you try and do is have a method that was able to randomly select where to sample, but also possibly had some local refinement going on simultaneously, so that was the idea of how to work on two levels at one time. So my thesis came out with this sort of randomized search methodology for global optimization. But it was still fairly theoretic in the sense that the types of multiple modal surfaces I was working on, were not from industry, they were just things I constructed myself. And then the question became, what if the shapes of these valleys shifted in time? So the idea was to have an ongoing real time process that could track shifts and particularly those shifts that were meant, but the valley that was the deepest went up on a valley over the hill went down, you don't have a direct path, and how do you keep a watching brief to – so a lot of it had to do with strategies for keeping a watching brief on what was happening in different parts of space, and then with some idea of not wasting time, because obviously if you've spent a long time on the hilltops looking at the next valley, you're not going to overall, minimize. So that was my thesis work, and then that's where I know about the work of King-sun Fu, so when I graduated in '68, I was attracted to go overseas, and it was fairly natural for me to say, okay, Purdue is a good place to go. I had never been outside Australia, except when we migrated initially, and so it was a great, kind of excited exploration to go to the United States and spend – I intended initially to spend a year, stayed two years, but ironically didn't actually work with King-sun Fu, worked with Ed Patrick instead. And that's where I got interested in pattern recognition.

Q: Who was your PhD advisor?

Ray Jarvis: Oh, it was Dr. Brian Leary. Leary, L-E-A-R-Y, who did not have a strong interest in my topic. But what he did have was faith that I could do it. So he would say, I'm not quite sure what you're doing, but I think you can do it. And it turned out to be the best attitude, because I didn't want too much advice. Quite frankly, I had an idea and I didn't want someone to divert me from what I wanted to do. So that part worked out well, because if I had had no real sort of drive, then the supervisor, not being specifically interested. So he was generous enough to say, I'm telling you now, I don't know much about it. I'm willing to act as your supervisor, because you need one, formally. I'll help where we can to get resources, but don't expect – and I took that to be very honest, and made the most of it, basically.

So that was good. So he would not be embarrassed to hear that I reported his non-specialty, because he was really up front about that, and I appreciated that at the time. But I find sometimes when I have my own students and they start drifting in areas where I'm not familiar with, I tell them outright, either to get advice from my colleagues who are specialized, or I tell them, look, if you want to move in that direction, I can't help you. It may be better for you to find another advisor. I think it's good to be up front, because it's not good for a student to believe his advisor is an expert, and find out later that his suggestions are almost wild guesses, and don't really reflect what's happening on in the field. And I think that's dishonest, so I'm – I really like the straightforward approach. So –

Q: And so when you got to Purdue, what was the situation there like, what kind of –?

Ray Jarvis: Okay, now what I found out was that Ed Patrick, who had a very strong theoretical background in pattern recognition, had a number of students working in this area. I got friendly with one of his students, and we worked together on a project or two, and one of them was to build an interface to their computer, to do some image processing. In those days, one didn't have sort of standard cameras, so this was a scanning method of collecting data and analyzing. So that was my first interest in the image processing side. And I also was very interested in various pattern recognition methodologies and Ed Patrick was very smart in sort of high level mathematics about probabilistic structures and the like, and he happened to have some good external grants. So it was really nice working with a group where there was funds for travel, and there was no concern about buying components, and things like that, so for me, it was rather nice. I did a little bit of lecturing as well, in computer logic, theory, and stuff like that, so it was a pretty nice time.

We wrote a couple of papers together, I got interested in an area called clustering, which is non-supervised pattern recognition, and I authored a paper, which I acknowledged his support as second author, on a clustering methodology, which turns out to be more useful than I ever imagined, because it was called, clustering using a similarity based on sharing near neighbors. Now the title actually tells you the method. So the method was just super simple. What you did is imagine all these points in, say, Euclidian space. You found, maybe K nearest neighbors, and K may be seven or ten for each point, and you looked at the neighborhood list of two points, and you said, these two points could be considered put together if they shared a lot of neighbors. It's a bit like saying, tell me your friends, you tell me your friends, if many of them are in common, you are probably friends. And it's a very, very simplistic computation, and it doesn't make any assumption about underlying probabilistics, so it's totally non-parametric, and by changing the size of the neighborhood and the thresholds, you could adapt to many different sorts of clusters, clusters that were long strings of points, they're sort of very globular and so on.

And then 20 years afterwards, and say, five years ago, I just looked up clustering, and I found everywhere they're referring to this JP algorithm, and I was, what the heck is JP algorithm? It was Jarvis Patrick algorithm, so I was very pleased to find it was in great use, and by pharmaceutical and molecular people who were working in the field on the edge of medicine and molecular chemistry. And they were using this as a standard method for clustering their different kinds of products and making this clustering available to their customers and the like. So it was a pity I didn't patent it, but you know, I was pleased to see it was used, and that's where the kind of mix came drifting out of global optimality into pattern recognition and image processing that finally led to my interest in computer vision and robotics. Yeah.

Q: Did you, while you were doing the clustering work, did you have a particular application for it in mind, or was it –?

Ray Jarvis: No, I think I was simply picking up the interest of the people around me. Patrick – King-sun Fu, was, international figure in the pattern recognition are at the time, so I think I was just trying to absorb the local specialties. At that time, I didn't really think about slight shifts, but certainly later on it was interesting, because my old interest in optimality and pattern recognition came together. So I picked up some of my thesis topics again, when I went to the ANU, and instead of just – oh, I guess it was mixed. So I guess instead of looking at the question of how do you find when the minimal points have shifted, in a way that you cannot track them, I developed this idea that you could have a pattern recognition algorithm that tried to detect when there was a sufficient change to warrant a re-exploration. So that's where the pattern recognition and the optimality came together. And then when I went to the ANU, I got more and more interested in collecting image data and analyzing it, so I drifted away from the pattern recognition side, but I did find applications for clustering in the image processing on several different occasions.

So probably that idea of mixing it came to its peak when I went back to Purdue on study leave. I kind of felt a bit guilty going back to the same place, but there were still some of the colleagues that were there before, and that was just a one year stint, so I went back, and this time, did work a little bit more with King-sun Fu and some of his students, and there the image processing and the pattern recognition came together quite strongly, and what I was interested in then, was what is now called image segmentation. So someone gives you a photograph, and you show it to someone else, and say, can you take a pencil and draw an outline of the major things in this – and then say, can you get a computer to do that? And it turned out, it turned out to be really difficult in those days, partly because they were computationally intense, took a long time to do, and I tried to use the idea of clustering, to group the pixels in the image into a conglomerations, and I developed a number of different strategies that were based on graph theory, basically, saying that you could imagine every pixel in this image as a leaf node, and what you're looking for is how to get the leaf nodes to conglomerate into larger and larger pieces, but not to cut across the boundaries of different objects, and the clustering helped there.

So by the time I went back to the ANU after that, I had this strong interest in the segmentation and the more and more concerned about how to do it faster, and you needed to do it faster if you wanted to use a robot, so I got interested in this idea of robotics as a means of testing the validity of the computer vision. And it was about that time I got interested in the 3D aspects. So until then, it was just 2D imagery. Now I visited, I think, SRI when Bob Bolles was there, I don't know if I met him that time, but he had some colleagues there, David Nitzan and others who had built a laser range finder. One of only one or two in the world at that stage, and I was fascinated by the use of time of flight laser ranging, and when I went back after that visit, I was determined to see if I could get some funding to build one in my lab in Canberrra, where the Australian National University is. And I applied for what's called an Australian Research Council grant, for mainly the equipment money. I had worked out how I could possibly buy off the shelf equipment, and I put together a budget, it was mainly about the costs of various things, and in those days, the Australian granting committee was split into smaller groups, and they often visited laboratories to talk to the applicants, and try and ferret out some of the background information and whether that was – it was in fact, a very good way of doing it, rather than just looking at what's written on paper.

And I met up with three of them, and they spoke to me very openly, and said, oh, well we partly came just to see your lab, but we are also doing this job for the Australian Research Council and we only think it's fair to let you know that we have got some opinions about your plan by some experts in – in that case, it was CSIRO which is a commonwealth research organization in Australia. I wasn't given any specific names, and they said, well, the experts there believe that what you want to build is not possible under – with this kind of, you know, under the plan you've got. And I thought, oh well, at least they're decent enough to tell me, don't have false expectations of getting the money, and had a nice chat, and they thanked me for the visit, and then they gave me the money. So now I was really under some pressure. They told me it couldn't be done. They said, have a go, and to my delight, I managed to get it working, and I used to boast this was the first laser range finder in the southern hemisphere, which it was in those days. That was sort of 1982, it was very really early, and that's where I really got interested in the third dimension, because that used to be the thing that stopped you doing the segmentation and the handling properly. There was no point just outlining figures on a photograph. What you really want to know how to recognize things, and to find where they were in space.

So about that time, I had this notion that robotic manipulation was perhaps the most honest way of testing whether your vision was operating correctly. Anyone could look at your sketches and say, hey, that looks like what I would do. But how do I know that that sketch or that segmentation has sufficient information to actually permit manipulation. And I thought, if you could actually show that manipulation was supported by this analysis, you then had a way of saying, even if you have objections to some of the philosophic ideas, I can demonstrate a practical outcome, and therefore, from an engineering point of view, that's a good and solid result. Then I thought, what I need is a robot. And again, with some funding, I had something like 20,000 dollars for a robot, and in those days, a decent robot such as the one made by Joe

Engelberger's staff, was called Puma, and this one was called Unimate, the small one is called Unimate 250. Those cost about 45 to 50,000 dollars, and I had 20,000 dollars.

The local agent for Unimation, I think they were called Unimation Incorporated in those days, said, why don't you ring Joe Engelberger? So with nothing to lose, I remember trying to pick my time, so it was sort of midafternoon, in the United States, and I spoke to Joe Engelberger, you know, the great – the father of robotics and all this kind of stuff, and he just said, oh, look, leave it with me. Maybe I can find a refurbished unit and send it to you. And indeed, good to his word, this box arrived with the Puma 250 at 20,000 dollars. So that was my first robot, and between the time I got it and the time I finally shifted to Monash, it was over a three year period, I started to do hand eye coordination. So I was using my laser range finder, and other methods, vision methods to allow robot to pick up blocks. So I scavenged a lot of the building blocks from my kids' playpens and stuff like that, and I had all these different colored blocks, and I could manipulate them and pick them off a table, and put it into different – and I was quite pleased with that. Again, the experimental side was seen as critical for me, rather than the more theoretical side.

Q: Were you in contact with the other people who were doing manipulation in Blocks World at the time?

Ray Jarvis: No, but I certainly – I knew the Blocks World Experiments in great detail, and in fact, taught them in my classes in computer vision, so I knew all the work of Winston and Guzman and the others, and also some work that was happening at SRI simultaneously. In fact, there were two schools of thought about how to do it. The SRI people were dealing with one aspect, and MIT was going in the other direction, and the two actually came together later. So I was very intrigued by the Blocks World Experiments, but I was really more interested in trying to exploit the physicality of the 3D analysis, rather than the notion of clustering vertices, which is the Guzman approach of sort of reasoning about block structure from the appearance of the edge data. I said, if I've got a 3D analysis of things physically existing, I can pick them up even before I know what they are, and what's more, I don't need to know which things are attached or not attached to other things, because, by picking them up, I'll find out. So it was – this is active – this is now called active vision. It's how to actually improve your quality by saying, the manipulator is allowed to change the world, and therefore, the change can expose things that you may not have known statically. So I was very keen on the idea of interacting.

And so a number of other projects that followed along this line, it was all about what I called post-recognition, manipulating first and then recognize. For example, picking up an object. I've got this object, now let me look at it. Now decide what it is, which is not what humans do, but from a robotic point of view, it seemed to be much simpler, because once you separated a physical object from a clutter, you now had it in isolation, and you could change the

viewpoint. So one of the projects we had when I had just established the Intelligent Robotics Research Center was, in fact, if your hand eye process by which we had a 3D scanner – this was not a laser scanner, but a stripe scanner, collecting 3D data, and then picking up unknown objects holding them up to the camera and then recognizing them. And that was pretty successful. And that project, about three years, with a few colleagues and so on, was one that we were hoping would be taken into an industrial application with a very large company we had as an industrial partner.

This company BHP, which is one of the largest mining companies in the world. They combined with a company called Billiton, BHP Billiton is quite huge. To give you some idea of the size, but of course they've grown since, that last year, they posted a profit of 22 billion dollars. So it's big stuff, and early days, they had a research laboratory about one and a half kilometers from my campus, so it was quite good to work with them. So we got interested in the hand eye coordination stuff, and also in mobile robotics, so they were talking about the possibility of using mobile robots in a steel plant, and mostly the sensors that we were developing, not so much either the demonstration with the hand eye coordination, or navigation, but we were developing a number of 3D sensors, stripe light and trying to speed things up and get quality. And, in fact, one of the outcomes of this research project, was a stripe light 3D system that was used in the steel mill to test the smoothness of the surface of a rolling mill sheet coming out. And BHP claimed that that application alone justified the money that went into the research project, so actually took one of, probably about five ideas, and at least used that within – in house. BHP at that time, did have a commercial wing, that was meant to take good ideas and commercialize outside, but they were a little bit hesitant, they didn't actually take this product, and I think that wing of BHP disappeared anyway, so we were a bit disappointed at what we thought were usable industrial ideas were not exploited, except for that particular range finder that was built.

Q: How much money was involved, if you remember?

Ray Jarvis: I think it was a million dollars over three years, so it was enough to employ about four people. Yeah, in the Australian context, that was considered quite large. And it was one of those very early, what was called a good – GIRD was Government Industrial Research Development, so it wasn't meant to be pure research, it was meant to be something you did with an industrial partner, and the hope was that by using public purse funding, you would encourage the industrial partner to exploit the market with the outcomes, and certainly BHP were one of the key industrial partners for quite a number of different projects.

Q: What year was this?

Ray Jarvis: Now that's a bit harder. I think around 19 – about I think it would be in about '91, '92, '93, about then. I'll have to double check, I'm not sure of the exact – yeah, it would be about then. About – a little later I joined one of the panels of the Australian Research Council, and I think that was about '92 to '96. So I think, yeah, this other project was before then. That was interesting, just to be a member of the panel, be like being on an NSF panel or some such thing, in America. So I got to learn, I guess what you could call the art of writing submissions, and how to do it and how not to do it with many, many examples. And I got good at writing proposals, I think, about that time. And so typically what this committee would do was collect 300 applications, farm them out to assessors, subdivide it into the panel members to act as the lead for discussion on maybe 100 samples each, and then incorporate our own responses with those of the referees, and come to a common idea about the quality of project, and then more or less democratically decide who was going to get the money. It was a fairly intense process, so the pre-reading maybe took a week or two, and then you'd meet together in a room, and you'd sit in that room for three or four days, and then take your results up to the next committee.

So I learned, I guess three basic things. One is, how to write a grant so that someone reading it gets the impact and the sense of it in the first page. That was something that's critical. We found some of the people who had very deep, narrow fields, tended to carry on about some theoretical issue that very few people were interested in, and you'd get through 20 pages, and you'd come out of that saying, sounded all right, but I still don't understand where it fits. So I got that idea that when I read many applications, and people make good points in the first couple of pages, I was much more positively disposed to reading the rest of the detail and fully understanding it. The second was, I really enjoyed working with those people, because they were all top people in their own field, and you got to know, I guess their mental processes, how they thought about things, how they approached things. So that was very – and the third thing was, I got to know a lot about the research community, because, in addition to seeing the proposals, we also did site visits to talk to people, saying, how's your project going, are there any bottlenecks? You know, have you the right sort of facilities, to do this work, and I found that useful, because you had a feel for what other people in the country were doing, where your own research fitted, and how, essentially, to tap the source.

So during the time you're on the committee, of course it's tricky, because your projects have to be – your proposals have to be sent to another committee who probably don't understand what's happening, to avoid that kind of in house, inside knowledge aspect. So that was a bit tricky, but fortunately, through that process, our projects didn't get lost in the system, and in 1996, I think, after I'd been – it must have been a little earlier. Yes, so I think I must have started in the ARC Committee, maybe '93, '94, '95, '96 or a little bit around – about '96, I was given a special grant, which was terrific, so what had happened is, I had written a project. That project was regarded as good. And so good that they said, here's some money, but you don't have to do that project, which is great, because now, you're saying, that's impressed us, here's some money, do whatever you like. And I thought, well I'm not going to work on that project, I'll keep that one, that's a good one, maybe I can use it again. And I used the funding for buying equipment,

basically, because I thought, I'll never have this opportunity again. So I had that flexibility for three years, I think it was '96, '97 and '98. Again, the dates are not all that clear. So I had three years in which I had complete freedom over, not a huge sum of money, but probably about quarter of a million dollars, buying things, so I tried to say, look, let me buy things that are going to be good for ten years, so I got lots of pieces of equipment that I used for 10 or 15 years.

Q: What kinds of things did you get?

Ray Jarvis: Robotic manipulators, vehicles different sensors, range sensors of various kinds, and so on. So I got myself kitted up, essentially, and was able to then do a lot of the experiment work from then on, so that was extremely valuable to have that freedom.

Q: I was curious, you mentioned that you got to see the lay of the land, what was going on in Australia in research? What were some of the topics and things that people were very interested in?

Ray Jarvis: Oh, okay, because the particular committee I was on, was looking at computer science, various parts of engineering, as well as some of the physical sciences, and you've got a quite a wide variety, so some things I remember was a very strong interest in one university's in solar cell manufacture, very early work on silicon cells, and they had pretty well the most efficient cells in the world for a while, and I remember feeling, isn't it too bad that the government isn't exploiting this? And it's early stuff. And I think, in fact they sold it overseas eventually. That's the sort of thing. There were other people doing pure artificial intelligence, you know, this sort of reasoning thing. I don't think the idea of Bayes Networks was there, but certainly there were people talking about more discreet, maybe extensions of the sort of stuff you saw in the Blocks World, but then Lisp was the big language at the time, so it was all about reasoning, about spaces that might have been involved with making decisions in industry, or selecting diagnostic stuff for drugs and the like, I remember that aspect.

There were also people working on micro-scale, I suppose VLSI circuitry at that time. That was big, everyone was walking around with – had been already walking around with a textbook, Carver-Mead textbook, was the bible for VLSI designs over a number of different projects building up in Australia, and that was expensive exercise, because unless you were getting your chips manufactured somewhere else, you had to have clean room facilities and it took a lot of expense. I remember that being one of the things I was quite impressed with at the time. And I suppose there were also, let me think. Yeah, I guess the solar work and the VLSI work, apart from robotics, were the ones that really impressed me. There's not – wasn't a huge amount of robotics happening in Australia in those days.

Q: What kind of things were people doing in robotics?

Ray Jarvis: The best known research project in robotics in those days were the sheep shearing project run by a friend of mine, James Trevelyan, T-R-E-V-Y-L-A-N, and in fact, rather annoyingly, for about ten years, whenever I was at an overseas conference, and someone picked up my Australian accent, they would say, are you James Trevelyan? And –

Q: You need a <inaudible>

Ray Jarvis: Are you James? So James was doing this sheep shearing project, and whilst this may sound very agricultural, it turned out that he was a master of mechanical design. In fact, you can ask any of the people like Osama Khatib, and others, and people know James as a very elegant designer of mechanisms. And so some of the things he designed for the shearing head on a robot arm were exquisite, and still have a huge respect for that work. And he wrote a book about this called *Shear Magic* I think it was, yeah. So that was what people understood to be the robotics impact in Australia about that time. There was – so my group started up at Monash, there was a little bit going on with, I guess a little later on, anyway, Peter Corke was in CSIRO in Queensland, so that group was growing. A little bit of work happening in New South – University of New South – but there were really only about five groups you could identify, each with three or four people in them, doing any robotics work.

And then the big impact was when Hugh Durrant-Whyte. You may have heard Durrant-Whyte took up his chair in Sydney University, and established the Australian field – AFI – Field Research Center, that was momentous, because Hugh came with a lot of very strong ideas, very, very strong industrial links to people doing port automation and mining, and built up his group to 40 or 50, 60, 70 people, in the end, and absorbed a huge amount of funding from both government and from industry, and in fact, in a way, other robotics groups, even mine, really didn't get a lot of look in, because he was so powerful in attracting the funding, and legitimately competing. So Hugh and I were hopefully going to set up a field research center together, but our first attempt had failed, because most of the money was going to big science. Our projects were not monolithic enough, right, so if you were going to spend ten million dollars, and you could buy one instrument for people doing gravity waves, or something, then you couldn't say, well, I could use it in smaller pieces, because then those other people would never get a look in. So Hugh and I were interested in writing that up, and we failed, and subsequently, when the call went out for special research centers, the money was relatively small, and I rang Hugh, and said, shall we go together? And he said, No, the money is too small, and we split, and what happened was that his group got the Center for Autonomous Systems, and we got a second ranking support for a center, but at half the funding. So we were called a sort of half center, but we did different sorts of things. So Hugh's center was big scale stuff, mining, port automation, then eventually into underwater and aerial on the big scale. Our stuff tended to be gadgets, small sensors of

various kinds. My colleague was interested in touch sensing and thermal sensing, others were in ultrasonics, and I was working in vision. So the things we did were really small scale, and Hugh was doing the big scale stuff. And the meanwhile, Peter Corke was starting to do big scale stuff as well. So I think Peter and Hugh worked together for a while. So the landscape changed quite a lot, to go toward the group that Hugh set up.

There's still some work happening in WA, and at University of New South Wales, and our Monash group had a reputation in sort of a limited size of operation. So Hugh was keen to use the phrase that Australian academic research community was punching beyond its weight, right? You know the term? So it means where a featherweight has moved up the scale. So it was true, because if you took the population and the small number, we were making quite a good impact. You know, we would find lots of Australians at international conferences, much higher than the proportionality would suggest. But nevertheless, we never got very strong recognition. Robotics has never been a national project in Australia. There have been funding for various particular things, but no one has ever come up and said, hey, this is where we really need to have – certainly there has been strong emphasis on alternative energy sources, on things like medical issues and mining and all this kind of stuff, but not specifically robotics. So given that we were all scrabbling for our funding from a bigger pool, having to compete against a bigger pool of relatively unrelated engineering and science, we were doing reasonably well.

Q: So what kind of research did you end up doing with your – you mentioned you had kitted yourself up at some point, and that was at Monash already there?

Ray Jarvis: Yeah, that's right, we were already there. So, I started – well because I was able to buy bigger pieces of equipment, I started to get interested in a lot of outdoor robotics. So two of the big things I bought both related to my visit to a laboratory in Finland. I can't, again, remember the exact date. What had happened was a colleague of mine, a Professor Aarne Halme, who was running quite a big laboratory in the University of Helsinki, had a conference that I helped with by reading some papers and became good friends, and he introduced me to some Russian engineers who had been part of the Russian space effort, and they had a large laboratory just outside of St. Petersburg. On this first visit, anyway, to Aarne Halme, there's a Swedish – no, no, there was a Finnish company that built a big tractor-like vehicle that was incredibly powerful and had a very low impact on the soil, and they were presenting it as a potential vehicle for robotic agriculture, because it didn't dig up the earth, and they were able to show that this thing could be driven with high precision and did lots of nice things. And I got interested in that machine. I finished up having one bought at slightly reduced price, around, I don't know, 50,000 dollars or thereabouts. And also I had visited the laboratory in St. Petersburg. It was just a few hundred miles train trip between Helsinki and St. Petersburg, and spoke with these Russian engineers, and my wife, who is half Russian, could speak Russian with them, and she thought it was wonderful.

So we did a deal over the making of a small scale Martian robot, called a Marsokhod. M-A-R-S-A – I've forgotten where the H is. Marsokhod you probably Google it, you'll find it. So I went there, did a deal with these engineers, and we finished up drinking vodka and eating sausages, about ten o'clock at night in the European – in the white nights area, it was light until 11 or something. And my wife had a wonderful time talking to these engineers in Russian. So it was a lovely sort of feeling, and then subsequently this robot was delivered without the motors, to Helsinki, where my colleague Aarne Halme fitted it up with Maxon motors, which were seen as the best at the time, and he then brought it as excess baggage to a conference he was attending in Australia, at my invitation. So in – not that big, it was this big. So this thing was something that was just fabulous, because it had an articulated body that was made of aluminium, and it had slightly conic wheels that had special serrated sort of fins on it, and this was made out of titanium and aluminium. And titanium cannot be welded, so every connection was a rivet, manmade rivet, and I was told that the Russians were really good at mechanical design, but it was better to get the instrumentation done elsewhere, so the combination was Russian equipment, and Finnish setting – setting up with the motors and the like. So that arrived and then I did a lot of work on that machine in a rough environment.

But this thing could climb objects as big as its own wheels, just go over the top, because it was like being able to articulate over the body, and I don't know if you spoke to Raja Chatila but he had a full scale one of these for many years in a project they called Eden, and he will tell you how wonderful this design was. He had a more sophisticated one that could extend the wheels as well, but I had a lot of fun with that over a number of years, so the big track machine, and the Martian rover with two of my key experiments and a lot of sensory stuff, and then trying to work out how to navigate, how to climb things, how to build models in the environment, and I extended some ideas I had on path planning, originally in about 1984, and used it over and over again in different circumstances. I may go on and talk about that. This methodology is based on an idea called distance transforms. Distance transforms started out as a way of analyzing binary images, and the two key people were Asrel Rosenfeld, and John Pfaltz, I think, were the University of Maryland, I don't know if they were there at the time, but this became famous stuff on how to analyze binary images.

And their interest was, if you're looking into a microscope, and you can see the outline of a cell, how would you make some measurements about the shape of that cell, and its volume and so on, and one of the ideas was to use what was called a grass fire, and when you light fire around the periphery and let it grow inwards, and you mark the path of this trajectory toward the inside, and that was called a distance transform, because if the outer layer was zero, the next layer would call one, two, three, and you got into the center, and you could do a lot of computation on those numbers, to work out things like the area, the amount of the – the ratio of the perimeter to the area and these sorts of things. And general shape, and you could actually take the skeletons, a kind of compressed version of this, and recreate the original and so on. So my idea was to turn the algorithm inside out, and say, how do I get this same type of propagation arrow to work on the outside of things which I called obstacles? So my breakthrough was turning

that algorithm inside out, and having something that propagated distance in such a way that it flowed around obstacles and it gave you a very, very simple way of global path planning. Because after you sort of had a start point, and grew distance out from that all around the obstacles, you now had all the free space, all the places the robot could go, with a number on it, saying precisely how many steps to the goal. And if you went down the quickest way, you were guaranteed a global solution.

So that was my idea in 1984, and I found you could extend it to any dimension. So I actually had a three dimensional one. In fact, one of the very, I think the first ISRR, I came, I presented the three dimensional version of that. And this algorithm has been reinvented several times, once or twice, even in my presence, you know, I've actually been in a room where someone has said, ah, I've discovered this, and fortunately some gallant colleague has gone up and said, hey, that was done ten years ago. So this algorithm, it turns out that you can extend it, so one of the extension that I did was with a student of mine, about five years ago, where I said, instead of finding paths simply from A to B, what if you had to go from A to B, but be unseen from most positions, most of the time, so what I call covert paths. So had applications in security, how to get from A to B, to be least detectable, and then, how to go from A to B, if you knew a sentry looking for you has entered that door, and is going to move probabilistically through the space, where should you go to find a hiding place, where should you go if you needed to escape from a hiding place, and then that extended onto, you are in a disaster zone, where should you look for victims of an earthquake, and again, how do you find paths that are weighted by the expectation of certain things. And it turns out, these paths can incorporate all sorts of other factors about being seen or not seen, potential for discovering things, and it also worked when you don't know everything, you could say, for the time being, I'll do this, and then I'll get a little bit more, and I can do that, and shift. So it's fast enough to keep on applying.

So this idea, I used for quite a few of my projects, and the simplicity, I guess I found appealing. I didn't really want to get into very, very complicated methods when this suited what I wanted. So even my outdoor work, this distance transform was one of the key underpinning path planners I was using. And I had some students working on related projects on this. So I'm still interested in the visual side, but for the outdoor projects, more in laser range finding than cameras, and the combination of how to build maps, how to plan, how to move. Now that was before this idea of SLAM had got any strength. The SLAM idea grew out of the need to go and explore new areas, and how to develop maps of new areas. And the only way you could do that effectively before that point, was to get out a theodolite or do a survey using standard equipment, moving very accurately from point to point, and taking readings of angles back to edges of buildings and objects and slowly but meticulously putting that together, and even in early surveying, this idea of closing the loop existed. So I remember mentioning doing surveying and stuff in my undergraduate, I remember there being a technique by which you started somewhere, you started to measure angles and distances, and when you went further and further round, the tendency would be small accumulations of error, because each depended on what you thought was your current location, that depended on your previous measurements and the idea in

surveying had always been that, when you came back to where you started, you'd say, this is meant to be zero, zero, but all my things are saying minus five, plus ten. So you knew you were in error, and then there was quite simple ways of distributing the error to make the corrections all the way along.

Now the current SLAM is really that technique, but at a much more sophisticated level, and probabilistic, the use of common filters and extended – it was all about how to distribute and propagate the error to minimize the – but you still have to complete the closure, because a lot of controversy is how you absolutely know for certain you're back where you started. You make that mistake, everything blows up. So a lot of work in SLAM is about that. Now by the time I sort of realized how important this area was, there – hundreds of people were already doing it, so I decided not to do it. So I went in the opposite direction, so I started doing what you could call the opposite of SLAM, so I bought a very – with a colleague, we bought a very expensive laser range finder, which you could put in the middle of a city square, and over a period of two hours, collect every detail up to 800 meters away, right, so you could just – the whole volume high precision, fractions of a – well maybe half a centimeter, and with a high resolution camera, get it all painted as well. So instead of building your maps incrementally, I said, well you just put this down, you collect everything at once, and you now have a sort of a cyber-model of that whole space. And if you could not see everything from one position, you could pick three or four positions, and then you can, with a – you can't automatically, but with a little bit of hand help, you can get the stitch together of all these things.

So we did this over a number of spaces around the campus and also sort of a property I finished up buying out of Melbourne to do some of this big field work, and so I sort of argued that in some circumstances, where you have either access to existing plans of buildings, or you're prepared to spend half a day collecting this information, if you're going to have a robot working there for ten years, the cost of collecting this is of no significance, so why do I need this complexity of a moving robot who is doing all this high calculations of common filters. So I just said, if you were prepared to do this once, and do it accurately, there was no need for that. So you know, that's – so I'm saying there's room for SLAM, but you could do it this way if you were being practical about an existing, relatively organized space. And then the notion was, okay, you've got this very exact map of things. Could you now put a robot in there, could find its way easily around this? And one of my students worked on the idea, you have a camera pointing up at a panoramic mirror, and it could see right around all in one hit. You could collect that image, and then you could hypothesize, where should I be to recreate that image from the cyber model. All right, so what you're really doing is saying, I've got an idea of what's there. Where must I be to get that same – so that was nice.

So you could actually say, I'm collecting this live data and I could do a search using particle filters to work out where must I be in this cyberspace for this same data to be visible, and that's where you are, and that worked pretty well. So you could just move this robot around with

its mirror and camera, and it could almost in real time, tell you where it was, and plot it fairly accurately. And it wasn't necessary, in fact, to have a robot, you could actually put it on a walking frame, and so this student actually bolted on a walking frame, and walked around like this, and you could see a map of where he'd been and so on. I thought, this is good. So later on, I got a second laser range finder, called a Velodyne, which is about the size of a human head, and can be bolted on the top of a vehicle. And that wasn't as accurate as the first one, which was called a Regal, if you want any details. A lot of people that bought these things. So what that could do was collect range data very, very quickly, up to 120 meters away, so it could collect one and a half million samples per second. So that gave the opportunity to knowing what's around you quite accurately, you know, ten times a second or something like that. You could spin it at about ten hertz. So it meant that you had live data from your vehicle, and it covered from about plus two degrees to minus 24, so it actually covered the volume that the vehicle may be running through, because you're not worried about overhanging branches above that, nor things, of course, under the ground.

So it allowed you to not only avoid things, but we used that range data to now hypothesize, where must I be in the cyber knowledge to get this set of range measures, and you didn't really have to use the 3D, you just use a slice. So you got the nice way of saying, I can avoid obstacles, but I can also – and we found that in an area of about 150 meters by 150 meters, you could find out where you were to 13 centimeters, in real – well ten times a second. So I'm still working on that, because I think, well, okay, if I have a vehicle in a bush land setting and I'm prepared to say nothing much has changed, of course one of the problems are, the trees grow, so, but if you use your scan about where the trunks are coming out of the brush, that's not going to change that much, as the growth above that point, so we found out, if you take something about a meter above the ground, that's reasonably stable. And I've been back to the same site six months apart, and it still works, and this sort of thing.

So I'm still keen to use this combination of pre-scan data and this fast scanner for localizing a robot, so I'm going back out to this property to do this sort of work. I guess that's probably a good story, too. I bought this property when I knew that I couldn't do this work on campus, because there weren't enough playing fields and stuff, and then students always come and gawk, and then it's dangerous, because this is a big robot, they could run over your feet, so there were all sorts of risk issues. So I got a bit frustrated, so I finished up buying a 20 acre property about three hours drive from Melbourne, it happens to be near some sort of foothills of some reasonably looking mountains and a fresh water lake nearby. It's a lovely spot, so we used it partly for – mostly for the work, but we also enjoyed going out there. And so the big work I finish up doing with the fire engines and all sorts of things were done out there, so what happened is that I went to the Country Fire Authority, and we had a project with them for five years, and so the big stuff was done out at this property, so in fact, even today, I have three large fire trucks out there.

Q: So you're not in danger of forest fires.

Ray Jarvis: No. The trouble is that I'm not sure that there's even petrol in the tanks. I've got to be careful. Yeah, they could be used for firefighting.

Q: So you know, maybe some of the folks in the last, maybe ten – because I think we got quite a bit early on. Maybe for the last 15 years.

Ray Jarvis: Well I had from 2003, to 2007, set up this center that I told you just slipped off the first rank when Hugh Durrant-Whyte got his center. And that was called The Center for Perceptive and Intelligent Machines, and I had partners in three other institutions. The partners in Melbourne were a group of people doing database analysis with artificial intelligence methodology, a group in Perth who were doing pattern recognition, and computer vision, and a particular individual working at the Australia National University who had spent 20 years studying vision in insects, particularly bumblebees. He was very well known for that work, and he had come up with some wonderful ideas about optical flow, which is a process by which you can work out three dimensional structure by working out how things move as you move through the environment. And he had picked out some analogies from insect vision to do this work. His name was Srinivasan, his – he's sort of a world leader in this field. He's since moved from – but anyway, it was his mirrors, in fact we finished up using, so amongst other things, he designed wonderful optics.

Now they were the people that we collaborated with, but the collaboration was very loose. It was very hard to bring those four different interests together on focus, and, in fact that was pretty much the downfall of our center, because the finances were subdivided too early in my opinion, but that was the deal that I had to do, and it was very hard to jog people away from their current research interests. They were individually doing extremely well, publishing well, getting recognition, so there probably was not sufficient incentive to be more centrally involved. My job was to actually try and stir up more and more collaboration. I guess, in the end, I have to say I failed, couldn't get this to work. So they're my – should be collaborators, rather than my – but most of my work I've done just with my students, and to some extent with my immediate colleagues. I hadn't, in fact, done a great deal of collaboration sort of on the ground, except in very, sort of ethereal terms of saying – talking to people at conferences and swapping a few things, so I found, because my work was highly experimental, it wasn't too good just swapping a theoretical paper. I actually needed people who had knowhow of how to build things, and right there on your particular device, so I didn't collaborate a great deal.

Q: So in terms of the one thing that Selma has been asking everybody, is in – because they're going to do a piece on education, too. What are some insights or some advice you might give people even at high school age?

Ray Jarvis: Oh yeah. I guess to start with, we found that there were many students in – undergraduates who developed a strong interest in robotics, partly because of a few demonstrations we could do, but also because typically computer systems people spend a lot of time sitting in front of a computer. And I think the prospect of actually seeing something move was kind of a refreshing thing. So we had no lack of interest, and for a while, I was teaching courses in pattern recognition, computer vision, automata theory, path planning, and they were usually a specialized group of students who were really interested in this. I was always excited of how students would come up with new ideas in small classes. So I was very lucky to have small, really interested classes, it made the job much easier.

Now some of our students who had an interest in graduate work, would come to us, but I think our group attracted more foreign students than locals, in the longer term. We had a lot of students come from – some Malaysian students, some Indian students, Iranian students, and some locals as well. So I think the second thing to say is, we told all our students, at least I told all my students, two things to remember, one is, don't read any literature until you've started to think about your project. That was controversial, because everyone said, hey, you've got to do a – I said, don't do it. I said, three things can happen. One, you'll find your ideas have been already done, that puts you off. The others, you find your ideas have never been done, but for good reason. And the third thing maybe that you feel that you've lost your ownership of something. I said, actually, it's better to think hard, ab initio, about a problem, and then when you've written down some ideas, then go and look at the literature, check it out, and that way you'll get either affirmation, or you'll be diverted. So the literature, I said, was actually corrosive. It would spoil your creativity. That was controversial but those students that did that, actually I think, got a lot out of that.

Q: Yeah, I could see that, because the other side is even just looking at a problem a certain way, and you start seeing – it could be just slightly –

Ray Jarvis: Yes, you may have found something new, and that newness is kind of gone, because you've been distracted. The second thing I think I told all my students is that in this center, we're only interested in theories that can be proven with physical experiment. We told them right from the beginning. Have good theories, but we want to see the thing move correctly, and that was always the, kind of motto of our group, in fact, all my colleagues believed in that, too. So that was good, because it meant that if you had someone who was just strong in computer science, but had done very little hobby type things, they got told, well you better find out how to make something, go and talk to the people in the machine shop, get to make your things, know how to put the wheels on and stuff like that, so our students generally were able to – were good at maintenance of their own projects, right, so that was the second thing, and I guess the third thing was, once you've got this idea and you've looked at the literature, sort of aim pretty high. There's no use just incrementing an idea by epsilon, you know, aim for something exciting, and even if you don't quite get there, well you'll get a better thesis than you otherwise would. So that's it.

Q: Yeah.

Ray Jarvis: Done?

Q: Sounds great, we really appreciate it.