

Pittsburgh, PA, USA
1:23:23
m4v

Chris Atkeson

An interview conducted by
Peter Asaro
with
Selma Šabanović

November 23 2010

Q: So, you can just start by telling us your name, where you were born and where you grew up.

Chris Atkeson: My name is Chris Atkeson. I was born in Washington D.C., our nation's capital in 1959, which is two years after Sputnik and I grew up in Washington, spent some time towards the end of the '60s in the Philippines and then moved back to Washington in the '70s. In the '80s I went to college at Harvard and then graduate school at MIT, actually Harvard was late '70s, so I was in Cambridge during that time. I moved to Georgia Tech in 1994 and I moved to Carnegie Mellon in honor of the millennium in the year 2000 and I've been here ever since.

Q: What did you study at Harvard?

Chris Atkeson: Biochemistry, I was a biochemistry major. I also got a master's in applied math. Harvard didn't acknowledge that computers were worth studying, so you had to sort of sneak it in in applied math at that time.

Q: And then when you went to the graduate school what department did you go to?

Chris Atkeson: When I entered it was called Psychology. When I graduated it was called Brain and Cognitive Sciences.

Q: And what did you specialize in, what did you study?

Chris Atkeson: I don't think they had internal divisions or titles or anything. My advisor was a neuroscientist, called Emilio Bizzi. He was interested in how the brain worked, obviously, but in particular how the brain controlled movement and how you coordinated movement. At that time robotics was starting to happen at MIT, I guess in a more visible way, and he reached out to two people, a guy named Neville Hogan in the Mechanical Engineering Department and a guy named John Hollerbach who was at the MIT AI lab and they collaborated trying to ask the question, "What can robotics tell us about how the brain works?" Because clearly the body is a robot, so the brain <laughs> has to be a robot controller and maybe some of the theory that's being developed for robot controllers could be useful, and that was the context in which I oper – I was in graduate school. First year I studied human movement, second year I – yeah, it's all coming back to me now – I <laughs> studied more human movement or human movement in a slightly different way, I'm happy to go into more detail on that kind of stuff later, and we also did some monkey studies and then at that point I decided, "Oh, let's go try to understand robots, they're a lot simpler," and that's what I did. So, my thesis was on the question of, "If you pick something up, how do you know how much it weighs?" The example I loved to use was, "When you reach into a refrigerator and pick up the milk, how do you know it's frozen?" Now, that's not a question I actually answered in my thesis, but that was a motivating question.

Q: So, what were some of the human movements that you studied early on? What kinds of movements?

Chris Atkeson: So, a big underlying question at that time was it seemed that robots or the way to make robots useful was to develop something called a general purpose controller, which meant that whatever you were asked to do you had some planning machinery that basically operated in the same way and just did it, and a signature of that kind of approach is no matter how you moved, you're always gonna do things in the same way, so maybe you move in straight lines or something like that. So, the obvious thing to do is to then look at humans and apply that kind of test and see when we ask humans to do different things, is there some commonality or are they coughing up some sort of special purpose solution to everything they do? And it turned out most of the time you ask people to move from point A to point B, and we developed some measurement system to do that, they move with a hand in a roughly straight line and more importantly with what's called a bell-shaped philosophy profile, so they start up in a particular way and they slow down symmetrically. So, it's a nice smooth movement. And that was actually an argument that whatever's going on in your brain, there seems to be a general purpose movement system, at least for move from point A to point B in a human with your arm. Later, the bell-shaped philosophy profile was extended to all sorts of things like jaw movements, tongue movements, leg movements, so it actually seemed that whatever's going on, it generalized across limbs as well.

Q: Okay, so what were some of the things you discovered in terms of the control structures or models that you then tried to apply to robotics?

Chris Atkeson: I'm going to focus on what I think is the last part of your question because I could parcel that as two questions, you know, how did the study of humans affect what we did with robot? At about that time there was a lot of emphasis on industrial robotics and doing things fast and so the idea was if you want to go from point A to point B, you want to hit your arm as hard as possible, turn the motors maximally on, accelerate maximally, get up to your maximum velocity, do what's called "coast," move at a constant speed, and then when you're close to your target you want to reverse the process and blast your arm with your motors in the opposite direction so you slow down as fast as possible and then you stop on a dime. Okay, the public image of robots at that time was that they move jerkily, like this <swishing sound> and if you ask anybody to imitate a robot, you'll see them do jerky movements and in fact I believe there's even a robot dance called "Do the Robot" which is all about that. That had nothing to do with robots, it had everything to do with this choice of movement strategy. Humans could move like that but they don't want to. They instead use this bell-shaped velocity profile, which means we move with a slow acceleration and in fact, let's see, there's something called position, you're changing position, there's something called velocity, which is how fast you're going, there's a derivative of velocity called acceleration, which is how fast you're changing your acceleration, then there's a derivative of acceleration. Guess what that's called? It's called "jerk," okay,

because when you change your acceleration rapidly you induce a lot of vibrations in the system and it makes it look like you have high jerk. What humans seem to do is whatever they're doing in their head, they seem to be moving as if they're minimizing the square of jerk. When you apply that to robots and get them to move in the same way, they look alive, they look like a human and they do not look jerky.

Q: So what's...?

Chris Atkeson: It's time to end this mindless discrimination against robots, they move just like we do, okay. <laughs>

Q: <laughs> If you can control the jerks. <laughs>

Chris Atkeson: If you just get the jerk under control, exactly, so to speak.

Q: <laughs>

Chris Atkeson: Okay, there's another joke I have to share with you. The higher derivatives, people ask the question, "Okay, if minimizing jerk is so good, let's minimize the derivative of jerk." You wanna know what that's called? That's called "snap." Can you guess what the two higher derivatives are? "Crackle" and "pop," how do you like that?

Q: <laughs>

Chris Atkeson: Okay, moving right along.

Q: Okay. <laughs> So, what was the first robot that you actually worked on?

Chris Atkeson: So, we actually have one of these still downstairs, it doesn't work, but it was a PUMA, the large size PUMA. I've forgotten what kind of number that is, it's like a, I don't know, maybe it's the 250 or maybe it's something, but each length was about that big and that's what I did and the first thing I made it do is I wanted to avoid having to have a hand. It didn't really have a hand, so I screwed on a weight to it and then I had it wave it around and it would tell me how much the thing weighed and what were its moments of inertia, in other words sort of what is its shape. It's sort of like going to a carnival and guessing your weight by having something grab you and shake you. So, that's what I did.

Q: And was that your thesis project or...

Chris Atkeson: It was part of the thesis. It turns out that we can take that idea of holding something in your hand and figuring out how much it weighed and then pretending that you're holding your forearm in your upper arm and we could move that around and you could now know how much your forearm weighs. We can do the same thing with the upper arm. So, and we were very interested in this idea that humans self-calibrate, right? You're born and no engineer measures you or anything like that, and somehow throughout your life you're always figuring out, "How much do I weigh? What does it take to move me around?" and so that's what we were doing with robots. We were strapping things onto them and moving them around, or just having them move around by themselves and having them tell us, "Here is a good model of how we work and then we can control ourselves more effectively."

Q: Were you particularly interested in making the robots be more humanlike or was there a different motivation for taking the humanism model for robot making?

Chris Atkeson: Okay, there's a basic split, there are two kinds of robotics people. There are people who want robots to be useful and there are people who want robots to be like humans, and those are pretty much disjoint sets of people and they fight with each other and I'm clearly on the side of, "I want to understand what it means to be a person," robots are a way for me to do that understanding and my idea of a good robot is it's like a person, behaves like a person, looks like a person. Isaac Asimov, and of course many, many other science fiction writers and movie generators, so, "Metropolis," the robot looks like a person. It's a lot easier, you have an actor be your robot so they naturally look like people, but a lot of these people believe that that's what a robot should look like and so they're on my side. In the 1970s I can remember as a graduate student, or I guess 1980s, a professor named Warren Seering in the Mechanical Engineering Department wrote an article for I guess the MIT alumni magazine, Technology Review, which basically said if you wanna build VCRs your robot should not look like a person, it should look like a VCR making device, and he was absolutely right but the politics of this message were really bad. It basically said, "Don't fund all these crazy people who want to build people or artificial people, fund people who are going to do something useful to society," because at that time the reason robotics was hot was because we were trying to deal with the Japanese and the perception that automation was making them much more competitive relative to us and our trade balance was therefore terrible. That fear has receded now and we have terrible trade balances with places like China, which has little or no automation, so there are other issues going on as well.

Q: Okay, so you mentioned your advisor...

Chris Atkeson: Emilio Bizzi yes.

Q: Emilio Bizzi. Okay, so your robotics work then got the degree in cognition or...

Chris Atkeson: Brain and cognitive sciences, that's right I was now accredited to go up to robots and say, "How are you feeling?" And if they said, "Well you're feeling bad," I probably, with a little extra work, could have been accredited to give them therapy.

Q: <laughs>

Chris Atkeson: People did do that in my department, they went out and got accreditation to give therapy, except to humans, not to robots.

Q: Alright. <laughs> So, were there other people in that department working with robotics at that time?

Chris Atkeson: In Emilio's group, yes, there were other graduate students. Tamar Flash, who's now in Israel and, you know, several others as well. I'd have to go through the records to actually figure out who was there.

Q: And were you working closely with people in the AI lab?

Chris Atkeson: So, John Hollerbach was in the AI lab and I was working pretty closely with him.

Q: But you were physically separated in a different lab?

Chris Atkeson: No, I picked myself up and walked across the railroad tracks to the AI lab, you know, I went wherever I needed to go.

Q: Okay.

Chris Atkeson: I was kinda like the guest who stayed too long, shows up, "What's that guy doing on the ninth floor?" You know, that kind of thing.

Q: And who were the people that were doing robotics at the AI lab at that time that you...?

Chris Atkeson: Well, there was a huge number, right. So, they had gone out and gotten funding to build a robotics group from I believe something called the Systems Dynamics Foundation and got a lot of money and there was a big effort at that time. To sort of track down the details of that you really should talk to Patrick Winston or Tomas Lozano-Perez, who are both still there, to sort of flush out that vision. I think Berthold Horn is still there but I think most everybody else has moved on to one place or another.

Q: But which ones did you interact with specifically from the people in the AI lab?

Chris Atkeson: Well, I had a fellow graduate student named Cheyenne [sp?] that I worked very closely with. I had some other graduate students, Dave Siegel [sp?] and SundarNarasimhan, I worked very closely with. I was a graduate student so I worked very closely <laughs> with other graduate students.

Q: Good. So, what did you do after you graduated?

Chris Atkeson: I became a professor, or assistant professor, at MIT in the Brain and Cognitive Science Department and also in the MIT Artificial Intelligence Lab.

Q: And so what was the kind of work that you started developing there when you were able to have your own lab and your own research direction? What kind of projects did you work on?

Chris Atkeson: I'm just pausing to think...since you're asking about things that were...

Q: <laughs>A long...

Chris Atkeson: ...a long time ago, 34 years. Okay, so again I'm motivated by things that we think are going on in humans. So, a lot of people have been very interested in something called content-addressable memory, and the idea is in a computer, if you shove things in the memory, it doesn't speak to you, you have to remember the box number that that thing went into if you're ever gonna get it back. It's like going to the lockers at the train station and shoving things into the lockers, if you don't know that you put your hat in A-52, you will never see that hat again. Worst case, what you have to do is open every door until you find your hat. There's a different kind of memory where you shove things into the lockers, but then you can speak to the lockers and say, "Hat" and A-52 opens, and the belief is that's a much better model of how biological memory works. So, we spent a lot of time thinking about different ways to implement that in engineering systems and how you would use it to improve control over a robot and there were other people who were using it in the general memory-based approach, or exemplar-based

approach or the instance-based approach, this was a very popular sort of meme at the time where they were using things like something called the connection machine, a big parallel computer, and they would just shove a lot of things into it and you would say, "I'd like to know about baseball," and it would figure out what memories are relevant to baseball and spit them out. The modern version of that, of course, is Google. Okay, so we were asking the question, "How could Google be useful for controlling robots?" about 30 years before there was a Google.

Q: And so what were some of the specific projects that you started working on to address this larger problem?

Chris Atkeson: Well, we would do simulations which always worked really nicely, and occasionally we'd actually try to drive a robot arm. I mean, I'm not sure I could fractionate it more than that, I mean those were the daily activities.

Q: Were you still using the PUMA or was it something...

Chris Atkeson: No. Okay, so as a graduate student I started with the PUMA but John Hollerbach was very interested in something called model-based control and the claim was what are called models of the robot, they essentially predict what forces to apply to make a particular motion, are too expensive to use in real time and he believed that was wrong and he believed that we could demonstrate that was wrong but we needed a better, higher quality robot arm to do it. At about the same time that CMU here was developing what's called a direct drive robot arm he thought, "That's a great idea, let's build one for ourselves," and that is the arm I used for a lot of studies.

Q: What was it called?

Chris Atkeson: We called it the MIT Direct Drive Arm. There might have been a "1" or a "2" after the name, I can't remember which.

Q: So, what were the differences in the direct drive and how did that allow you to do more precise movement?

Chris Atkeson: Okay, so this gets a little technical. The PUMA robot, and if you look at, for example, the Japanese humanoid robots of today, each joint has a motor and a gear or, essentially for our purposes, equivalently something called a harmonic drive and then the joint. And what that gear does is it has a gear ratio of a hundred or more and it's like your car or your bicycle where you put it in very low gear and you pedal really, really fast and the bicycle goes really,

really slowly or the car goes really, really slowly, and the big problem with a heavily geared arm is it's very stiff. It's easy for the motor to move the arm part, it's almost impossible for the arm part to move the motor. That is completely different from humans, I can grab your arm or shake your hand, whatever, and it's relatively easy to drive people's arms. We're back-drivable, whereas these humanoids and the PUMA is not. It also means that you don't really need to accurately control the forces or what the motor is doing because the gears, there's a lot of sort of slop in what's called the dynamics, so models are much less important. That's the advantage of heavily geared arms. That's one of the reasons that humanoids have been so successful in Japan, but if you want a robot that acts like a person, you have to get rid of those gears and so that's what we did, and then you can easily grab the arm and move it around but the downside is it's very weak because it doesn't have gears. So, we had to compensate for that by putting the thing on steroids and having really big motors. If you want I can show you, I still have the motors.

Q: Yeah, we'll take a look after.

Chris Atkeson: It's kind of frightening.

Q: <laughs> So, who were some of your first graduate students when you started teaching?

Chris Atkeson: <laughs> I should have brought my CV.

Q: <laughs>

Chris Atkeson: Man, this is embarrassing. Okay, I had a very successful undergraduate, David Reinkensmeyer, I think he was the first guy to sign on. He's now a professor in the University of California system, I think Riverside or Irvine, I'd have to go check. So, he does now prosthetics and rehabilitation, so he shared this interest in people and that's the direction he ended up going. I had a Chinese woman named Ying Zhao, she was interested in machine learning, which was another thing I was doing at the time and I jointly supervised her with a statistics professor. She did very well, I must remark that she basically defended her thesis, had a baby and moved to California basically all within a week which was kind of amazing, and she ended up working for IBM out in California. I had an Egyptian named Sherif Botros, he also was interested in how humans work. He ended up working for a consulting company in Boston. I had a guy, Eric Abouaf, but we did some early learning work, something called task-level learning with him. He ended up saying, "Okay, I'll take a master's" and went to business school and I imagine now he's making a lot of money but somehow he's forgotten <laughs> to send me the checks, so anyway. I really need to have my CV to get them all and I apologize to everybody I'm forgetting. I had two very successful postdocs, one is Andrew Moore and he was interested in machine learning. He was interested, came to MIT interested in these memory-based models and he

actually, after a couple of years with me, ended up as a professor here and now he's working for Google. So, you know, he is moving out into the real world. Another very successful postdoc was Stefan Schaal who you, I believe, talked to so I don't really need to say a lot there. He was also, I think, fundamentally motivated by how people work and his work now reflects that. I had a guy I co-supervised, Paul Viola, interested in sort of how we see, which you might ask the question, "What the hell am I doing advising somebody about that since that's not what I do?" but somehow it worked out. I think that's all the MIT folks, but again I really need to check my CV.

Q: Did you start incorporating vision into the system?

Chris Atkeson: No.

Q: No.

Chris Atkeson: We ruthlessly cheated. We needed to know where stuff was, but we couldn't wait around for the vision people to solve the vision problem because they never will. So, what we did is we marked things and there are two ways to do that. One is very simple, you paint things bright red or bright green or something, and then you look in the color channels of the camera and the bright red shows up in red and the bright green shows up in green, its vision is pretty simplified. I actually had an undergraduate who worked with me, she developed special purpose hardware to do that really fast and they went out and started a company, it was successful so that's nice. Another way you can mark things is bicycle safety tape. So, you'd wrap things with bicycle safety tape and you'd shine a light at it and it would reflect right back at you and the camera was right next to the light and so it would be really bright, and that's how a lot of motion capture systems they use to make movies like *Avatar* work like that. We used it in real time to do vision. And, you know, thirty-something years later they still haven't solved the vision problem and I still use those tricks to do vision, it's very sad.

Q: What other kind of sensors are you using to control movement?

Chris Atkeson: Sometimes I need to know when I'm touching something and we don't have a good robot skin and we don't have good robot tactile sensing, so a lot of my work has used what are called force-torque sensors. So, you put something in between your hand and the rest of the arm and that measures the forces in different directions and something called the torques in different directions, so it's called a 6-axis force-torque sensor. So, that's a sensor I use a lot.

Q: What about measuring the position of the arms and things like that?

Chris Atkeson: I'm sorry, I sort of took that for granted. After this is over I'll show you our current robot. It, like everything else, has a position sensor in each joint which, we work with a company called Sarcos which is now Sarcos-Raytheon in Utah, they love potentiometers so that's what it has. If you wanna be compliant it's often very useful to have force sensing on each motor or actuator, so we have that as well and I'll show you that. You have an inertial system and a bunch of gyros in your head and that's what makes you sick if you, you know, move in a way that doesn't match your vision, for example. So, we <laughs> not to make the robots sick, but to help the robots know where they are and which way is up, we put in the same sort of system, a vestibular system for the robot, so accelerometers and gyros.

Q: So, what were the first robots that used those or, did the ones that you were using <inaudible>?

Chris Atkeson: Everything I'm talking about has been around long before I existed.

Q: But when was the first time that you used those kinds of sensors in your robots? So, like...

Chris Atkeson: I mean the PUMA robot had whatever sensors it had...

Q: Okay.

Chris Atkeson: ...and we used a force-torque sensor on that. The force sensor on every actuator, it's an idea that's been around a long time, the humanoid we have now has it.

Q: So, what were some of the other applications that you were working on at MIT before you left? Like, what...

Chris Atkeson: Okay, like I said there are two kinds of people, people who wanna be useful and people who just wanna be human. As much as I'd love to do applications, here are the tasks that I do, paddle a ball up and down, throw a ball into a trashcan, that kind of stuff. So, I would hardly dignify that with "What applications am I gonna do?" I don't do anything.

Q: What about juggling?

Chris Atkeson: Okay. So, batting a ball is a simple form of juggling. Again, I wouldn't really call it an application, <laughs> it's not gonna make money or though, yeah, I suppose we could get it to make money. Yeah, so we did a couple things.

Q: Can you elaborate?

Chris Atkeson: <laughs>

Q: <laughs>

Chris Atkeson: Let's see, I'm trying to remember. So, I talked about the ball juggling, there was ideological warfare about the ball juggling, would you like to hear about that?

Q: Yeah, please.

Chris Atkeson: Okay. So, we did a ball juggler which was a classic example of sort of the – I'm gonna use classic twice in the same sentence – but a classic AI paradigm which is we see things, we reason, we do things. So, in the ball batting case the vision tracked the ball and said, "Okay it's gonna land over here, get over there, tilt your thing like this so you can get it back and then hit it. Oh, it's coming down here, move over here," okay? Now, we were very proud of ourselves and the system worked. It is true that if you had a crappy vision system and you guessed wrong as to where it was gonna go, you're gonna do something screwy and the ball's gonna fly away, true enough. So, there was a guy down at Yale named Dan Koditschek and he, I don't know how well, like, he's now a department head at UPenn or a dean or something like that, you can certainly track him down. How can I characterize his view? His view was, well he – we're interested in dynamic systems and this business about predicting where things are gonna go, that's really bad. It's a little hard for me to explain why he thought that was bad because it involved having a model of the world and all that kinda stuff. He'd rather have a system that had less internal cognitive abilities. So, he and his students developed a system where yes, you watched the ball, but you didn't think about how to hit it. Instead, you just, he talked about reflecting the ball through a mirror so as the ball comes down the hand comes up. And so his students developed this system and then they basically said to us, "Na nanananana," our thing is better.

Q: And how is it better?

Chris Atkeson: Well, you know, they claim first of all the formalism behind it is more attractive and I think there was an implied claim that it was more robust to noise. Well, we <laughs> had to deal with this so we sat and thought about it for a while and said, "You don't really need to see the ball at all because you have a ball coming down, if there's a ball that's too high, has too much energy, it'll come down later, right? If it has too little energy it'll come down sooner. So, what you wanna do is hit the balls that come down sooner, harder, and hit the balls that come down later, softer, and then they'll all have the same amount of energy and that's what

you want.” So, that says as you move up, you slow down. So, you’re going fast and the early balls get hit hard and then you slow down and the later balls get hit softly. So, we made our robot do that and voila, it could juggle balls without actually seeing them, it was totally blind. So, this was our retaliation and one of those guys from the Koditschek team came by, gave a talking to my team and we dragged him up to the lab, we said “Look, here’s our robot. You make a big deal out of juggling or batting two balls simultaneously, we can bat as many balls as we want simultaneously.” Okay, why? Because you don’t have to look at any of them, we just have separate paddles, a row of paddles, and then they’re all moving the same way and balls are bouncing on each one of them. So, we showed him juggling three balls and everybody had a good laugh and we moved on, but nevertheless <laughs> that was our little ideological warfare.

Q: That’s great. And did you actually learn anything about, was this outside of the view of how to understand human...?

Chris Atkeson: So, a deeper fundamental question is to what extent do biological systems choose strategies that minimize sensing versus what I call the classic strategy, you measure everything, typically with vision, you model it, you reason about it and then you do something. Whereas this, slow down, you don’t see anything, you don’t reason about anything, and the stuff works. If we were to sort of say what happened later that sort of fit that paradigm, there’s something called passive dynamic walking, which is building these walking robots that don’t measure anything, don’t reason, they just do the task. So, the deeper intellectual thread was how much sensing and reasoning do we need or want to do?

Q: What are some other uses of this manipulation and grasping, like where do you see this principle popping up again in your work?

Chris Atkeson: So, there have been people who have studied walking that looked at what’s called swing leg retraction. So, when your leg swings forward, they actually say that when you hit the ground you should be moving backwards for the same reason. If you come down early the leg’s out far and that slows you down more, if you come down late the leg’s here and that gives you more energy. So, that’s the same kind of phenomenon.

Q: Were there any other things of note that you did at MIT before you moved down to...?

Chris Atkeson: <laughs> It’s like asking, “How is your life summed up here?” you know, okay. Okay, I mean, I know you wanna hear this story so I’ll just tell it to you.

Chris Atkeson: So, we were interested in the control of a very high degree of freedom systems, okay, that makes them really complicated, hard to sense, hard to reason about, hard to push

around. So, cloth is a great example, and so I had a student get a robot to crack a whip because that's a very complicated, high dimensional system. If you do it right, it's kind of fun, you can break the sound barrier, make a lot of noise, and he did it. And then he came to me and he was saying, "Well how do I show this off?" And he said, "If you're in a circus you'd put a cigarette in your mouth and sort of sit there and whip that cigarette out of your mouth." <laughs> I said, "There's no way that I am going to permit you to put yourself anywhere near that thing. What kind of medical insurance do you have?" Okay, so we compromised, he wore a motorcycle helmet with a cigarette sticking out and the whip knocked it out. Somewhere I have that video, I'm not sure where it is at this point but it's a very amazing thing. Let's see, what the hell did I do at MIT? I think that's a pretty good summary of what I did. I mean, again...

Q: any other parlor tricks? <laughs>

Chris Atkeson: I could look, you know, in academics, right, I hand you my CV, it's my list of all my papers and there it is, that's what I did.

Q: Any other good parlor tricks you can recall?

Chris Atkeson: I need to get reminded of them more.

Q: Okay, and so when did you leave MIT?

Chris Atkeson: 1994.

Q: '94, and you went to Georgia Tech?

Chris Atkeson: Yes.

Q: What kind of projects were you working on there?

Chris Atkeson: Well, a similar sort of thing. We kept pushing to make this memory-based thing more real. There was an intellectual war about neural networks at the time. Neural networks were very popular, the statisticians were horrified because they're hard to analyze and formalize and I'm sitting here with these memory-based models and I think the fundamental property that the memory-based models have, which the neural networks do not is, there's something called interference. If I teach you how to play tennis and you get really good at it and then I teach you how to play squash, chances are you'll get confused, that's called interference

and your ability on tennis will go down. Well, memory-based systems don't have that problem, right, because they remember one, then the other and as long as you can figure out which task you're doing, your memories are all there. Neural networks suffer horribly from interference, train them over here, train them over here, they basically forget everything about here. So, I was fighting that fight at that time.

Q: So you were against the neural networks at that point?

Chris Atkeson: That's correct.

Q: Why did you decide to go from MIT to Georgia Tech?

Chris Atkeson: I got an ultimatum. My then girlfriend, future wife, got a job at Georgia Tech and said very clearly and slowly, "Either you come to Atlanta or forget it," so there you go.

Q: What were some of the differences between MIT and Georgia Tech at the time, in terms of the environment?

Chris Atkeson: It's not as hot, MIT is not as hot. I mean, they're very different places. MIT is, I think, a much bigger place, I might be wrong about – I mean in some ways Georgia Tech might be a bigger place but at least in the sort of the world that I was in, there was a whole lab of MIT AI lab for artificial intelligence and at Georgia Tech there was a group but it was like three people or four people or something. And the Computer Science Department was similarly much smaller, there was no comparable Brain and Cognitive Department, so it was a smaller place, absolutely.

Q: What was the robotics situation like when you got there and were there a lot of roboticists and robots running?

Chris Atkeson: There were a couple people. There was a guy in mechanical engineering, Wayne Book, who was doing a lot of stuff. Ron Arkin was doing a lot of stuff, but it wasn't organized. One of the big wins of the MIT AI lab was that there was – so MIT had four robotics or robotics relevant places, but the MIT AI lab was a big concentration. Now, ultimately CMU beats everybody with having a much bigger critical mass in the Robotics Institute but anyway, in terms of comparing MIT and Georgia Tech, you know, they had critical mass in the AI lab. Now, that dissipated later and you can ask them why that happened and interesting enough, Georgia Tech after I left said, "Hey, we got to really get our act together." They hired a guy named Henrik Christensen and now they have a robotics, I don't know what they call it, center or

entity or something. So, you point to the places that have made the commitment to robotics, Georgia Tech is definitely on the list.

Q: But it wasn't when you just got there.

Chris Atkeson: No, no, that was a long time ago.

Q: So, who was there when you got there? You said that there were a few people that were working in robotics.

Chris Atkeson: Well, there was Ron Arkin, my future wife, Jessica Hodgins was there. I'm sure there are other people I'm forgetting.

Q: Does she work in robotics also?

Chris Atkeson: She did, you're gonna really have to talk to her to get her story, but she was a student with a guy named Marc Raibert.

Q: Is she still working in robotics?

Chris Atkeson: I think the most accurate description is she's come back to it.

Q: Okay. And she's here on the faculty?

Chris Atkeson: Yes, her office is over there. She's in Europe right now, so.

Q: Okay. So, what did you start doing when you came to Georgia Tech? What were some of the projects you worked on?

Chris Atkeson: As an academic you just move, you don't really necessarily change what you do.

Q: Right, but what kinds of tasks did your robot do? Did you bring your robot with you? Did you have to get another one?

Chris Atkeson: I think I'm a little more abstracted, at different times I've had different robots, there isn't any "the robot" or, you know, a lot of the stuff I do in simulation. One of the things, in fact, one of the things that happened around the time I moved to Georgia Tech was I started spending a lot of time working in Japan and in fact what happened was that my collaborators in Japan invested in a lot of robotics stuff which I then worked with or helped them work with.

Q: Who were you collaborating with in Japan?

Chris Atkeson: A guy named Mitsuo Kawato at a place called ATR.

Q: Okay. And so you were working with the dynamic brain robot, or DB?

Chris Atkeson: Yes, <laughs> as a matter of fact I was. That was actually not the first robot we worked with. What we did is we went to this company, Sarcos, and said "Okay, how are we gonna get to the point where we have something humanlike?" which at that time was defined as two arms and a head and some kind of torso and Sarcos, I believe, wanted to reuse technology they had already developed. They had a robot that they had used for actually, I believe, power line repairs or something like that and it was pretty bulky and heavy but it did what we wanted and I think Stefan has it now. So, we started out with one arm and we used cameras to figure out where stuff was and we did a lot of good stuff with that and then DB was actually the next generation.

Q: And so did you have more of an input in terms of how to design DB to make it amenable to the kinds of things that you were studying, or...

Chris Atkeson: Have you ever worked with an artist?

Q: No.

Chris Atkeson: Or commissioned art for your house?

Q: I'm not that rich yet.

Chris Atkeson: Yeah, okay.

Q: <laughs>

Chris Atkeson: So, you walk up to the artist, you say, “You know, I really think it should be more pink.” Typically the artist says, “Yes, of course,” and then doesn’t do what you want. So, there is that element to our relationship with Sarcos that they had a vision, we certainly gave them input and then they tried to make the best of it that they could but they did have their own vision. And they were also basing it on an existing technology, it’s very expensive to make technology out of nothing.

Q: How did you initially develop your relationship with Sarcos?

Chris Atkeson: I inherited that from John Hollerbach, he had worked with Sarcos on something called the MIT Utah Hand . There may well have been previous projects before that but the MIT Utah Hand was, you should talk to Hollerbach and/or Steve Jacobsen about this, but it was an attempt to really push hand technology.

Q: And so Jacobsen was also...?

Chris Atkeson: He was head of Sarcos. It was his company.

Q: Right. Okay, and so DB was also, was he partly the artist that you’re talking about?

Chris Atkeson: Yes, I mean you should really talk to them. My understanding was that they had done what are called animation figures for a couple of companies and DB was a modified version of those animation figures, and after we’re done here I’ll take you over and I’ll show you some of the technology, help you understand where it came from.

Q: Mm-hm. So, what kinds of things did you get DB to do? I mean I’ve seen it, for example, playing – what are those things called?

Chris Atkeson: Air hockey.

Q: Yeah. <laughs>

Chris Atkeson: Okay, so that was a student of mine, Darrin Bentivegna.

Q: Mm-hm.

Chris Atkeson: I mean there's a whole list of stuff, if I had my laptop I'd start running videos. I was interested in tasks that people do, dynamic tasks. We had done ball batting and that was nice but when I say "juggling" that isn't what people think of. So, we got a step closer to that by doing what's called three ball juggling in what's called a shower pattern, a S, I always get this confused, <laughs> anyway it was symmetric and the balls were crossing over rather than the asymmetric pumping with one hand. Yeah, exactly. Now, there are some facts of life, one fact of life is hands are heavy and we didn't have any. So, I had to cheat on that, so I used kitchen funnels. So, it's not quite where I wanted to get to but it was close and, you know, the thing juggles three balls. We were interested again in the role of sensing, so we did it blind and that worked because of the kitchen funnels, and with vision and that also worked. Surprisingly it was more reliable blind than it was with vision, which was a surprise to us. There was another task called devil sticking, which is you hold two sticks and you sort of bat a center stick back and forth, which is a very nice task to study. It's clearly dynamic yet you never grab anything, you're always on the move, so we did that. Stefan did a whole host of things which I hope he told you about. There was some learning to balance pendulums, learning to do something called swing up, so there was an interest in learning. I think that's what I remember at this time.

Q: How did you meet Kawato?

Chris Atkeson: He's, I think, about five years older than I am. We share interests in how a robotics or an engineering approach can be used to sort of understand how the brain works and because we basically had that interest we met each other at conferences and things like that.

Q: What were some of the conferences that...?

Chris Atkeson: Oh man...

Q: ...you jointly attended?

Chris Atkeson: I believe, ah, it's so easy to be caught in an error here because the conference might not exist at the time. I believe there are two important conferences. One was NIPS, Neural Information Processing, that at the time was pretty much always in Utah at Snowbird, and then six months later there was something called Snowbird <laughs> which was a very similar conference. I don't know what it was really called and I think for a while both of us were going to that.

Q: Who were some of the other people that you collaborated with when you were at Georgia Tech?

Chris Atkeson: I'm trying to figure out how strongly to apply the word "collaborate." You have collegial interactions with all sorts of people, you know, you say things to them in the hall. I think in the robotics end of things I did my own thing. I had a huge collaboration with the Japanese folks, spent a lot of time there. Because I didn't have a lot of hardware at Georgia Tech and also I was kind of interested, I developed a new research area for me that you could argue is not robotics, or you could argue it is robotics, and that is, I think the simplest thing to call it is intelligent environments. It's mostly about sensing, it's about using sensing to figure out what people are doing and how you can help them, and I had a lot of collaborators there and in fact we ended up building a group. A key guy early on is Gregory Abowd and we had Aaron Bobick join us from MIT, Beth Mynatt, I'm spacing on some of the others but if you send me an email I'll give you a more complete list.

Q: What were some of the intelligent environment projects that you worked on?

Chris Atkeson: I think the main one, or the one with – okay, there were two main ones from my point of view. One, so I teach for a living, right, and I would like to make that work better and one of the things that's always amazed me is that a classroom lecture can hold people's attention, but that exact same lecture videotaped is incredibly boring. What's the deal? You know, do we need to smell people to be engaged? It was a total mystery to me. So I, in the end, went totally hog wild on let's capture everything, let's pump it out to the web, which was just happening at the time, and let's allow students <laughs> to skip lectures and then feel they were really there as best we can, this sort of max technological approach. So, the culmination of this effort, I would wear the robe of knowledge, which was a bright orange robe, it's like a highway safety vest, and that would make it easier later to figure out where I was and what I was doing. We would have four cameras in the room, some zoomed in on me, some zoomed in on the boards, we had electronic boards so everything was captured, and then I would wear cameras looking out so you could see what I was looking at and then there was a camera looking at my mouth so you could read my lips. Yeah, <laughs> and the amazing thing was people still came to class when this crazy guy was up there, you know, with sort of technology exploding. If you've watched the "Star Trek Deep Space Nine," it was like the Borg with the thing sticking out your ear and all this other stuff.

Q: Did people go to the virtual one?

Chris Atkeson: Yes, at least I think they did. You're asking a very interested evaluation question. To be wholly honest I was never interested in evaluation, I just wanted to do it. It turned out, actually, that people's parents would go to the web and use the camera on my head to figure out if their kid was in class and if they're not then bad things would happen.

Q: So you opened a different kind of privacy related, <laughs> a different can of worms.

Chris Atkeson: It had never occurred to me that that was an issue.

Q: But if you were interested in this issue of presence and what makes people feel like they're –

Chris Atkeson: Really there.

Q: Right. They're really there. I mean did you evaluate it in any way?

Chris Atkeson: Oh, I had other people I collaborated – that's what collaborators are for, right?

Q: Okay, uh-huh. So what happened?

Chris Atkeson: What do you mean what happened? Did it work?

Q: Was it any good?

Chris Atkeson: Everybody was very enthusiastic. Yeah, of course it worked. It's doomed to work. I'm very cynical about this kind of thing.

Q: Okay. Yeah, no, because it seems to me like it would be this huge information overload.

Chris Atkeson: Oh no. Everything's cross-linked, right? So you could sit there. It's like the old movies where the little dot is bouncing on the words and you sing along. You're just watching a lecture and you've got a video here and the guy's talking. You got zoomed in on what he's writing and the little words are skipping by. You can make a very integrative presentation. It doesn't smell and maybe it's the element of fear. At any moment the lecturer could keel over. That's what makes it exciting. Or they can get angry or do something unpredictable. I don't know.

Q: There's also the other people. Ask you a question.

Chris Atkeson: Really? They shame you into staying awake? Yeah, maybe. Maybe. You have teams of people watching this web thing. Okay, so that was one thing. The second big project, and this was actually a really big project, is we built a house and it was very fortuitous. The people interested in wifi and radio were very interested in having a test bed that would allow them to figure out how to design better wifi and radio for computer networks and we latched

onto that and said, well, this would be a fantastic place to do experiments on people with sort of futuristic technology, how they reacted to it and stuff like that, and if we could get rid of the lawyers, we could even have them live there. It turns out the liability issues go through the roof if they spend the night. So I don't think –

During the time I was there, we never had anyone legally live in that house. We had a lot of people live there illegally, but nobody legally lived there.

Q: What was the house called?

Chris Atkeson: The Aware Home.

Q: So was this one of the first living labs? It's a type of living lab.

Chris Atkeson: Actually if you look historically, there have been a lot of people who've built houses for research and if you're willing to include world expos and stuff like that. People build houses to demo stuff all the time. There's been a long history of solar power houses and stuff like that. So I'd be very reluctant to say this was the first of anything. I think at about that time I think a lot of universities jumped on the let's explore how futuristic technology can help old people and I have no idea. Who was first? I have no idea.

Q: And so your focus there was also on the elderly.

Chris Atkeson: Yes, yes. It was very explicitly we're going to help improve the lives of old people.

Q: What were some of the sensors and technologies that you incorporated in that lab?

Chris Atkeson: So I spent a lot of time building infrastructure like how would we put a hundred cameras in a space like that and collect all the data and things like that. We had vision people who would actually do the processing on the data. We had a set of people interested in interfaces and actually probably doing very simple things that turned out to be very useful. So one thing that'll get you into a nursing home really fast is if your adult children lose faith that you are safe, okay? Like you burn a hole in the floor. So the house I bought here had a big burnt hole in the floor and very clearly the guy who lived in it, he was a widower, he lived alone. I don't know how old he was. He was clearly shipped off to a nursing home the minute that happened. So a very simple thing you can do is have a house that reports to the adult children things are okay.

Mom got up this morning. Mom seems to be doing the normal pattern of room-to-room movement and stuff like that. You can go a step up from that. You've seen on TV the life alerts which rely on you have to press a button and then help comes. Often old people are too confused or too injured to actually press that button. Maybe the house could press the button for you. It's like OnStar.

Q: So were there particular insights that you got?

Chris Atkeson: Did we learn anything?

Q: Did you learn anything about people?

Chris Atkeson: After spending approximately a million dollars of the Georgia taxpayers' money, did we learn anything? Yes, I think we did. Are we go for launch?

Q: Yeah, yeah.

Chris Atkeson: Okay. One of the things we struggled with was the issue of privacy and as far as I'm concerned, people had very strange attitudes about privacy. If you walk into a room that has a lot of cameras bolted to the ceiling that you could see, a lot of people react very negatively. They don't like it. Interestingly enough in Georgia Tech, every single classroom had such camera because they had vandalism or security problems. We lived in a world that had everything on a camera. So we spent a lot of time trying to figure out how are we going to get people to be happy to have monitoring technology in their lives because if you want a servant to take care of you and the servant is deaf and blind, the servant is not very useful. Well, in the classroom example I gave you where we recorded everything, the students were totally happy to have a gazillion cameras because they saw an immediate benefit. So if you can sell people on an immediate benefit, it's a much easier thing and the TSA's problem is they can't convince people there's an immediate benefit to these strip searches. Oh well. A second issue is people's perception can be sort of shaped. For example, if I set up some cameras on the roof here looking at you, you'd have one reaction. If I rolled in a mobile robot that had eyeballs, the same cameras, you wouldn't react at all. Why that is? Well, I speculate that you're sitting there thinking, well, the robot needs to see, so of course it's going to have cameras and then you project all sorts of humanlike characteristics like it'll forget and it doesn't store stuff and stuff like that, but that is nonsense. The cameras in the ceiling and the cameras in the robot are exactly the same and have exactly the same system behind them, yet people's willingness to have them in their lives is totally different. So privacy was a big issue. A second issue was just how are you going to handle this mountain of data and we developed techniques to do that.

Q: What were some of those techniques to manage all your data?

Chris Atkeson: They mostly had to do with figuring out what the big problem was and then compressing. Worrying about the big problems and not the small problems and then compressing like crazy and there's a lot of analysis you can do in real time and then you compress the vision data, for example, or you throw it away.

Q: Were there other kinds of sensors besides cameras or do you have any kind of <inaudible>?

Chris Atkeson: Microphones were very helpful. Okay, so everybody in our day and age knows what GPS is and is very excited about it, but they don't often realize that GPS doesn't work indoors. Plus it's not very accurate. We would love to have a GPS system that worked indoors and everybody had a little tracking thing on them. That would make life so much easier. So we spent a lot of time basically buying systems that promised that and trying them out and then discovering that they had problems.

Q: Were there any controllers in the environment as well or just sensors?

Chris Atkeson: Okay. There were things like dim the lights. I don't know what level of control we got over HVAC, but our interest wasn't really in that aspect of smart houses or controlling those things. Our interest was in monitoring the people and understanding what they were doing and there actually was no robot in the house either.

Q: And so did you figure out something about human dynamics that was particularly interesting?

Chris Atkeson: So again on the privacy end of things, this is actually a project that happened here at CMU. We asked the question, well, how simple can the sensors be and still figure out who is who and what they're doing and most people are familiar with motion detectors and that doesn't bother them and they're happy measuring when doors open and close and when cabinets open and close and when the toilet flushes and stuff like that. So these are referred to as binary sensors and then you can ask the question. Well, with all these binary sensors, can you still figure out who's who and what they're doing and the answer is yes, if you use sophisticated machine learning techniques. So in reality, we probably don't need to put cameras in there.

Q: So you just learn people's patterns and behaviors?

Chris Atkeson: Correct, and that'll allow you to recognize them. So I had a student named Daniel Wilson and he instrumented his house with a bunch of sensors for his thesis work.

Q: Is this the Daniel Wilson of the goofy book? <laughs> What was it called? "How to Survive a Robot Uprising." Yeah.

Chris Atkeson: No, no. You are totally missing the bus. This is the Daniel Wilson who is my most successful and probably most famous student. In addition to the fine books that you mention, he has written four or five other books. His latest book to be released in 2011 will be made into a movie directed by none other than Steven Spielberg. So as a measure of success, he's far more successful than I am or any of my other students in terms of reaching people and getting a message out. Now I agree with you, it is unfortunate that the message mostly revolves around robots attacking people. That's not good press, but okay, that's his thing. He's doing it.

Q: Or the robots being actually quite easily incapacitated with small tricks.

Chris Atkeson: Well, okay. That is what he took away from being in the robotics institute and it's not inaccurate. It's not hard to incapacitate a current robot.

Q: So how did you get to CMU from Georgia Tech?

Chris Atkeson: I followed my wife. They called her up, said, "Can you help us get computer graphics going again?" and in the end we came and she did that and here I am.

Q: So when you got to CMU, you mentioned first you were at MIT where there was a lot of robotics. Georgia Tech, not so much, but you were still able to work on it, and then here you came to CMU and then there is obviously a great deal of robotics activity here.

Chris Atkeson: Right. The thing –

Q: How did that affect your –?

Chris Atkeson: Okay, the things CMU did that these other places didn't do, although Georgia Tech has now done it, is they made a commitment to robotics and said instead of distributing it on a gazillion different departments, computer science, mechanical engineering, electric engineering, aeronautics, operations research, wherever else, we're going to concentrate in one place, and they made that decision I guess 1979 or something. I don't know exactly, but that was

fairly early on. So clearly getting here, suddenly, wow. A lot of people are doing robotics and not only does that make you feel good, but there are a lot of people you can get help from or talk to or colleagues and things like that.

Q: Who are some of the ones that you interact with the most?

Chris Atkeson: So at my age, one of the very important activities is helping young people get going. So we have a system of older people mentoring younger people. I mentored somebody named Yoky Matsuoka who is now at the University of Washington. I mentored somebody named James Kuffner who is now at Google. I'm mentoring somebody called Hart Macgyver who we just hired recently. He's still here. And I spend a lot of time talking to them. In addition to the sort of humanoids work, I also because of the Aware Home work, got involved in something called Care Media which was a spinoff of something called Informedia and it's basically digital libraries. So they were interested in if you collect buckets of data, how do you make sense of it, and the Care Media was well if you collect buckets of video and audio and data and stuff like that from real environments, how do you make sense of it, and our emphasis there was going into nursing homes and helping take care of people with Alzheimer's Disease and stuff like that. I helped get together what I'll call a training grant for people in that area from the NSF and more recently there's been a bunch of us who have done something called it's a National Science Foundation engineering research center which is one of their big deal grants and something called quality of life technology which is again how can we use technology to help people.

Q: So you are actually now more on the application side in a sense as well.

Chris Atkeson: I guess. I mean I still also have an activity in humanoid robotics.

Q: And for the humanoid robotics, I mean did Japan end up being this big collaborator because of specifically you knowing Mitsuo Kuato or is there some reason that you're doing this specifically with Japan?

Chris Atkeson: You got to know who you're collaborating with. You got to have personal connections.

Q: Right, right, but I don't know. There are humanoids in a sense in the US as well, so what is it—?

Chris Atkeson: No there're not. There's nothing like it.

Q: But the Sarcos humanoid that they're using is obviously –

Chris Atkeson: We have one. The other two are in Japan. They're really – I mean okay, all you folks out there in TV land, I don't want to insult anyone, but there really is not that much going on in the United States in terms of humanoid robotics. It's nothing like Japan and Korea, I have a lot of interaction with Koreans coming here and Korean companies and they're working hard to catch up and they've got a lot more than what's going on in the United States.

Q: What are some of the centers of robotics research in Korea?

Chris Atkeson: It's very confusing. So they have something called KAIST, Korean Advanced something, and then they have something called KIST, both of which have efforts. They have a lot of big companies there, some of whom have efforts, and then there are other places as well.

Q: I'm still curious. So Japan obviously has a lot of humanoid activity in terms of designing them there's Waseda, there's all of these places, yet in ATR they're using the Sarcos humanoid. Why not –?

Chris Atkeson: That's probably the only Japanese lab that's using a US thing.

Q: Right. Why did they end up using the US thing? Is it mainly your connections or is it something –?

Chris Atkeson: I think it's partly due to my connections, but it's also due to – we wanted a humanlike robot and these heavily geared electronic things are just not. It's very difficult to make them humanlike.

Q: Okay, so I mean –

Chris Atkeson: They're useful, but they're not humanlike.

Q: And so in your estimation, DB is kind of the most convenient platform to do the humanlike?

Chris Atkeson: Okay, so DB's been retired.

Q: Okay.

Chris Atkeson: And I'll show you.

Q: The new one.

Chris Atkeson: The new one. Why do we have it? We can make the joints compliant. We have higher max velocities at the joints and we have bigger forces at the joints, all of which, by the way, are still not close to what humans can do, but they're closer than any other robot I know about.

Q: Okay. What do you consider the biggest accomplishments or contributions that you've made?

Chris Atkeson: On the algorithm side, the work in memory-based learning. I have, let's see. A lot of people my age are managers, okay? They don't do anything. They manage projects. Early on I figured out that that didn't make me happy, so I've always tried to maintain some activity where I actually program and unfortunately it's very hard to have time to work with an actual robot, so I program. So I'm actually very proud of my work in planning for dynamic systems which has sort of been my side activity to keep me happy all these years and I'd be happy to launch into a spiel about it. It's like having a drunk guy in a bar start talking to you. You can't shut the guy up. But only if you want to. So I'm proud of that and I'm also proud of the stuff that we've done with these various robot systems which the key difference between them and everything else is these are compliant systems and so in some ways that makes it more humanlike. In some ways it makes it harder, actually, because if something goes wrong, everything goes wrong.

Q: I'm planning for <inaudible> system. I'll take the bait.

Chris Atkeson: Okay. There's been tremendous progress in what's called robot motion planning and getting complicated robots to move slowly through complicated geometry and that's a tremendous accomplishment. What we don't know how to do is find the best path under those situations and also how to be more dynamic, move faster, forces matter, so I've been spending a lot of time trying to figure out how to do that and I think the best way to do it is to combine a lot of greedy planners together in such a way that when you're done you actually find a globally optimal answer. And what does that mean? It means a greedy planner means you go straight to the goal. So if I'm here and you yell fire, I run straight to the door there. But if I'm there and you yell fire, it means I run straight into the table. That's not a very good plan. But if we had a bunch of people all run to the door, some of them would make it. Then those guys, we could think about how they catch up with the people who made it. So in the end, by combining a

lot of fairly stupid planners, we can get a plan for everybody that's pretty good. So that sounds good, but the challenge is to make it real.

Q: That might be the final question. So our final question is if you were giving advice to somebody who's interested in getting into robotics these days, what would your advice be?

Chris Atkeson: So I think we have to separate out the robotics part of that question from the, "What's a good life strategy," part of that question. I think the good life strategy is you should do something that you fall in love with. You should just do it whether it's a good idea or not. It could be a terrible idea. Just do it anyway. That's my generic advice to young people. I mean there's something we didn't talk about, about robotics and that is suppose you walked in the door and said, "I want to make fundamental progress. What should I work on in robotics?" I think we have to recognize that robotics is a kind of screwy field in the following way. I drive a 20-year-old car. There hasn't been that much progress on the mechanical side of robotics. We have been building new mobile robots. We use something called the Barrett arm from Barrett Technologies. That was an arm that was developed when I was a graduate student, okay? That's the beginning of the 1980s. So if I have my math right, that was 35 years ago. We're installing a 35-year-old arm in on our robot and saying state-of-the-art. Contrast that to computers, which is sort of the other half of robotics. The idea that you would have a 20-year-old computer or a 35-year-old computer, it's a joke. We throw away computers after a few years because they're totally obsolete. So on one side, on maybe the sensing and intelligence side, you have this mind-blowing explosion of resources. The computer stuff is just blasting away. On the other side, we have slow or almost no progress on the mechanical side. So someone coming to me for career advice, I'd say well, you've got to recognize that. It's unlikely that you're going to make 100 percent improvement on the mechanical side because lots of smart people have been beating themselves to death with these same systems and it hasn't happened. On the other hand, you could be an idiot on the intelligence side, but because computers get faster, pretty much whatever you do is going to be better, whether it's a good idea or not. So I'd just point that out to them.

Q: Is there anything else you'd like to add?

Chris Atkeson: I don't think so.