RODNEY BROOKS INTERVIEW PART TWO

Hello again, Ars Technica readers. This is the second installment of a three-part interview with Rodney – who is one of the most influential people in the worlds of robotics and AI. If you haven't yet heard part one, there's a link to it on the page where you found this player, and I strongly suggest that you go back and listen to that installment before this one.

And with that - back to my conversation with Rodney Brooks.

Rob Reid: Tells us about Baxter and Sawyer, what you created.

Rodney Brooks: We started building robots, which used a technology that had been developed at the AI lab. Gill Pratt and Matt Williamson who had been a student of mine, they had developed something called a Series Elastic Actuator. The Series Elastic Actuator lets you know how much force you're applying on a joint and measure any external forces.

Rob Reid: On every joint?

Rodney Brooks: On every joint. There's a lot of mathematics around it to make it work, but we put those into our robots. Instead of going for rigidity, we had softness. There's a spring in there. The spring actually absorbs energy very quickly, 'cause springs can compute faster than anything. Then, the computation clicks in at 1000 hertz, and dissipates the energy in a few milliseconds. It's a combination of software and hardware, and exquisite control of forces.

Rodney Brooks: Our robots, if you go up to them and you get in the way of them, even if they're moving fast, they realize in milliseconds they've hit something they didn't expect, they just stop, so please don't do this out there in the real world, but I put my hand in front of our robots and have them hit, and they're safe. There's that aspect, but then, using that force perception, when you or I put a nut on a bolt, we don't locate that within a 10th of a millimeter and then go and put the nut in exactly that position, which a traditional robot would have to do. Instead, we go there and we tilt the nut slightly, and put it on the edge of the bolt, and then tilt it over, and then start turning and feel whether it catches.

Rob Reid: We do it by perception and proprioception.

Rodney Brooks: Yes, and we use these springs in the joints, Series Elastic Actuators with all this computation.

Rob Reid: To replicate that?

Rodney Brooks: To do that. Now, if you go into an electronics manufacturing plant in China, 30% of the labor is taking circuit boards, putting them into a tester where they're on alignment pins, making sure it settles down on the alignment pins, shut the

tester, run the test, and on a screen it goes red or green. Our robots can use exactly the same fixturing that people use, and our robots have cameras in them, and they look at the screen, and see whether it goes red or green.

Rob Reid: You said in one interview that I heard that basically for a lot of people in factories, what they're doing is they're waiting. They're tending to a very, very expensive piece of capital equipment. Because it's expensive, you don't want any downtime.

Rodney Brooks: So, they're not allowed to go to the bathroom for two hours.

Rob Reid: They're not allowed to go to the bathroom, they need to feed things in in a precise way, and then they need to wait for the light. They might be sitting there for 10 or 15 minutes, because that equipment is so expensive that you do not want to have the downtime. You can take over jobs like that. I know, because I've seen the demonstrations. It's incredibly easy to train the robots, because you're moving their arms in space.

Rodney Brooks: Yeah, although we're not teaching them a trajectory.

Rob Reid: You're teaching them a task, right?

Rodney Brooks: We're teaching them a task.

Rob Reid: Which is an important differentiation.

Rodney Brooks: Come to some place and now close your fingers until you feel such and such of force. Now, if the coordinates system varies, then either find the coordinates system by force, or we can say use vision here, and there's a screen which looks like the head of the robot. We actually have eyes on there by default when it's running, so it can look surprised when something is too weird.

Rob Reid: It also looks where it's about to reach.

Rodney Brooks: It's the same as the human worker. They don't do things that surprise you, because they look before they reach. We don't need to train the workers about this, they pick it up.

Rob Reid: A really interesting anecdote is when you went to Automate, which is a major robots trade show.

Rodney Brooks: Enormous trade show in the McCormick Center in Chicago. 2013, there were no uncaged robots. None at all. We said we want to have our robot uncaged and they were like oh, no you don't. Oh, no, no, no, no, no. That's not safe. Our VP of product management said "It is safe, it is safe," so he went with the safety people to our robot and then let the robot bash him in the chest time after time

after time. The safety people for the trade show, okay. Just this once. You've shown that it's safe.

Rob Reid: One uncaged robot [Crosstalk 00:35:30].

Rodney Brooks: One uncaged robot. That was 2013. I was there last year in 2017, there must have been 2000 uncaged robots on that floor.

Rob Reid: Thousands of uncaged in just four years?

Rodney Brooks: It's just changed.

Rob Reid: That's amazing. Then there's also this thing with PLCs in connecting to other automated equipment. If you're going to do that with a traditional robot, you need to start talking to this infrastructure, right, called a PLC. Can you give us a quick description of what a PLC is, 'cause it's an amazing legacy.

Rodney Brooks: Yeah, so, Programmable Logic Controllers. They were an invention of a company in Bedford, Massachusetts in 1968. They were amazing. Before that, control of industrial processes with electromagnetic relays. Now, a lot of people have probably never heard of them, but it's a coil wire wound around a metal core that pulls down the little lever, and that makes and breaks contact.

Rob Reid: And, you hear a click and everything else.

Rodney Brooks: You hear a click and that's what phone systems were built on. That was the technology through to the 60s. Phone systems were not then built on computers after that, they were built on special purpose electronics. PLCs, the PLCs and all their specs, there are two abstractions they use. One is the coil. There are no real coils, but that's one of the abstractions, and the other abstraction is a 16 bit number.

Rob Reid: So, they're still mimicking to this day, PLCs still mimic these mechanical relays from decades and decades ago.

Rodney Brooks: Yeah, and then it gets worse. You say, that's just the old companies. No. When I was writing a blog post in September, I went to TESLA's website, the Bay Area factory just to check, and there they had ads for PLC engineers, for the manufacturing system for TESLAs.

Rob Reid: So, the automated infrastructure is old?

Rodney Brooks: Yeah, it's old. Around that same time, I was with one of the major suppliers of this technology. I was talking about allocating some software updates. Every three months we try to have a new version of software. They said we try to do three versions every 20 years. Anyway, so, we realized that this was a friction point for people who hadn't had automation before they bring in a robot. Now

they're going to learn about PLCs? They're going to learn about this 60 year old technology? 50 year old technology? We made it so that for a lot of things, not every single case, are using PLCs. For a lot of the common cases where you want to connect a robot to other machines, use the same behavior trees on the robot to take over the role of what a traditional automation engineer would have done with PLCs. We got rid of them.

Rodney Brooks: That's how I like to think about it. I like to think about going to out customers and seeing what is the friction of deploying the robot. We've seen many pieces of friction. PLCs was one of them. You buy the robot, then how the hell do we connect it?

Rob Reid: When you think about the near to intermediate future of robotics, what do you see on the horizon that maybe most folks would not see?

Rodney Brooks: I am worried that we will not have enough assistance for eldercare. I want to be clear on this. I'm not talking about companion robots. That's not what I think is important. If you look at the demographic inversion, we see it happening in Japan. They've got so many more older people than young people already. They're going to 30% of their population over 65 within a few years and it's going to get worse.

Rob Reid: The number of working people who support a retired person has famously gone from double digits to very low, single digits in many questions.

Rodney Brooks: Yeah, and it's going to continue, and it's going to continue in the US. Why in the US? Well, we've relied on immigrant labor forever. We're not encouraging immigration anymore, which is bad for farms, and it's going to be really bad for eldercare. The baby boomers are about to get really old and there's a lot of them. I'm not saying that we want the robots to be companions, but what is important to people is their independence and their dignity. The first thing that is bad for old people is when they can no longer drive. I think not self-driving cars, 'cause I'm skeptical about how quickly they're coming, but the driver assist features that we now are putting in our modern cars.

Rob Reid: They're important.

Rodney Brooks: They're important, because they let the elderly drive safely longer. Their standard of living goes down incredibly when they have to go to managed care. They can stay in their house longer, they're mentally fine. But, what if they can no longer get into or out of bed by themselves? In my mother's case, that was a deciding factor when she went to managed care, because she couldn't do it by herself anymore. But, what if you had a device that a person could talk to? We've now got the speech systems with Amazon Echo and Google Home. Grab me here, do this, do that, and let the person maintain their independence of getting into and out of bed by a robot that can deal with the soft body of a

person, not hurt them, use understanding of forces to manipulate the person under their command.

Rob Reid: That's something that I have a personal understanding for, because I have a close family member, we have 24 hour care. That is a one to one mapping of a caregiver to a care receiver, partly because this person does struggle to get in and out of bed, and struggles to get in and out of the bathroom, and so forth. But, were it not for that, until very recently would have no trouble living independently, and we can't map people one for one, caregiver to care recipient. Also, pushing lots and lots of people to managed care, because two or three relatively simple tasks, it's a relatively big deal. How long do you think it will be before we have robots that can help with these very intimate, relatively straight-forward, but when you get down on a biomechanical level, very complicated functions like assisting in and out of bed. Is that within a decade, do you think?

Rodney Brooks: We're starting to see significant research in Japan. Now, when you go to a traditional robot trade show, which is full of [inaudible 00:41:06] robots, there will be a whole section of research institutes showing off their not quite ready for Primetime, yet, robots to do these sorts of tasks.

Rob Reid: Oh.

Rodney Brooks: So, they're working like hell on it. They're trying to create the market at the same time. We see a few places in the US. There hasn't been much research money going into it, but I'm having dinner with some entrepreneurs tomorrow as it happens who are doing something, I don't know what they're doing. They probably won't tell me exactly, but I know it's something to do with eldercare with robots, and they are people that have a track record of doing stuff. So, we stand to see it just at the edge happening.

Rodney Brooks: Now, whether that's going to be ready in 10 years for mass market, that's why I'm worried. It's going to be 20 years or maybe 25 years, but I think we'll see a flow in bench capital and other things into this in the 5-10 year timeframe.

Rob Reid: It will be a colossal need and a colossal market. Markets abhor vacuums, so, hopefully there will be radical and unexpected progress there. Now, you mentioned self-driving cars. That is the very intersection of AI and robotics, which are two things that you were almost foundationally early into. You wrote a fabulous blog post on the first of this year, in which you made dozens of predictions dated with years by which things in robotics, self-driving, AI, and so forth can happen. First of all, I just want to salute you for having the courage to do this. Most people, even self-proclaimed futurists, tend to make their predictions vague and leave them highly undated. You boldly put years on pretty much everything that you made.

Rob Reid: Now, a few of your dated predictions pertaining to self-driving cars, you said no earlier than 2022 you expect a driverless taxi service in a major US city, and this is a critical factor. Dedicated pickup and drop off points in restricted areas, so, highly restricted taxi service. You didn't even see that happening until 2022 or more likely thereafter.

Rodney Brooks: Yeah. I think we'll see people claiming to have demos, but they will not be real. It reminds me of a great tweet that I saw about a year ago about San Francisco actually. The passengers knew they were getting into a driverless Uber, because there were two people in the front seat instead of one, but I think that's the point. We say there are driverless cars and all these things, but they're not really. In Phoenix, there's going to be some things, but whether it's real or demo mode is a different question.

Rob Reid: You were very specific in your predictions. You were saying an actual, functioning taxi service with lots of restrictions, designated pick up and drop off points, restricted part of the city, 2022. I think you said no earlier than 2032 will we see a taxi service with arbitrary pick up and drop off in a major metropolitan area. Do people in the field consider these to be timid suggestions? Do you get pushback from people in the industry?

Rodney Brooks: I don't get pushback from heads of major automobile companies. I don't get pushback from heads of places that are really trying to do it for real. I get pushback from people in Silicon Valley who think that everything is just going to happen quickly. They see Waymo cars driving around Mountain View and think that's general.

Rob Reid: These would be investors, journalists?

Rodney Brooks: And, tech bros.

Rob Reid: And, tech bros. Got it. Why do you think these people and others are overestimating the rate of development in this field?

Rodney Brooks: I think they're making a bunch of mistakes. I asked them when did the first car drive down a freeway for 10 miles at 50 miles an hour. They know that the Google cars did that in 2004 or 2005. It was actually done in Germany in 1987.

Rob Reid: Wow.

Rodney Brooks: When are we going to get the first car, hands off the steering wheel, feet off the pedals, drive coast to coast in the US? Yeah, well, it actually happened in 1995 with the Navlab Project from Carnegie Mellon University. My point is, everyone thinks, oh this is just a [inaudible 00:44:50], this is going to happen quickly. It's been around a long time to get to where we are. I have now demonstrated to them that their scale is wrong, their start point is wrong. It has taken a lot longer to get to where we are now.

Rob Reid: That's a really important subtlety, because when you look at a Waymo car today and you think all this started four years ago, it seems so rapid. But, when you realize this started in 1987, it's a completely different scale.

Rodney Brooks: Anyway, so, there's that point. The second point is everyone's concentrating on the driving. That it can drive here, it can go on this freeway, it can do this. I think that's the easy problem. I've heard major manufacturers of cars' lawyers say, "Our car will never break the law." Well, in my neighborhood in Cambridge, Massachusetts, I cannot drive anywhere, any day without breaking the law.

Rob Reid: Give me an example.

Rodney Brooks: I have to cross a double yellow line everyday. I cannot drive around. Everything is so crowded.

Rob Reid: There's also a lot of one way streets where you park.

Rodney Brooks: There's a lot of one way streets and about every three weeks, I have to drive the wrong way on a one way street, because there is a temporary blockage there.

Rob Reid: It's the only way out?

Rodney Brooks: It's the only way out. I can sit and wait for six hours-

Rob Reid: Literally.

Rodney Brooks: ... For the road workers to clean up and leave.

Rob Reid: So, if you've got a self-driving car that will never break the law, suddenly you've got a very unhappy passenger who is sitting there for six hours?

Rodney Brooks: Well, self-driving cars are probably going to have speech interfaces. Is the passenger allowed to tell the car to break the law? What if the passenger is a 14 year old that has been sent to soccer practice by their mum? Is the 14 year allowed to tell the car override what the car wants to do?

Rob Reid: Somebody without a driver's license, yeah.

Rodney Brooks: Exactly. Or, what if it's a dementia patient? Are they allowed to? Dementia patients often have very strong opinions about things, which are wrong. I think there's all sorts of issues whenever you get picked up by an Uber in a very urban environment or Lyft. They often pull into a no parking zone and stop. They pull into a bus zone and stop. You and the driver know that's okay.

Rob Reid: Yeah, that's a necessity. You're dead right. Almost every Uber or Lyft I've gotten into, there is some tiny violation of the law that's going on. That doesn't really bother anybody, it's the way it works, but it's a violation of the law.

Rodney Brooks: Now, are the cars allowed to violate the law on their own, and where's the limit? Then, what happens if something goes wrong once? Then, they get sued for violating the law.

Rob Reid: I guess the car industry had to deal with this when they decided to allow cruise controls to be set for 75 miles per hour back whenever cruise control came out. But, that's a lot different.

Rodney Brooks: Still, the driver in the car, whose the one whose liable?

Rob Reid: Who set it. Precisely. Yeah.

Rodney Brooks: Is it the passenger in a Mobility as a Service car, are they the one that's liable now? I think there's all sorts of issues and you can say that's just details, but that's what takes a long time for things to happen.

Rob Reid: There's dozens of those details, there are hundreds of them.

Rodney Brooks: There are hundreds of those details. Let alone the fact that driving most places in the US is very different from driving in Mountain View or Phoenix. Driving on the east coast, driving in a highly congested city, very different style of driving. There's many, many more problems, but it's the social interaction. Then, there's the pedestrians. In a crowded, urban environment such as I always refer to Cambridge, Massachusetts-

Rob Reid: Or, all of Manhattan, or-

Rodney Brooks: Or, all of Manhattan. There is an interaction between people and drivers all the time. Social interaction. All the time at every intersection. Whose turn is it to go?

Rob Reid: You're catching each other's eyes, you're making sure people are paying attention.

Rodney Brooks: Even in a highly lit area, you are interacting with people in a car. What if that's a Mobility as a Service Car, which doesn't have a passenger in it to get angry, and you know that for liability reasons, it ain't going to hit you. Now, you don't have to be polite to it anymore. Why be polite to that car with no person in it? Just step out in front of it. Or, if you're a kid, hey, let's go chase some cars. It will be a fun game.

Rob Reid: Car chasing would actually be kind of fun if you were a kid, maybe even as an adult now that I think of it.

Rodney Brooks: These are the sorts of things that aren't in the algorithms right now. The car is pulling up to an intersection. Is that person chatting with their buddy on the corner or are they about to want to walk across the road? Then, if there's a

painted white line there, is it legally the responsibility of the car to stop and let the person go? You have to make a judgment on what their intent is in the next second or so.

Rob Reid: Or, if somebody's waving you down. Is it a police officer? Is it somebody of some authority in an orange vest? Is it some kooky person? Is it a carjacking? How do you make those distinctions? In San Francisco, there's often temporary handwritten no parking signs. All these little edge cases mount, and mount, and mount, and then you start thinking that this is every single driving situation that we're in.

Rodney Brooks: I said all this sort of stuff in my blog and someone said, "Oh, you just don't understand. There aren't going to be one way streets anymore, 'cause all of the cars are going to communicate with each other, and they're all going to decide the optimal routing, so we won't have one way streets, so problem solved." That may be in the 2070, but that's a long time between now and then, so there's going to be a mixture of human drivers and self-driving cars for 20, 30, 40 years.

Rob Reid: It's got to be a slow, steady thing. The day the first self-driving car ships, there will not be any special lanes or regions for it.

Rodney Brooks: Silicon Valley, because of Facebook, Google, Twitter, are used to doing code pushes a few times a day, deploying a new version of the code, because the cost of deployment is almost zero. Because, every time you fire up your browser, you go somewhere, the code downloads. That's very different when it's a piece of capital equipment.

Rob Reid: I've had two cars as an adult. I had a 1991 Nissan Centra and then I bought another car in 2003 that's still serving me very well. There will be a lot of non-self driving cars when the first ones get out there and there's also these unbelievable opportunities for the 1% to misbehave. You had this great blog post that I think was a little bit playful, but you had this phrase. "Virtually parked cars." Describe what a virtually parked car is.

Rodney Brooks: This came from Cambridge again. I [inaudible 00:50:31] jump into a UPS store, and there were no parking spots, and I thought man, if only my car was self-driving. I could jump out, it could cruise up and down Mass Avenue-

Rob Reid: And, be virtually parked, yes.

Rodney Brooks: Virtually parked by driving and then be ready for me. As soon as you get that capability, people are going to abuse it up [the hill 00:50:48].

Rob Reid: Another argument against handing the steering wheel to an AI today is really top of mind for me, 'cause as I was just driving over here, my brakes failed, but my steering continued to work, and I had to choose between running over a nun

and two thieves that were on the other side of the road. As I was confronting this-

Rodney Brooks: Is one nun worth-

Rob Reid: Worth two thieves, yeah, yeah. 'Cause, it's two lives, but this person is probably a better person than the two thieves. As I was confronting this difficult choice, it occurred to me that it would be very hard for an AI to think through this incredibly common situation that most drivers confront on an almost daily basis. Okay. For those who are listening and don't realize I'm being playful, that certainly did not happen, but this is an [inaudible 00:51:30] of something called the Trolley Problem that a lot of people have given thought to. How would an AI choose between running over one nun or two thieves, and do we want to hand that problem over to a non-human decider. What's your take on the Trolley Problem?

Rodney Brooks: My take on it is it's a great intellectual thought problem, but in practice, it doesn't happen to you.

Rob Reid: No, I've never faced this [inaudible 00:51:55].

Rodney Brooks: You've never faced it. What happens when there's an accident coming is you're trying to put the brakes on as fast as you can, and that's about all you can do. The only piece of practical advice I've ever been given about this is in driving in northern Maine. If a moose is crossing the road, swerve towards its rear-end, 'cause it's going to keep going the direction it faces. That's about as much processing as you're going to do.

Rob Reid: But, it is funny, 'cause when I've sat around with brilliant people in tech, brainstorming about this stuff over dinner, I've probably had 5 to 10, maybe 20 times as many discussions of the Trolley Problem as I have had of the problem that most pickups and ride shares involve double parking. Going from robotics, where the convergence of AI robotics to AI, you wrote a great post back in September called The Seven Deadly Sins of Predicting the Future of AI about cognitive distortions that people commonly fall into when they're thinking about the future of AI, and also the risk that a hypothetical Super AI might pose to us humans. I'd like to start out of with one that is called Amara's Law, which is that we tend to overestimate the effect of the technology in the short run, and underestimate the effect in the long run. Why is that significant and how does that play out when people think about AI?

Rodney Brooks: We saw that with computers. In the 60s, people were worried that intelligent computers would decide that people were useless and launch nuclear missiles to kill all the people. We didn't have any computers in the 60s that were remotely capable of that. We still don't. You know, I'm a Trekkie, the original Star Trek, or even The Next Generation in the 90s, you look at what the

computers were predicted to be doing in 300 or 400 years, that looks laughable now.

Rob Reid: [inaudible 00:53:44] your interstellar travel, but the computer couldn't do things that any computer could do today?

Rodney Brooks: Exactly.

Rob Reid: In that case, certainly underestimating what the longterm was.

Rodney Brooks: Yes. We've seen that.

Rob Reid: I would say one thing for me that's seared into my memory is mobile computing. I first moved into technology in 1994, I went to Silicon Graphics, as I mentioned. Apple had just released the Newton, all of the really smart tech-forward people in my company, which was almost all of them, were carrying around something called a Sharp Wizard, which was a personal digital assistant. At that point, although it was a little bit prior to Primetime, I would certainly have guessed that within three or four years, there'd be portable devices that were at least as useful as Outlook. I would have been off by a factor of several hundred percent, but there's no way I would have imagined that there would be well within my lifetime, this super computer in my pocket that can summon cars, and do all the crazy insane stuff that we do with our smartphones. So, people I guess you would argue, are systematically overestimating how much AI will be able to do in the near term, near term I imagine being what, decades or something like that?

Rodney Brooks: I think people are overestimating it in two years, in three years. Oh, it's going to do everything, it's going to change everything. The counter to this is they are underestimating longterm, but I think their timeframe for the longterm is also too short. I think it's going to be hundreds of years.

Rob Reid: Now, there's an interesting distortion here that hearkens back to what we were talking about with self-driving. I think a lot of people are looking at AI and they're seeing these radical breakthroughs that we've had in the last few years with machine learning. There's a sense that we've gone from zero to 90 miles per hour in a very small, single digit number of years. Just five years ago, identifying what's in a picture was something that would defeat any system. Seeing that in the real time that it's happened, the very low, single digit number of years, one could certainly think "Oh my God, this is going so rapidly," but you've been working in AI since the late 1970's, and you know that its history goes back to the mid-1950s, so it's kind of like the starting point problem.

Rodney Brooks: Yeah, let's talk about deep learning.

END INTERVIEW ELEMENT OF PART TWO

Bit of a cliff hanger there, huh? Sorry about that - but we've come to the end of the second excerpt of three from my interview with Rodney. As you could probably tell from that final sentence - we're gonna talk about deep learning tomorrow.

As mentioned before, if you can't wait to hear the rest of the interview, you can just head on over to my site, at after-on.com. Or, type the words After On into your favorite podcast player, and scroll through the episodes to find this one, which originally ran on March 19th. Or, you can join me tomorrow here on Ars, when we'll continue with third and final part of this interview.