

Chicago, Illinois
1:08:37
.mp4

Ron Arkin

An interview conducted by
Peter Asaro

16 September 2014

Peter Asaro: If you could just start by telling us where you were born, where you grew up and went to school?

Ron Arkin: I was born in New York City, New York – Manhattan. I was raised on Long Island – Roslyn, on the North Shore. And I did my undergraduate degree at the University of Michigan; my master’s degree at Stevens Institute of Technology, in Hoboken, in New Jersey; and my Ph.D. at the University of Massachusetts in Amherst.

Peter Asaro: What did you study as an undergraduate?

Ron Arkin: Two things. I have two degrees in chemistry, which many people don’t know. The first was at Michigan. It was a BS in chemistry. But, also, I got a dual major in applied mathematics, which shows my age, because they didn’t have computer science there, at the time. So I had a good advisor who told me that, “You’re taking all these mathematics courses, so maybe you should take a few more.” And I did. I spent an extra summer term there, and they gave me the dual major, which gave me the insights into the computer science career that I have. But I also ended up eventually working in the chemistry industry for a significant number of years, first at a dyestuff company, and then as a pharmaceutical R&D chemist, both in R&D, with my bachelor’s, and then I got my master’s while I was working full time. So I like to call myself a “serendipitous roboticist.” I kind of stumbled into it for – in an odd way; not like most folks, who have had this passion for robotics from their youths. I’ve had a passion for science from my youth, but not for – specifically for robotics.

Peter Asaro: So your degree from Stevens Point was...?

Ron Arkin: Stevens Institute of Technology. It was in organic chemistry.

Peter Asaro: And then what made you decide to go back to do a Ph.D.?

Ron Arkin: That’s an interesting convolution of events; or a “configuration of events” is a better way to put it. What happened was I was working in this R&D laboratory, working with chemicals that were labeled “known or suspected carcinogen” on a daily basis – fumodes [ph?], of course, and the like, as well, too. And I loved the joy of discovery of chemistry, which was one of the beautiful things about chemistry, a heavily plowed field. But I also recognized, “Maybe there’s not a future for me if I keep doing this.” And there were accidents occurring in the plant that I was working at, because we took it from research bench, all the way out into process in the plant. And I kind of said, drawing on, “Yes, I was a hippie.” I actually had long hair at one point in time. And so we decided to move up to New Hampshire, my wife and I, who’ve been married for 41 years now, to investigate what kind of work we had. I had bought

50 acres of land up there prior, on a previous visit, so moved up and took a job at a small college which is now defunct. It was called Hawthorne College, in Antrim, New Hampshire; taught there for, I think, close to 10 years, with only my master's degree in chemistry. A series of events occurred where the Florida Institute of Technology came up and bought the college as a campus for them, because they had a flight program, and it was boring to fly in Florida, and we had interesting terrain up there. And since it was a Institute of Technology, they were interested in starting up computer science, and as I had some background from applied mathematics, I volunteered to set up the program. I served as department chair, and after a while I started thinking maybe I should get my Ph.D., if I'm going to be a department chair. And so the school was very accommodating, let me go take the time off – not completely. I was still teaching, and the like, at a – it was a teaching school – and went and got my degree. First, I went to the University of Lowell, and I took a robotics course there because it fit my schedule, not because of any deep-seated commitment for robotics. Then I found – and I was disturbed because I had to take the GRE, even there. And I said, "I'm a chair of a computer science department. Why do I have to take a GRE?" Well, I took it, and they give you a couple of places to apply to, and I – UMass Lowell was actually – UMass – excuse me – Amherst was further away, but it was faster because it was freeway. I didn't know that at the time, so I gave my score. I was accepted because I applied there, and when I was interviewed there, they assumed I was going to be a roboticist because they saw I took a robotics course at the University of Lowell, and they lined me up with all these folks at UMass – Wendy Leonard and others, as well, too, in the AI group – and it stuck. <laughs> So... <laughs> it wasn't a conscious choice, and I would say I'm fairly unique in the field for two reasons. One is the serendipitous aspect of being led. I kind of feel like Jonah being vomited up on the shores of robotics, from a calling from God, if you will, perhaps. So that was one of the reasons, and the other is my scientific background. I have formal training – extensive formal training – in the scientific method, and I think that's something that, in the computer science area in which I am housed within, not as many of our colleagues have that kind of basis for a deep scientific exploration and understanding that they – as they should. And so, again, those two things, I think, distinguish me from some of my colleagues.

Peter Asaro: And so when did you do your Ph.D. work?

Ron Arkin: I graduated in '87. I was able to complete it in three years, which was unusual, but I had four kids living at home at the time, so I had unusual pressure to do so. I spent – I would go down after the courses were done, and I was able to stop teaching for a while. I would go down after all the qualifiers were done, in a year, year and a half. I would go down on Monday, spend Monday night on campus, come back late on Tuesday, work from home on Wednesday, go down Thursday, stay overnight, come back Friday, and then rinse, wash, repeat, or whatever the order is, for about a year afterwards, as well, too. So it was a bit grueling, but it was very satisfying, as well, too.

Peter Asaro: And who was your thesis advisor, and what was your thesis on?

Ron Arkin: My thesis advisor was the late Ed Riseman, at the University of Massachusetts. He was extremely well known in the computer vision area, and I worked in what's called the Visions Laboratory, which dealt with interpretation of natural scenes, dealing with very tough problems, and that's where much of the schema theory that has permeated my work has come from. I also like to say that Michael Arbib was my mentor, in this case. He was there when I first got there, but he left for USC shortly afterwards. But his work has had a profound influence on the way in which I do robotics.

Peter Asaro: Can you explain that a little more? Was it some of the cybernetics that he taught?

Ron Arkin: Yeah, it was the metaphorical brain and other things: the relationship of biology to intelligent machines, and his formal language that he's developed and used in many ways, from everything from explaining how the brain works to how manipulation works, called schema theory, which permeated UMass at the time, as well, too, in many different groups in – the AI group was the lingua franca, if you will. It made a big difference in the way I thought about these particular types of problems.

Peter Asaro: And how were you applying the schema theory to robotics at that time?

Ron Arkin: Well, at the time, the behavior-based paradigm was emerging. If you look at the early work by Brooks, and others, as well, too, it tended to be rule-based. Schemas themselves were concurrent asynchronous processes, which could be multithreaded and run independently. There are also things – those are the motor schemas. Perceptual schemas were the same, so it became an interesting question of how you schedule and manage these particular processes, kind of like, one could argue, the brain does, as well, too, with all these different processes, and how do you make sense of it. My thesis title, if I remember correctly, was *Towards Cosmopolitan Robots: Intelligent Navigation in Manmade Environments – Extended Manmade Environments*. So we're looking at inside a building, and outdoor structured sidewalks and the like, as well, too, but not the wilderness or something like that. The key factor, I think, which distinguished my work was – and the advantage of using schema theory, was the ability to create what I call hybrid architectures – hybrid deliberative reactive architectures. The early work in behavior-based robotics, I would say, threw the baby out with the bathwater. All the high-level reasoning and high-level planning no longer played an effective role. My work recognized that maybe we do have a cortex, and there is some value in being able to deliberate about certain things, and the schemas themselves gave us the ability to instantiate change parameters, reason over these things in ways that other forms of reactive systems, if you will, were not as easily amenable to. So the real focus of my dissertation work was hybrid deliberative architectures. The schema theory got

most of the press because that's what made the robots actually do the cool stuff that they were doing, but the deliberative side was what configured those particular control systems.

Peter Asaro: Can you give an example of how would that play out, in terms of a robot making deliberative decisions?

Ron Arkin: Sure. We used – initially used A-Star [ph? - usually "A*"] path planning over maps of the world, both that were *a priori* available – available ahead of time – as well as those that were sensor-acquired. So we had both short-term and long-term memory. Those would create legs for a robot to complete. If I wanted to go from here, for example, to a plenary session, and I had a map of this, I could generate automatically, using a planner, a set of waypoints. But for each of those individual waypoints, instead of following those in a rigorous, "I'm following this trajectory," we would use behaviors, the reactive component, which would be selected by the deliberative system and parameterized by the deliberative system, based upon the context, to be able to move from one location to the other: goal-directed while avoiding obstacles, including a little randomness, which we called "reactive grease" at the time, and other aspects which would be able to complete and accomplish the mission in ways that are taking advantage of *a priori* knowledge when it is available, but still being able to function in the absence of it, which was, I think, an important contribution.

Peter Asaro: And did you implement it on a physical robot?

Ron Arkin: Yes, we did. We had a – one of the very first – I think, the Denning robot number two [ph?]. We called it HARV, was the acronym which we came out. Initially, we called it HARV because it was short for Harvey Wallbanger, which was an interesting cocktail, as you may know; and the other reason, we had – dutiful computer scientists must come up with an acronym for H-A-R-V, so we called it HARV because it was "Hardly Autonomous Robot Vehicle," at that point. So we'd come a good ways, and the whole field has come a good ways since then, but it was hard. It was very hard in those days to be able to, given the limited computation and perceptual power, to be able to make these machines move in real time, in a safe and effective manner. And I think some of the architectural capabilities and the decomposition – the parallel distributed processing, which was present in that – helped the field move forward.

Peter Asaro: And did you stay at Hawthorne after you finished your Ph.D.?

Ron Arkin: No, I actually left Hawthorne after I completed my qualifiers. So I knew I couldn't be working full time and getting my Ph.D.. That just is a impossible situation. So I quit about a year and a half before then, and devoted myself fully for my Ph.D. studies at that point in time.

And then I started looking, as people do when they graduate, looking for an academic position. Interviewed at a number of places. Georgia Tech was the one that finally landed me, made me the right offer. My family was very happy to move out of the cold winters of New Hampshire, where we were living at the time, and I've been there now 28 years. It was a good decision. One of the things that has kept me there all this time is the strong influence on interdisciplinary research, which, as any real roboticist knows, that is absolutely crucial. We can't know everything in every field, and we need to find the right people to be able to team with and draw on when we need additional knowledge. So I went down there. I was the first roboticist in the college in the computer science group. There was one other roboticist that I can recall there at the time, which was Wayne Book, in mechanical engineering. He's now since retired. He was working on flexible arms and the like. So a lot has transpired. Now we have almost 60 or so different roboticists, a Institute Research Center in robotics, a Ph.D. program in robotics. I won't take the credit for that, but at least I got to witness it happen, as well.

Peter Asaro: Can you tell a little bit about the history of how those things emerged? Did they come out of the CS Department, or more [redacted] projects between the different [redacted]?

Ron Arkin: The center itself and the Ph.D. program came out of the Computer Science program, whether it was the College of Computing or Information and Computer Science, and I have to give credit to the hires that we made – the young faculty, or the Young Turks, as I would describe them – folks like Tucker Balch and Frank Dellaert, who said, “We need to do this. We need to do that.” So, both having spent time at Carnegie Mellon, and wanting to create a similar structure to be able to do that. We also had, at the same time I arrived, a president who was the former provost, Pat Crecine, at Carnegie Mellon University, who pulled out the Computer Science Department from the College of Science and Liberal Studies, and made it a stand-alone college, which was amazing, I mean, because that was the model that they had at Carnegie Mellon. And it gave us a lot of freedom and opportunity to try things; basically, enough rope to hang ourselves with. And so now that we didn't have to answer to the Dean of Science, and we had our own deans and the like, we could explore things in new and different ways. The younger guys got us old guys worked up a little bit, so we moved forward with it. And then we – after time, the first thing was a – I believe the Ph.D. program came first. Then I had the opportunity to do what technically I should call a professional leave with pay, but effectively it was a sabbatical in the Royal Institute of Technology in Stockholm. I spent a year with Henrik Christensen, working with him, watching him, in his second year there, setting up the Center for Autonomous Systems that they had there, helping him with that. And through a series of events later, I was even serving on the board of that center, but he called me up once, one evening, and said, “I'm leaving.” And he was very upset. He was not happy with the way the Swedish government was treating them at that point in time, and he said, “I think I'm going to Germany.” And I tried, actually, at length to talk him into staying, because I was an Advisory Board meeting, but he convinced me that he was going to go, after – I don't know – 45 minutes, an hour, or something like that, on the phone with him, or at least it seemed like that to me. And he – I said, “Well, maybe you should think about applying here.” And I did. He was well received, and now has

become the center director for what was first the Robotics and Intelligent Machines Center – RIM. We were going to call it the Center for Robotics and Intelligent Machines, but our provost, Jean-Luc Chameau, who is French, noted that that “CRIM” is “criminal” in French. <laughs> So that would probably not have been the best one, although we thought of wearing striped T-shirts, or something like that, to show our commitment in many ways. But we – after we got that center going, after a number of years, we kept growing the mass, and eventually, under our new Executive Vice President for Research, Steve Cross, we were able to stand up what’s called at Georgia Tech an Institute – an Interdisciplinary Research Institute. There are only, I believe, eight to ten of those on campus, where it’s basically the fields where the entire Georgia Tech –

<break in recording>

Ron Arkin: – going to put priority in research, for hiring and other aspects: space, things like that, as well, too. And so Henrik deserves the credit for doing that, leading the National Robot – the National Roadmap for Robotics, among other things, as well, too, playing very, very well on the national stage, things I know he has talents for, which I saw firsthand before, and they are not my particular strengths to be able to do that, as well. So I think the aggregation of people was right, the timing was right, and we’re very pleased with where we are right now.

Peter Asaro: So we can go back to when you first arrived at Georgia Tech. How did you set up your lab?

Ron Arkin: Well, it’s interesting. My very first lab was a – it was a point problem for mobile robots. The Dennings were big in those days. They were about this big, and about this high. So I got a Denning robot, and my lab consisted of a small faculty office, at that point. So it was really just a point problem. <laughs> The robot would go like this, and – so then we moved into the corridors, and eventually I recognized that, “That’s not enough.” But there was no space in the college at the – or the school. It was actually a school at that point in time: the School of Information and Computer Science at Georgia Tech – ICS. So I started looking around for Interdisciplinary Research Institutes. There was one – Oh, gosh, I should remember the name of it. Manu – no, it wasn’t the Manufacturing Research Center. It was something like that – a precursor for that. And so I moved into a different building, surrounded by engineers. They were happy to share their space with me, and so I actually had more laboratory space. Then, after a while, they put up a new building on campus called the Manufacturing Research Center. I got a really nice lab at that particular location, with adequate space for students and offices and for robots. But then later, the best lab I got – I still have it right now – was in the Technology Square Research Building, another new building we built, across the freeway, which was a new area. It’s in the basement, and I think it’s one of the most spacious laboratories in all the research – mobile robot lab research in the country. So we can reconfigure the space in interesting ways, we can do human subject studies down there, and the number of robots has

proliferated dramatically in many, many different ways. Initially, I did a lot of work, after the architecture stuff, in teams of robots and multi-robot scenarios, and so we were able, through NSF funding and DARPA funding, to buy significant numbers – four, five – of different types of robots to show we could maintain formations and do coordinated activity, and create software where average people could be able to program these kinds of systems.

Peter Asaro: And did that sort of build intellectually on the schema work and the behavior-based work?

Ron Arkin: Yes. My – there's an actual – a very natural progression to the kinds of things that I have done in my career, initially starting with single robots in deliberative reactive architectures. Then we started looking at teams of robots, first working mostly in the reactive side, and people were asking, "It's so hard to make a single robot do anything. Why are you working with teams of robots?" Well, the real answer to that question was, "It changes the way in which you approach the solutions to problems." It's like distributed computing; you can do things in a fundamentally different way with distributed sensing and distributed actuation. So it made good sense, and **Amiah Matarich [ph?]** was a colleague of mine, as well, too, at that point in time – not at Georgia Tech, but working in this particular space. And then after that, years later, we started looking at human robot interaction, recognizing that teams not only mean teams of robots, but teams of robots and people. So we were concerned about how we could integrate human beings into these kinds of teams, and leverage their deliberative capabilities with the reactive capabilities of robotic systems. And then, after that, came the quality of the relationship, which spoke to the ethics work that I later did. In other words, well, if humans and robots are working together as teams, how should they effectively work with each other as teams, and what should be allowed and what shouldn't be allowed, and under what sets of circumstances? So that's been the progression. I'm still doing things now in formal methods, trying to establish guarantees for the particular robotic missions that we develop for humans before they deploy them, in particularly dangerous situations, such as counter weapons of mass destruction. We've spent a lot of time looking at heterogeneous robot teams, where you may have small robots and big robots and flying robots, and how we can deal with the issues afforded by heterogeneity, in terms of doing things better and differently, and how they perceive the world through different sensors, and how you can reconcile those sorts of things, and getting back to a theme that permeated all of my research, which goes back to Arbib and others, as well, too, which is the bio-inspired work. We're always looking to nature for the ways in which intelligent machines can function, because one, they are proof of existence of intelligence, and I use "intelligence" in the broadest sense, not just limited to human intelligence, but equally to ants and their domain, and wolves and frogs, and all sorts of species have provided us the ability to analyze and apply these techniques; and two, to get a better understanding how we can back-project some of this into the biology community in certain cases, as well, too. Generally, I don't do that, but I am not reluctant to let biologists say certain things about the relevance of our work, whether it's from a cognitive perspective, or whether it's from an ethological, or – which is an observation of animals in their natural environment – point of view, how that can influence their studies, as well, too.

And it's always been extremely fruitful for me to work with folks in other domains, and now that has been more recently extended to lawyers and philosophers, and others, as well, too. That's the thing I find most exciting about the field, is the ability to grow in new dimensions. And getting back to my chemistry roots, chemistry has been a relatively well-plowed field for hundreds and hundreds of years, going all the way back to alchemy. Robotics is... I won't say virgin territory, but it is a wide-open frontier. And I've often said that in the early days, including myself, many of us were more like cowboys and cowgirls, just showing proof of concept of anything, and if you make a robot get from here to the door, "Wow! That's an amazing thing!" The field has progressed tremendously since then, in terms of the requirements for experimental methods and other things, as well, too, which I think is very, very good.

Peter Asaro: Just to touch on some of your points about the ethology and looking at animal behavior, you did some work with Sony on the AIBO with dogs?

Ron Arkin: Yes, I did. I worked with Sony for 10 years, initially on AIBO – we have a patent on that work – also with QRIO, which is the small humanoid that they had. I developed a model for them. I think it was – I studied human behavior and I studied dog behavior, and that's the important thing that's quite often neglected, is understanding the environment in which these systems reside. You have to understand, what do pet owners really want, and how do they interact with these systems? And what are the underpinnings of a dog's – a trained dog's – behavior, or domesticated dog's behavior? These 183 behaviors were compiled into a model, using methods from Timberlake and Lucas, cognitive psychologists, to be able to deal with some of the coordination mechanisms. The Scott and Fuller eth – who were ethologists, spent their lives studying dogs. And others, like Fox and others, provided the basis for the enumeration of the kinds of behaviors. We developed this fairly complex architecture to be able to allow these systems to act more like dogs than might be thought otherwise, and we included things like humping your leg and chewing the carpet – the maladaptive behaviors, as well as the normal ones. Sony did not implement those, <laughs> but – and I can't say I blame them – but I tried to be as complete as possible in this. We did not publish a lot. There are a few papers out there – two or three – and there's a patent I mentioned, as well, too. Part of this was the IP issues, and I worked with them as a private consultant, which made that feasible under those circumstances.

Peter Asaro: What was the patent actually for?

Ron Arkin: The patent was for ethological and emotional architecture, if I remember correctly. I mean, patent language is incredibly obscure, but it dealt with both the organization of the behaviors that were present within AIBO, and the way in which what they referred to as the "instinct emotional model," which was more ethologically correct referred to as "motivation," but I lost that battle. As to how these things could coordinate and manifest their behaviors in interesting and meaningful ways that hopefully would align with the owners' expectations, avoid

such things as what's called "behavioral dithering," where you get behaviors going back and forth. If you're hungry and you're thirsty at the same time, so you take a bite of food and then a drink, and then a bite of food and then a drink, as your internal thresholds go down, or should you be able to have some hysteresis in this? It stays depressed, and – but part of the trouble with obesity is sometimes that depression stays down, and people don't understand that, "I'm still hungry, I'm still hungry, I'm still hungry," and it should've quit earlier. So we have to find ways in which to set those particular parameters and the like, so that's what that was involved with.

Peter Asaro: Interesting. Did you do other consultancies with companies?

Ron Arkin: I did a significant number of other consultancies. I can't say I remem – that was the one that was the most durable. Again, they bought virtually all of my time that I had over the 10 years. Again, as a faculty member, you're allowed the equivalent of one day per week to consult, to do that. I didn't – I wasn't able to spend all that time on the project, but up to that particular point in time. And I've worked with Lockheed Martin and other companies, as well, too. <inaudible> I wasn't – I can't remember. You'd have to check my vita to get the specific ones. Georgia Power, I remember, there was some work; recently completed some other work for a small startup – incipient startup; a variety of – Evolution Robotics, as well, which is now part of iRobot. So a variety of different samples, I guess, but none was to the extent and the duration of the Sony work. Now, I've worked with Samsung, as well, too, but that was through Georgia Tech, partly through the Samsung Technology Advanced Research Center that we have there. I have two patents on emotional affective models of behavior for robotic systems, with my former student, Lilia Moshkina, and some others that were involved in the project, dealing with traits, attitudes, moods, and emotions, and how we can deal with time-varying affective phenomena, such as in coded-in circadian rhythms and moods. There were specific responses based on your own predisposition of your personality, which would be genetic in your particularly case, and how these things modulate themselves, both in the presence of time, specific things that you might have attitudes towards: particular individuals versus particular objects. You might like some food, and you'll be happy to see it, or you might dislike some food, and if it's put in front of your plate, you'll be unhappy to see it. These attitudes also are present, as well, too, and we developed what I think still, to date, is the most complete model for affective phenomena, which integrates all these aspects. We couldn't find a single cognitive model that has spanned this particular space, so we had to synthesize these, and we're still using that work now in some work in the management of early-stage Parkinson's disease and care, in a National Science Foundation grant that we're using.

Peter Asaro: Were you ever part of a startup, or try to do your own startup company?

Ron Arkin: No. There's been points where I've almost gone into those. I know I'm not a good businessman, and I think the key to successful startups is finding the right business guy to work

with. I have received great joy out of curiosity-driven research. I have been well compensated over the years, in any case. I admire those who have the wherewithal to be able to do that, but to me, it might have been more a distraction, as opposed to a goal, in my case, and I have enough money, I guess, is one way to put it at this point in time. I'm comfortable with what I have. And I think we are – there's many different ways to achieve impact, and one of those is through the university and private consulting, and other things, as well.

Peter Asaro: And who are some of the grant people who've supported your work?

Ron Arkin: Oh, gosh, all sorts. The National Science Foundation was the first. I still remember the joy of getting my first grant. That was just a – driving home, I was just all smiles <laughs> on that. Took a while, but, as you learn, you have to deal with rejection as well as success in this field, and if you can't do that, you shouldn't be in it. DARPA – I've had multiple DARPA grants; the Army Applied Aviation Directorate; the ethics work from the Army Research Office; as I mentioned, Samsung. I've had work with the Department of Energy. I'm going to give short shrift on many of these. ONR has been a major sponsor of me over the years, and also NAVAIR, out of Patuxent Air Force Base. Army Research Laboratories, as well, too; the MAST CTA program has been part of it. And I know I'm neglecting some at this point, but a lot of DOD sponsors, because most U.S. roboticists find that an easier pool, and they also do provide pure, basic research, as good as Six One [ph?]. Six One money is as good as NSF money, in many respects. I do not have a security clearance – although some people tried to tempt me to go after that, as well, too – because I like the freedom to be able to talk to people about anything, and I think that's paid off handsomely in my ability to have these discussions with folks such as yourself, internationally. And... so that's basically the sorts of things. I can't say that's complete, but it is a representative example, and if you check my Web page, you'll find a more complete version.

Peter Asaro: So in many ways, the military-funded research was more freely accessible than Sony or private companies?

Ron Arkin: Oh, far so, far so. I couldn't talk – I signed nondisclosure agreements with both Samsung and others, as well, too, so – and, in general, the DOD [ph? - also "DoD"] work is, for the kind of stuff that I do, is wide open. There have been some cases that – we've worked with, I guess, it was Lockheed Martin, or other companies, as well, too. When you're subcontracting, there may be some provisions which make you a little more careful, in terms of what you can say, due to nondisclosure agreements, but by and large – I started working in the manufacturing domain when I first got to Georgia Tech. That's where the money was, and I've often said, "I'm happy to play in anybody's backyard, if you can provide the resources to do these sorts of things." And my goal has never been to build specific applications, per se, but rather to explore the fundamental questions underlying machine intelligence and robotics, and how we can use

those in ways that have not been fully understood previously. Now, whether that's in dealing with internal organization of action and perception, whether that's dealing with coordination of teams of robots working together, whether it's dealing with the importance and the role of communication and these sorts of agents, whether it's dealing with human-machine interaction, as well, too, I like to think that we explore solutions in the more general way. And, as I said, if you're paying for it, we'll consider the application that you're considering, as well, too, but deal with the fundamental advancement of knowledge, which I think is what academics should do.

Peter Asaro: You mentioned your work in Stockholm. Have you had other major collaborations outside of Georgia Tech?

Ron Arkin: Yeah. Again, the other sabbatical I took – “professional leave with pay,” to be more accurate – deals – was – first four months were at Sony. As well, I was in Japan for that. But after that, I spent almost a year in Toulouse, France, which is a wonderful place to spend almost a year. It's a wonderful place to spend many years, I would contend. It – at CNRS _____, working with **Raja Chatila** and **Rachid Alami**, and others, as well, too, it was a wonderful time to be able to finish my book, among other things, as well, too, and to deal with the questions of human-robot interaction – in some cases, flying robots that they were working with. I grew greatly because of that, but I have to admit the French lifestyle was an extreme attraction, as well. But Raja is a – not only the current president of our society, but is a – has an incredibly long and rich history in robot architectures, which is something I am very much interested in. And Rachid has worked extensively – and continues to work – in human-robot interaction, which had common interests with the sorts of things that I were doing. So that was the other one. I have worked with the folks at the University of Pennsylvania on multiple grants in ONR – **Vijay Kumar**, for example; have worked with the colleagues at University of Southern California – **Gaurav Sukhatme** – as well. And I – the list goes on and on. I'm currently working with a colleague at Fordham University who I went to school with, did my Ph.D. – **Damian Lyons**, as well – in some of the formal methods. We both have the same schema theoretic framework, which we got from that particular time. The work – I'm working now with folks at Tufts University – **Matia Shoits [ph?]** – on the National Science Foundation work that we're doing for early-stage Parkinson's care management, using robotic systems. So the real key, I think, for a successful roboticist, is finding best-in-class to partner with. Many times, those are people that are right around me at Georgia Tech. **Magnus Egerstedt** and **Tucker Balch**, and others, as well, too, are great partners which I've worked with and continue to work with. But in many cases, you have to look both nationally and internationally. If you want to do the best research, you have to find the best people.

Peter Asaro: And what about some of the grad students or Ph.D. students who are continuing work in robotics?

Ron Arkin: Sure. My very first – I just had lunch with her today. My very first Ph.D. student was Robin Murphy. She has just recently completed a book in disaster robotics, is at Texas A&M now, as a chaired professor, working heavily in search and rescue. Very proud of her accomplishments. Another student was Tucker Balch, who is now a colleague. He is at Georgia Tech – a professor there. I have other students: Alan Wagner, a more recent graduate. We've done some very interesting work together in robot deception. He's at GTRI, the Georgia Tech Research Institute, as a research scientist there, as is Zsolt Kira, who was a – previously was located at Sarnoff, doing work up in the Princeton area, and had decided to come back to the Georgia Tech Research Institute, as well, too, for the environment it provides. Eric Martinson is another student, who is now at Toyota, out in Mountain View, working on particular projects for them. I should know more. Erika Rogers. Erika Rogers was my second student. She ended up as a professor at Cal Poly San Luis Obispo; since retired. Now, she – when your students retire before you, what does that say? So... <laughs> Cal Poly San Luis Obispo, and has done fascinating work in human-robot interaction and the like; was probably the first student I had in that particular space – actually, working in the medical domain, initially, in the interpretation of imagery, and developing very significant IRB protocols to tease certain factors out of how people process imagery, from a cognitive point of view. And I feel I'm missing several others. Lilia Moshkina has worked at the Naval Research Laboratories, now is out with her husband – it's interesting when your students marry each other, and find each other in the laboratory; that's the trouble, maybe, with having a basement laboratory, I don't know – but are now happily married, and she has several kids out in California. I don't know. Those are the ones that leap to mind. I think there's a couple of others, as well, too. But a postdoc I had, I worked with, was Jonathan Cameron. He's at – oh, here's – I knew there was others. Jonathan Cameron is at NASA, working on, last time I talked to him, was aerobots, which are these robots that float up and down in the heavy atmospheres. And Khaled Ali is a student of mine at the Jet Propulsion Laboratories, as well. He was part of the Rover group that is out there. I'm not sure exactly what he's doing at this point in time. But you have to check my Web page if I left anyone out, and apologies to those that I might have missed. I feel bad about that. But I have not produced a very, very large number of students. Some people produce large numbers. I have been careful, I guess, in terms – I'm not saying those are not careful if they produce a lot, but I try and – I've produced a relatively small number, compared to many other people.

Peter Asaro: And now you're also... is it vice provost, or assistant provost?

Ron Arkin: No, no. It's not vice provost. It's associate dean. I'm Associate Dean for Research and Space Planning in the College of Computing. That's enough. Don't give me the vice provost position. I don't need that, too. <laughs>

Peter Asaro: And has that allowed you to further develop the robotics research _____?

Ron Arkin: No. I don't – again, when you're an associate dean for an entire college, you – especially for research – you can't have favorite children, in this case. And it has enabled me to interact with – on behalf of the IRIM Interdisciplinary Center that I talked about, at the EVPR's level, and intercede accordingly. I'm spending a lot of my time right now – we have a new high-performance computing building being built on campus, which is going to be very large. So I'm working with the big data people – data visualization folks – and trying to assist the research enterprise within the entire college, and part of that is the robotics side, but I wasn't hired to just forget about everybody else and focus on that. That's part of the job.

Peter Asaro: And have you also held leadership positions in the IEEE and other parts of –?

Ron Arkin: Yeah, and in several places. I've served two terms on the IEEE Robotics and Automation Society Administrative Committee, the equivalent of Board of Governors. I was one of the three founding co-chairs of the IEEE-RAS Technical Committee on Robot Ethics. I have served on the Human Rights Committee for IEEE Robotics and Automation Society. I think I'm still serving right now, if I'm not mistaken. And – we don't have a whole lot of business, but it's interesting when we do – visa issues, and other things like that, that come up. I've also served in a different society, the IEEE Social Implications of Technology Society – Society for Social Implications of Technology. I served on – one term on the Board of Governors. I am now a Distinguished Lecturer in the IEEE SSIT Society. I served as the chair of the Distinguished Lecture Series, as well, too, for a while. So yeah, I put in a lot of time with the IEEE, and I'm glad I did. It's an important service that you do. But you do have to balance it with all the other things that are going on in your life, professional and otherwise, as well, too, because it can be consuming. But I think it's an important – volunteerism is an important aspect of being a professional, and giving back to the community, I think, is an important thing to be able to do.

Peter Asaro: So, you want to tell us more about your work on the ethics?

Ron Arkin: Well, there's different aspects of the work that I've done in robot ethics. The first and the most significant project, I guess, was the smallest DOD grant I ever had. It was less than \$300,000 for three years. It came about in a very odd way. I was at a Army Research Labs planning meeting, where they were laying out how – “What should we invest in, in the future?” And I had been thinking about all these issues in the background, and I said, “Well, what about robot ethics?” Well, they didn't write it into their program, but the chief scientist – the chief roboticist at the time – said to me, “Why don't you submit a proposal, okay?” So I was on sabbatical. I wrote it in Japan. I sent a proposal in, and I got it back. It was kicked back to me, said, “Who do you think we are – DARPA?” Because I was asking for like three quarters of a million a year, which – that's what I was used to in these projects. These were more NSF-like. So I rewrote it, scaled it back, and cut it down to 300,000 – just under \$300,000 – for three years,

which, as many people know in this field, is a relatively small amount of money, so it was partly a labor of love, as well. And so I got their funding. The first thing we did was a survey, which was part of our statement of work. We wanted to get an understanding of how people felt about military robots. So we surveyed different demographics and tried to understand how the general public – but the most important demographic, or the most significant responses we got, were from the roboticists themselves. It was very interesting, seeing the responses, and people weren't happy with lethal autonomous systems, whatever that meant. We did define it rigorously, actually, in the survey, so that people could get an understanding of that. And then the second part of the survey was that how, if these systems were going to be created, we could develop software which could potentially enable them to perform better than human war fighters under these circumstances. So I created an architecture – delivered a reactive architecture, actually building upon pre-existing architectural work – integrated in several new components. One was something called the “ethical governor,” which is a bolt-on piece of software that basically restrains the system to perform within predefined bounds, as dictated by the relevant laws of war and rules of engagement. We had another component, which was the ethical adapter, which dealt with one of the emotions – and we had a lot of work under our belts in emotions, at this point in time – looking at the moral emotion of guilt, and how that can be used to restrict the behavior of the system when it recognizes the fact that what it's actually doing is different than what it is expected – it thinks it should've done, in those cases, which would restrict the use of weapons under those circumstances. The other one was a responsibility advisor. It became clear, in my discussions with philosophers, among others, that it was crucially important that we be able to attribute responsibility to human beings in this particular process, so that no one could hide behind the robot and say, “The robot did it,” and we made explicit, in a prototype, in a proof-of-concept type of system, how we could establish that particular connection through a **gooey [ph?]** – basically, in two ways. One was pre-mission, where, in other words, before you send the robot out, you better say, “I'm taking responsibility for sending this robot out to carry out these particular actions,” or don't do it. And second is also in the use of overrides, where the human can step in and either – equivalent to hitting an emergency stop button: “Don't shoot!” Or do the reverse: Allow the human to say, when the robot is resisting firing, to come in and say, “Oh, you should shoot,” and take responsibility, and apply the lethal force under those sets of circumstances. Another component that we had was the behavioral design, which actually tried to integrate the ethical behavioral control within the behaviors themselves. We did not get as far as I would've like to in that case, but I was significantly underfunded to be able to do that. But that doesn't diminish the work that the other pieces did. We developed and tested scenarios that we found from real-world events, as much as we possibly could, and showed – at least, to my satisfaction – that certain things are potentially possible using these kinds of systems, which has led to great debate worldwide. And, like I say, as I've often said, the discussion that my research engendered is as important, if not more important, than the research itself, and I have never made the claim that, “This is the way to accomplish these particular goals.” I say it is one possible way, but I also say that others should be investigating this to find if there are better ways to be able to do this. So it's just a first foray into this particular space, and it has led to all kinds of interesting – and sometimes bizarre – interactions, through commentary and the press, being in front of the United Nations and having

the ability, as you did, to speak on behalf of what you believe is the right way to approach these particular problems, presenting in front of the assembly of the International Committee of the Red Cross, which do the Geneva Conventions, to talking in front of pacifists who say the only ethical robot are the – a truly ethical robot would throw down its gun, to begin with. So I have waded way up to my – in some cases, up to my neck, in terms of the controversy associated with that. That's because I feel it's important, and passion – I feel passionate about it, because I do believe it can lead to the saving of noncombatant lives, which is a view not necessarily shared by all; or even if they did believe that, there may be attendant risks associated with that, which would say that, even if you could do that, it might not be the right way to go. The debate is important, it is ongoing, and it needs to continue. And I tend not to be prescriptive, as some other folks may be. I'm just saying this is one possible thing that can be done. I have stated, in print and in public, I am not necessarily averse to a ban, but I think we may be able to do better with respect to noncombatant casualties, and that's my hope. That's my hope, that we can potentially save lives through the use of this technology, and that's where many of us part company, in terms of the risks versus the potential benefits, which still are unknown. And we're still wrestling with the definitions of what is an autonomous robot, let alone a lethal autonomous robot. What is meaningful human control? Who wouldn't want meaningful human control? Everyone wants meaningful human control, but at what level, and what is the appropriate restrictions that should be put into those systems? Those are the sorts of things. And I've extended that work more recently into the work that I'm doing with management of early-stage Parkinson's, leveraging the ethical adaptor and the ethical governor into that work. I think I may have mentioned this earlier, but the point being that we're trying to deal with a new problem, which is a five-year NSF grant that we have, joint with my colleagues at Tufts University, in terms of maintaining dignity between human-human relationships. Most of the time, you hear about human-robot interaction: It's between the human and the robot, and how are we going to interact? And we want to maintain dignity in that, as well, too, but quite often, when humans interact with each other, in certain cases the relationship goes poorly for whatever reasons, and in early-stage Parkinson's disease, due to this phenomena called "facial masking," where people lose the ability to effectively express emotion on their face, and their voices become raspy, so they lose the ability to express emotion through their voice, a phenomena referred to as "stigmatization" occurs between the patient and the caregiver. Even if the caregiver knows that that's the case, they just don't feel that this person is being responsive or caring about what they're feeling, and the patient becomes more ashamed and embarrassed by this lack of congruence between the beliefs of the two. And so we're hoping that through the use of a robotic companion, or a presence in the room, we can enhance the environment through the use of nonverbal cues by this robot, through both kinesics and proxemics body language and facial expressions, among other things, and said distance of separation, that will assist in restoring and preserving the dignity in this relationship. So we're leveraging some of the work in this particular space, expanding the moral emotions, as I mentioned earlier, from guilt, now to incorporate shame and embarrassment, and also, potentially, empathy, and foster empathy in the caregiver, and try and reduce shame and embarrassment in the client, in this case.

Peter Asaro: So you mentioned some run-ins with the press and the media, and I know that the ethical governor has gotten a lot of media attention, and so has the robot deception work that you did. If you could just talk a little bit about what it's been like, trying to get robotic research out into the public?

Ron Arkin: I always say that if, after interacting with the press, whether it's a film form or whatever – and hopefully this, as well, too – if we get 80 percent of it right, we're doing well. There's no way it's going to be 100 percent right. The best article I ever had was 95 percent, I would contend, which was from the *New York Times*. They came and visited our laboratory twice. They did fact-checking. They sent a photographer down. They invested heavily in that. In many other cases, things are – fact-checking seems to have been thrown out the window in recent times, which is a bit disconcerting, but I guess it's the rush to press. And even worse is when articles are done by the game of Telephone, you know? They read one article, and they yank something out of that which looks controversial or significant, throw a picture of the Terminator on the page, which I give talks on how to not build a Terminator. I'm not interested. I don't think anyone should build a Terminator. The Terminator's the wrong thing to be building. No one should be building that thing, but, nonetheless, here we have those sorts of things. So the way in which a kind of – it's like a fire running along the ground takes root. But the important thing is, even with that – you know, I could take the choice and say –

<break in recording>

Ron Arkin: – because of the misrepresentation quite often that occurs from the work, sometimes deliberate from others, but independent of that. As I've said, the discussion that the work engenders is very important, so if it gets people thinking about it and talking about it, and trying to understand and get their head around that, even if it is sensationalized, which is a problem with robotics; sensationalism occurs. Getting to the deception work, very interesting things have happened in that space, as well, leading to articles. We created – a simple article appeared in the *International Journal of Social Robotics*. We had many limitations, saying this is not the final word, this is the first steps, more research needs to be done before we can say anything really substantial about the – and then some articles say, basically, the sky is falling. Chicken Little is here, and it's the end of the world. It – our names will live in infamy because we taught a robot how to lie, and the like, which seems a bit extreme. And the same time, we win an award from *Time* magazine as one of the top 50 inventions of 2010, in this case. It wasn't really an invention, either. It was just a paper describing an algorithm, per se. So, I mean, it can go off in both directions. Some cases, like the *New Scientist*, got it right, and talked about how important it is to get a deeper understanding of theory of mind, which is what we were using, to be able to understand when and when not to deceive. But the other side of the coin is when we started doing work in squirrel deception – nonhuman models of deception. We did some studies of the Eastern gray squirrel, and we found that they basically used evasive behaviors to hide their nuts from other squirrels, so that they would patrol empty caches when someone was observing

them, and we got lots of press on that, as well, too. But everybody thought that was really nice, because squirrels are pretty, and you put a picture of a squirrel on the page, everybody loves it. I was just saying, there's so much... nonsense, maybe, or irrational reactions to these sorts of things, the way it's approached, that it's a little disconcerting at times, but you get used to it. You try and be selective, in terms of who you're talking to, and you try and find what the spin is on the article before you engage, and you don't talk to everybody, so you don't get burned more than you would normally get burned. All those things are important lessons, and I would contend that media training is probably a good idea for most roboticists that go into this form. And this even happens when I'm asked to comment on someone else's article, as well, too. Even then, things are taken out of context, in some cases. So dealing with the press is an art in and of itself, and most roboticists are not trained or ready to be able to do that. So, just as fair warning, if you're going to do that, be prepared for things that you might not normally encounter, and at all costs avoid ambush reporters, as well, too, who lure you in with the illusion of one story, and they're actually trying to tell another. The most interesting anecdote was when *The Daily Show* called me up, wanted to do a piece on the military robots – the ethical robots, as well, too. I'm a big fan of Jon Stewart, but they told me that Samantha Bee, who did a field piece, was the one that we would be speaking to. The producer of the show called me up. They wanted to fly me up to New York. I got like three calls from the producer, and I said, "I love your show. I also know how Samantha Bee does field pieces and the like, as well, too, so keep me laughing and enjoying these things, but I can't participate in these sorts of things." So it's important to know, at least from my point of view, that not all press is good press, contrary to whatever that old adage is, that, "As long as they're talking about me, it's a good thing." I don't think that's necessarily the case, and it affects public opinion and the like. But, like I said, in certain cases, where we're dealing with life-or-death issues and intelligent robotic systems and lethal systems, it is important that we engage in the discussion in a public manner. And I've been involved in numerous debates, and I will continue to be involved in numerous debates, to make people think about these particular questions. I'm not the one that should be telling them what to do, and no individual, and perhaps no individual organization, should be able to tell them what to do, but we should present our viewpoints and let people make a decision as to what they think is the right way to handle this new technology. And that's not just in lethal systems. That's in sex robots, or "intimate robots," as I prefer to call it; healthcare robots. There is an – I – there was a very interesting study from the European Commission which came out two years ago, which showed what kind of robots should be banned. And they interviewed 27,000 people in 27 nations in Europe, and I believe the number was 47 percent said that elder care and child robots should be banned, which just blew my mind, and basically that indicates to me that we, as a community, are doing a very poor job of educating people as to what the risks are of certain types of systems. Military was down like at 17 percent or something. Healthcare was like 30 percent. I don't know the exact numbers, but I'm just saying. When I saw this, it was actually in a healthcare workshop that I was attending. I was just shocked, and so I said, "I got to see the report," and it's freely available on the Internet, as well, too. Most interesting.

Peter Asaro: Okay, so what are, for you, the biggest sort of theoretical questions or technical issues facing robotics in the coming years?

Ron Arkin: Well, there's two different aspects of that. What is – what are the issues preventing large-scale dissemination of the ideas that we in academia and elsewhere are creating? And as I've often said, many of them are not the intelligence of the robotics questions. It is things like safety engineering, reliability engineering, power. These systems, if we're going to get broad public acceptance, have to be as reliable as our automobiles, or laptops, at the very least, even if they do crash from time to time. You need 24/7 operational capabilities if you really want them to work. That's important, but that's not so much the robotics questions. Systems integration is absolutely crucial to be able to put these things together in meaningful ways. Better understanding of human behavior, again going back to my roots of ethology and neuroscience. If we can elucidate more how the brain functions – which we are beginning to make significant progress, with new instrumentation that is available – how can we put that to work in intelligent robotic platforms? That's an important aspect, as well. Computer vision remains daunting; it has for decades. Significant progress is being made, but we're still far, far away from anything ever approaching general vision. Ethical reasoning, I think, is an important aspect, as well, too. Keep in mind that I don't believe we need full human moral faculties to do the sorts of things that I'm advocating in this particular case. I'm dealing with compliance rather than reasoning. Those are important aspects of the problem. Machine learning is something I'm more worried about, perhaps, than most. I don't want robots necessarily to learn in **unexpectable [ph?]** ways, except in bounded environments, as well, too. Tremendous progress is being made with new technologies in machine learning, new algorithms in machine learning. Where they will lead, it's hard to say. And then it gets to the questions of regulation for domestic use, whether they're for drones and privacy issues, and all these other things, as well, too. If we can create it, will we be able to use it is an important aspect, as well. So those are the kinds of things I focus on. The good news is CPUs and processors and the like – new sensors – are rolling out at a highly significant rate. We can expect new sensors to be game-changing things, like the Kinect and other sensors were. **Flashlight-R [ph?]** will make a difference, as well, too, and the cost will come down, so new applications will unfold. So there's hope for many significant advances. But the field, as I think I mentioned earlier, but the field goes in fits and spurts. Quite often there is a fundamental advance or a paradigm shift that occurs within a field, and then, for the next 10 years or so, people catch on, move that forward, and make steady progress with it, enhancing it, and then another paradigm shift goes on. Some of these things were behavior-based robotics, then probabilistic robotics, and now we're in the phase where we're making continuous advances, but slower, and so I always wonder what the next big computational or intelligence breakthrough will be, which will change the paradigm. I don't believe it'll be a positronic brain, a la Asimov, but I do believe they will continue to occur, and they are completely and utterly unpredictable, as scientific research often is.

Peter Asaro: So is there any other projects or work or things that we missed that you want to talk about?

Ron Arkin: Oh, there's dozens of projects, but the most recent ones are the Parkinson's stuff. We're doing work in the deception stuff that we talked about, as well, too. We're looking at other deception right now, as well, too, which is deception for the benefit of the mark, so that you can tell me after this, "That was a great interview, Ron." <laughs> Whether it's true or not, I'd probably feel better because of it. And the same sort of thing if someone is bleeding out on the floor. You don't want to say to him, "Oh, you're all messed up." You'd rather, perhaps, try and reduce the risk of shock occurring by saying, "Help is on the way, somebody get help," to do these sorts of things. There are times, if you're a non-Kantian – if you're a consequentialist, in this particular case, where you're judged by outcomes – lying may be – or deception may be – appropriate. In other schools of ethical reasoning, that is utterly and completely false, and that's part of the reason why ethics is so hard. So there are four projects that I'm actively involved in. One is dealing with mental rotations in primates, and the formal methods, as well, too, for the Defense Threat Reduction Agency, in counter weapons of mass destruction, where we may be – it's not a question of if, it's a question of when, we find ourselves under attack by a biohazard, a chemical weapon, a radiological weapon, or a nuclear weapon, and we may have the opportunity to deploy robots – or robot – to address that. I refer to our robot, in this particular case, as Jack Bauer, from the 24 series. We're trying to write Chloe, I guess, to make sure that before Jack goes in there, there's a high probability of success, or don't send him in at all.

Peter Asaro: And we always wrap up with this question, which is, what's your advice for young people who are interested in a career in robotics?

Ron Arkin: Get a solid foundation. It's a question of how young you're talking about, but if you're in high school or younger, pay attention to the basics. Pay attention to mathematics. Pay attention to chemistry and physics. Don't be seduced by the quick pull to creating cool little robots that are moving around. You can do that. That's okay. But you need to have a deep understanding of the fundamental principles upon which robotics, which *is* a science, is based upon. Pay attention to that, and work hard, and you have to have a passion for it. If you don't have a passion for robotics – and almost any field, for that matter, as well, too – you probably shouldn't go into it. Don't – although you can make a good living in it, it shouldn't be about money. It should be about joy, fulfillment, and creating new knowledge.

Peter Asaro: Great. Well, thank you very much.

Ron Arkin: You're most welcome, Peter. Thanks.

End of RonArkin.mp3