AI and You

Transcript

Guest: Elizabeth Croft, part 1

Episode 144

First Aired: Monday, March 20, 2023

Hello, and welcome to episode 144! My guest today is Elizabeth Croft, Vice-President Academic and Provost of the University of Victoria in British Columbia, Canada, and expert in the field of human-robot interaction. She has a PhD in robotics from the University of Toronto and was Dean of Engineering at Monash University in Melbourne, Australia.

The field of robotics has a very interesting relationship with artificial intelligence, in that it is something we are clearly nowhere near as good at as we would like to be, and here the kind of standard that we are going off is from the numerous presentations afforded by fiction of one medium or another, where robots are presented as being capable on a scale ranging up to indistinguishable from humans. (Think *WestWorld*.) Clearly that capability has far more to do with what is (a) entertaining and (b) affordable to makers of movies and television programs than it does with what is realistic in engineering, but it has also implanted a deep expectation in our collective subconscious. And today we see a lot of presentations of robots in the media in either research demonstrations, prototypes, or actual use, and we have all been treated to displays of what Boston Dynamics robots like its almost humanoid Atlas robot can do. So there is a modern updated expectation that we are just around the corner from robots forming some integral part of our society that might go up to the scale of market penetration of the hordes of robots in the *I, Robot* movie, but we don't really have any good ideas of what to expect and when. Even if your interest in AI is purely disembodied, and you have no interest in a robot making your breakfast or fixing your plumbing, there is considerable reason to think that generally intelligent AI depends upon an understanding of our physical world that will only be attained through embodiment, in other words, gained from smart robots. And we have explored that several times on this show.

One of the things I mention near the beginning is the DARPA Grand Challenge, which was a contest launched by the Defense Advanced Research Projects Agency in the United States in 2004 to encourage autonomous vehicle development, reckoned to be one of the starting points of the modern AV era. Anyway, let's get right into the interview.

Elizabeth Croft, welcome to the show.

Thank you. It's a pleasure to be here.

So, for the benefit, especially of people that are contemplating career choices or changes, how did you get into the field of robotics? How do they enter your world?

Oh, okay. So I started out in my undergraduate degree as a mechanical engineer and I really didn't know where I was exactly going. I was actually very fortunate that one of the academics at the university that I attended, which was the University of British Columbia, even while I was in high school, he was very interested in talking to me. He was a friend of my parents and one of the things that was really wonderful is he saw me not as a teenage girl, but as a smart kid with possibilities, which I thought was really very cool when I look back, that in a time when women weren't really thought of as being potential young engineers, he actually saw me as a potential young engineer. So that was very cool. He spoke to me a lot. He was really into thermodynamics in jet engines and he gave me the plans for the Boeing 767, particularly around the jet engine development that he was working on. And that first of all opened up my mind of the opportunity of engineering and how fantastic that you could design and build these amazing things. Further forward, I started to get very interested during my undergraduate degree in biomedical engineering. And I had a co-op work term at TRIUMF, which is the TRI Meson factory in the UBC endowment lands at the time, still is there. And the job was working in remote handling, which was using robots to take and remove targets from the underground beamlines so that people don't have to go in there and be exposed to radiation. And so I got involved with this group and the reason I was so interested in the group was actually because they were also developing robotics for assistive purposes. They were working with the Neil Squire Foundation at the time, and they were designing a robot that was kind of dual purpose. One part of the purpose was to do these beamline target pickup and removal, like place and removal. But the other purpose was to design a robot that someone could use if they had lost mobility and it would extend their mobility. So I thought that was very interesting and it kind of moved me from being very excited about biomedical technology into, oh, there's this whole other place where you can do good in the world and it's called robotics. And then just at the time that I graduated, I actually wasn't working in the industry for a couple years doing something else that was really super fun, which was motor vehicle accident investigation, a whole other story. But then I had a scholarship to go and do master's and PhD studies. And at that time, artificial intelligence was sort of getting into, you know, it was in *Wired* magazine or whatever the precursor of that is, I can't remember. I just remember seeing artificial intelligence on the front of a magazine that I was reading and going, "Ooh, what is this?" And starting to read about that and going, "I need to do something that allows me to get into this area." So I chose to do a master's degree on neural network control of autonomous guided vehicles. So it had the mechanical electrical stuff that I love and the programming I love, but I got to learn about neural networks and I read all about them and I learned to design and program my own neural networks because, at that time, PyTorch was not a thing. So I read the books about what a neural network was and I turned that into programming and I designed a neural network that could control a little HEV and that was super exciting. And then I did a PhD. I decided that I was actually interested in robotic arms and so I went off to Toronto and I did a PhD in optimal motion control for a robotic manipulator. So I had this kind of background. And then from there I came back to UBC as a fresh, wide-eyed assistant professor working in industrial automation. And that was very interesting but I really got pulled back all the time: The limitations of robotics and the limitations of AI, it always seemed to butt up against how does it interact with people, because that is where the actual interesting things that at least I thought were super interesting happened. And so I kept on getting interested in human-computer interaction and human-robot interaction. And when I found the human-robot interaction community, it was like, "I have found my people." There are psychologists here, there are computer scientists here, and they like engineers. Oh my goodness, these are the good people. These are the people that I can talk with and that we can explore ideas together. And the interdisciplinary interest in that community was so refreshing to me because most engineering communities, most scientific communities are very siloed and they want to go really, really deep in one thing whereas I like the depth of being a very good mechanical electrical engineer, good programmer. But I liked the psychology aspects, and the AI and psychology have such an interesting intersection there. And I liked the fact that people do the darnedest things and how people interact with robotics and with AI is actually fascinating. So that's how I got here, I guess. And that kind of led off to a whole bunch of other interesting things.

Wow. So much to talk about there. When you moved into the more cross-disciplinary area, did you find that there was more female population in that area or more acceptance? In the area that you're talking about, it has traditionally been difficult for female engineers. What was your experience in that regard?

Absolutely. Definitely both the human-computer interaction field and the HRI field, there are more women in those areas, and that is nice. I mean, it's good. I think it's probably still more men, many more men working in robotics and even in that field than women, but it's more balanced, I would say. And I think it has more openness to other disciplines. Because you're bringing in people from psychology and neuropsychology, which does tend to have more women in it, so you get those cool interactions. So yeah, I think that's one of the things that does make it just sort of naturally attractive to me.

I just wanted to pick up on something you mentioned earlier and make sure I understood it. You were talking about robots that went and fetched targets from beamlines. Are we talking about particle accelerators?

Yeah. So TRIUMF is a trimeson factory. It's a cyclotron, it's a particle accelerator. And when you work there, you have to wear a dosimeter. And then if you go into certain areas, they're hot. And if you go into the beamline, that's very hot. And then if you get too much dose, then you can't go to work the next day. So sending robots in rather than people is a good HR practice.

How big was it? Are we talking Death Star scale beams?

The TRIUMF is a big particle accelerator. I mean, I'm not a physicist and I work there as a mechanical engineer but it's a significant and very famous particle accelerator located just next to UBC and people from all around the area go and do experiments there. And they do a lot of stuff important for mining applications, important for medical applications, and important for pure physics. And some brilliant, brilliant people work there.

The work you were talking about doing with neural networks and autonomous vehicles, where was that relative to the DARPA Grand Challenge?

Oh, that was before. This was the 90s. I'm old. So this was just when there was a researcher who was doing No Hands Across America, which was the first... I'm blanking on the name, but there was a great researcher, I think they were from Carnegie Mellon, and they put basically - it was not a 386-type computer - a very simple neural net and it drove a van across America. I think 95% of the time it was on a highway and it was the yellow line in the middle and the white line on the inside, and it was driving and basically, it was using a neural net. So it was during that time when we were still trying to figure out how many layers, and it was basically reinforcement learning. People were starting to think about all sorts. There were all sorts. Clustering was happening and people were thinking about recursive. And this was around 1990 and the hype cycle had gotten really big again and then it kind of died off. And now it's really big again. But that was the time when we were doing it.

And there was also a researcher who, a few years before that, had been doing that in Paris. And it's always fascinated me that we had that kind of capability then, and that if you look at that in the light of today's capabilities, it seems impossible that a car should have even been able to stay on the road with the hardware that we had back then because it is so hard now, not to get it to stay between the lines, but to do everything else, like avoid pets and things.

Yes. But you've hit upon the problem. Because the problem is not staying between the lines. Ninety-five percent of it, and that's what that guy from Carnegie Mellon showed us, 95% is really actually pretty easy. When you are driving, how much of your brain are you actually using for most of the time? It's all the edge cases. It's all about the dark person wearing black, carrying a black umbrella on a rainy night, darting out in traffic. Can you perceive that? How do you know that is actually a thing or a person, not a shadow or something else, or a black paper bag or something? So I'm not surprised that we could do that then because we don't need a lot of computation for those things that are pretty consistent and regular. In fact, probably we could have designed a controller, like just a basic controller that would've done what the neural net did. It's just not that complicated. It was interesting. The thing that we need to improve on - and this is the whole thing about machine learning versus general AI - is that ability to understand and to respond and to get these edge cases that are so hard and so complicated and so infrequent that how the heck are you going to learn them? But we respond to them; we get on the brakes. And sometimes we don't, as I learned when I did motor vehicle accident investigation. So the difference, I think, between 1990 of course and now is, we've got GPUs, so we can do so much more. So we're really, really good at the easy stuff. We're getting better at the more frequent stuff because we can do segmentation and we can have more structures where we can respond to different rates of information and we can kind of segment things.

So we've pushed out further into that fat tail of edge cases, but we still don't know how long that tail is.

The tail is forever. How long is a piece of string?

Right. And it seems very reflective of tropes that we encounter all over the place in AI that we can solve the narrow cases easily, but that the useful things lie out in this world of generalization where we just don't know the scale of that. And if we carry that now into the field of robotics, tell us about your current line of research. Maybe we can start from the level of how you envisage this being used first of all. What sort of robots doing what sort of interactions with people is the kind of thing that drives your work?

So - for a long time, not always - what I have been interested in for quite a while is really the part where a person and a robot get into an interaction. So, for me, this really started out with the very most simplest interaction I could think of, which was object handover. And it's a really nice interaction to study. And a lot of colleagues have also jumped on this and studied it as well. It's nonverbal. Because there's a whole bunch of people doing language models, which is fantastic. And it's physical so it's a dynamic interaction, and we as people are pretty darn good at it. We hand stuff back and forth to other people all the time and we never drop things. Well, except in Olympic races in the baton.

But I've been thinking about this before our talk and it's actually harder than it looks in that if you look at how often young kids drop things that you're handing them or vice versa, then it's one of the skills that we learn relatively later compared to a lot of hand-eye coordination tasks, it seems to me. So we're talking about humans passing things to robots or robots passing things to humans or both?

Yes, both. Absolutely. And I guess the other thing you mentioned is that you're talking about hand-eye coordination, but handing things over you can do with your eyes closed. So the other thing I liked about this one is because I could take vision out of the picture - at least in my first studies - I could really zero in on what was the relationship and what were the rules of engagement that we had to understand about this interaction, and then pull that back into what that means for robot design and control. This is sort of my biomimetic phase of life where I was thinking, "Let's look at what people do and then see if that can translate into how we think about robots." So we made a baton, which was basically an ATI strain gauge, two boxes with force sensors on it, and my grad student, Wes Chan built this. And then we gave it to people and we got them to do handovers over and over again. So we had a huge data set of handover. We knew time and we knew force, grip force, and we knew load force. We knew who was holding the baton because that was all captured by this baton. Then we plotted the data. Actually, I just remembered sitting in the lab with Wes displaying the data on the screen and Mike and Chris were there. We were looking at the data and we're going, "Ohhh." And it really showed itself the rules. The rules are, the giver is responsible for safety. They do not let go until they know that the receiver has accepted it. The receiver is responsible for timing because it depends how quickly they take the load that the giver will let go. So those rules, we didn't make those up. Those rules appeared in the data. Those rules came from what we saw out of the data for that baton. And then what we did is we took that, and Wes made a very simple controller that basically controlled the force that the robot's gripper was, based on how much load it was experiencing. Because as soon as the receiver took the load, the robot let go. But we actually saw in the data that the giver does not let go until the receiver takes even just a little bit more of the load. They kind of tug it a bit. You know that they've taken it. And then we, once again, tried that with a whole bunch of people. And we tried different parameters of this and that's how we kind of tuned it. It was the best demo we ever had. People would come in and the robot would hand a bottle of water to people and everybody went, "Yeah, that feels really natural." So that was our first kind of foray into this interaction about this. And then we then did stuff about, okay, we've got the handover, let's get back to other cues; because this is really a cue, right? This is an implicit cue about when I'm ready for you to take it, and when I'm ready to let go. These are these implicit cues. People have looked at the cues like, where the bottle is, where I'm handing it to you. So that's a cue. And then the other cue that we were looking at is where is the robot looking. Now, robots, do they look? Where people perceive in their minds that the robot is looking, and we did a study where we put a whole bunch of people through, you know, they come in and the robot hands them the bottle of water. And we had three cases. There was, the robot looks at where the bottle is, the robot looks at the bottle and then looks at you to sort of say, “Hey, here’s your bottle,” and then there was what I like to call the teenager handover, which was the robot looks down and hands you the bottle, you know, “Here’s the car keys.” So we tried those out. What was very interesting is that while people liked and had a preference for the robot making what we call eye contact - people really felt that was engaging - people cued on the robot looking where it was going to put the bottle. And actually launched into, and we had a light curtain so we knew when they were reaching, they launched into the reach sooner when the robot was looking at where it was handing the bottle. So this was really interesting because what it told us is, hey, there are implicit cues that we use that if robots understood them, then they would be better at interacting with us. So this got me thinking about, and we’ve done a whole bunch of other stuff since then, but what I’m really interested in is human and robots’ subtle behaviors and how information flows from those subtle behaviors, those cues. What information? How do we get it? Because there’s a lot of it. People use a lot of body language, a lot of subtle cues. I talk with my hands, and people use their eyes. If you’re talking to someone and they’re not looking at you, you know you’ve bored them to tears. So there’s all of these cues that’s important information that we use in interaction, and robots to be useful to us need to be able to both pick up on those cues, but then also to be able to send those cues themselves. So I’m interested in information transfer.

I’ve got so many notes on that. And it’s so important because even if we were to have something like a robot barista in some robot Starbucks handing you a cup of coffee, there are so many things that play into that to have that handover. Now, let’s just look at the vision question for a moment. So the robot has vision and what I’m taking from what you’re saying about where it’s looking is that that’s more of a function, not for the giver, but for the receiver to know that the giver is engaged properly with this interaction. Because robot vision doesn’t have a central field of view, right? Humans, we have that because we just couldn’t process the amount of data if the entire field of view was at that level of resolution, so we’ve optimized for that. But these robots, the “direction that it’s looking in,” with air quotes around as you’re doing there, is not a function of its vision system, it’s a signal that it’s decided to give. It’s picked a point in its field of vision to signal, “I’m looking here” as a communication mechanism.

Bingo.

There was a concept called the OODA Loop introduced years ago by someone who was, I think, in the Air Force, and I inserted another letter in there to make it an OOPDA Loop because that sounded good. But also to say that our cycles of interaction are Observe, Orient, *Predict,* Decide, and Act. And it seems that that comes into play a lot in the sort of things that we’re talking about that when you were describing this handover moment where we’re feeling, has the receiver got this thing by the variation in the load on our fingers, that that is going through some sort of predictive cycle. Do you look at predictive cycles in this kind of interaction?

Absolutely. In fact, one of my students is right now working on prediction and production of robot motion. Looking at and predicting what humans are going to do. And this gets into this whole question, you know, the papers by Anca Dragan and Sid Srinivasa back about 10 years ago, all about legibility and predictability. Because we are constantly as people predicting what's going to happen next. Somebody moves to the left, we move to the right and then we do the, "After you, Alphonse" thing. But we're constantly adapting, as you say, looking, predicting, sensing, thinking, making a decision, making a move. And we're just constantly doing this. So being able for the robot to fit into that loop, because it's got to do that and it needs not to be out of phase, right? Because you could get into these weird things where if your cycle is wrong, you could actually do the thing where everybody moves to the right on the sidewalk, everybody moves to the left of the sidewalk and then we run into each other sort of thing. So how do we understand what the other party understands and how do we respond to that?

And you're uncovering these cycles of human behavior in these interactions. To what extent was that already investigated by psychologists, like child psychologists?

So absolutely there's a lot of foundational work that's in the psychology realm. And really, I'm interested in how far that goes over into the robotics realm because there is also, and this is quite interesting, at least I think it's interesting, that there's a different set of expectations for robots. So we're thinking about, well wouldn't it be better if a robot was more legible? In other words, seeing what the robot is doing, I can predict where it is going or what it will do next. So that's legibility. Or is it more important that the robot is predictable, which means knowing what the robot is doing, does its behavior make sense? So one of them is motion to goal or action to goal and the other one is goal to action. And so people like robots to be predictable, but that does not make them legible.

Can you expand on that? Because I've almost got it, but maybe an example would make it clearer.

Okay. So the robot has, and this is from Anca Dragan's paper. So the robot has three objects in front of it and it's going to reach and pick up one of those objects. And those objects are clustered together. If the robot just moves directly at the object, it is difficult to understand which one it's going to pick because they're all clustered together. However, if the objects are oriented in a way that you could come from the right and get the one on the right-hand side, or you could come from the left and get the one on the left-hand side, then coming from the right, looping around, would make your motion more legible because you could see that the robot is reaching for the right-hand glass or the robot is reaching for the left-hand glass. But that is a legible motion because it's more expensive, right? It's not the most optimal motion; it's going around. So that's legibility. Predictability is, I know that the robot is reaching for the right-hand glass, and it goes straight for that glass.

That's the end of the first half of the interview; This one has been split up again into two halves because the interview was so interesting that it ran long enough that we are going to put it into two episodes. I really liked Elizabeth's evident passion as she talked about how she got started in robotics, and I can tell you that she had to have a lot of it to make it as a woman in that field of engineering at the time she was entering the field.

In today's news ripped from the headlines about AI, I've talked so many times about various incidents and how they illustrate the penetration of AI into the mainstream of thought and application that it's practically redundant now; AI is here, it's here to stay, and it's everywhere. So I'll leave that commentary out this time and just point out that I couldn't resist telling you that Rentokil, the world's largest pest control group, in other words, people you call to get rid of cockroaches or ants, is, and I quote, "piloting the use of facial recognition software as a way to exterminate rats in people's homes." Rentokil said it had been developing the technology alongside Vodafone for 18 months. I know, I also am wondering what Vodaphone has to do with rat extermination; that was not specified, let's keep going. *The Guardian* reports that the surveillance technology, which is already being tested in real homes, tracks the rodents' habits, and streams real-time analysis using artificial intelligence. A central command center can then help to decide where and how to kill the rats caught on camera. I know, I know, the jokes write themselves. Maybe this is how Skynet really gets started. One little blown fuse, rats get switched to humans, and it's hasta la vista, baby.

One obvious takeaway from this is the learning that rats have different faces. If you were thinking up until now that they all looked the same, well, check your privilege, apparently it's not just their mothers that can tell them apart, but now AI as well. Rentokil's chief executive, Andy Ransom, told the Financial Times: "With facial recognition technology you can see that rat number one behaved differently from rat number three. And the technology will always identify which rat has come back, where are they feeding, where are they sleeping, who's causing the damage, which part of the building are they coming from, where are they getting into the building from, whether it's the same rodent that caused the problem last week." Perhaps there's some risk of people getting attached to the rats if the AI gives them different names; I don't know if that's a possibility. If you're looking to be the first on your block to have robot rat terminators, get in line. Rentokil is targeting "cities of the future" in countries that could soon experience a pest population boom, such as China, India and Indonesia.

Next week, we'll conclude the interview with Elizabeth Croft, when we'll talk about robot body language, dealing with a squishy world, and ethics for robots. That's next week, on *AI and You.*

Until then, remember: no matter how much computers learn how to do, it's how we come together as *humans* that matters.

http://aiandyou.net