AI and You

Transcript

Guest: Elizabeth Croft, part 2

Episode 145

First Aired: Monday, March 27, 2023

Hello, and welcome to episode 145! Today we conclude the interview with 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.

Last time, we talked about how Elizabeth got into robotics, autonomous vehicles, and introduced her current research, which I find fascinating because of how many questions and lines of exploration it opens up with respect to how we will use robots, which is object handover between humans and robots. It sounds obvious and elementary, but you could go in so many places with this, like how we look for affordances on objects, like handles, or whether an object should be thought of as a single object or multiple ones, like if it's a coffee mug that contains coffee or a jar with a lid. And there's so much more to this simple act of passing something from one actor to another that we'll learn more about in the conclusion of the interview with Elizabeth Croft.

My two daughters are both left-handed and I think there's something also in here about the difficulty that they face with objects that I didn't realize how many objects are designed to be right-handed. Saucepans. Actually, I made a note about cars when you were talking about directions the robots were looking in because there's a car in Japan that they put eyes on - just the sort of thing that the Japanese would do - that can swivel around. And they found that when pedestrians are crossing in front of it that they feel a lot safer when they can tell what direction they *think* the car is looking, even though it's got 360-degree vision.

Does the car see me?

Yes.

Yeah, is essential if the robot sees you, then you feel more seen, you feel more comfortable. So we did a study with AJung Moon who was the student who did this. She's brilliant. And what she looked at is turn-taking with robots. And so she had three cases. One of the cases was the robot's blind, so it's just doing its thing. So the worker, the human and the robot are both accessing the same facility. It's a pot of stuff and they both have to access it. And the three cases were the robot was blind; the robot would immediately stop if the human came into the area; and the third case was the robot would do a hesitation, would do a little jerky motion when it saw that the human was there and then it would pull back. And the interesting thing there is, again, predictability and legibility. The people felt that actually the blind stop was very predictable and they liked that. But they also saw the hesitation as very communicative. It said, "The robot knows I'm there. The robot is paying attention to me and the robot's being polite about letting

me have access first." So although it was not more efficient, it communicated something about the robot's understanding of the human being there. And so I think there's this thing about cues that communicate the robot's awareness of the partner that it's with.

So this body language of the robot, if you will, has communicated to the human that the robot is aware of the human; at least the human interprets it as that.

Absolutely.

To what extent does that have to be borne out by other actions? Because that's there purely as a communication mechanism. The robot is aware of the human to whatever extent it's aware of the human. But if it's not sufficiently aware of the human and it gives these cues anyway, does there come a point where the human thinks, "Oh, I can't trust this thing anymore"?

Well, that's a really important question. And I think this is the thing. Communication builds trust, or it can destroy trust. Because if you say one thing and do another thing, that's a destruction of trust. Whereas if you say something and then your actions support that, then you build trust. So communication is that first part of building trust because it's explainability, I understand you. So if I understand you and I know what you've said or what you've communicated, and then your actions follow up, then it is the pathway to building trust. And it's absolutely true that if a robot hesitates and says, "I know you're there," but then goes in, then we have the problem of, something has been communicated and the robot is like, "I don't really care. I know you're there--"

"And I don't care."

And so I think it's very important that communication is purposeful and that it is consistent. And it really goes along with the kind of things that Rod Brooks will say about features on robots, is if you have a feature and it has not got something that is communicating what it does -- He was saying, "Don't put buttons on robots that are not useful, that don't do anything. Don't put eyes on a robot if those eyes don't tell you something." But then when he built the Rethink robot, he put a panel on it with eyes, because he wanted people to know that the robot was aware, saw what state it was in, so it was quite useful. So these communication cues have to be purposeful and they have to be followed up with action and they have to build trust.

And these interactions, especially object handoffs, are going to be such an important part of our world. I mean, it's so fundamental. You can't imagine robots being that useful if they're not even able to hand something to a human. And yet we live in a squishy world. Like if I'm handing you a jelly donut, it's a very different interaction from if I'm handing you a coffee mug. You're looking for different affordances and your senses and these things, some of them compress in different ways. And we have expectations, which is notoriously hard to deal with. Robotics when it was in my classes, was all about degrees of freedom and joints and propagating transformations because you could deal with that. You didn't know what to do with something that could deform as it did this. That would have been a horrible sort of calculation. But that's

the world we live in and the world that these robots are going to have to deal with as well. How do we teach robots about squishy things?

Yeah. So that's pretty interesting stuff. I think that's where robots have to be able to explore the world and gather information and force sensing and tactile sensing is really important for a robot to be able to push on stuff and measure the stiffness. But the other piece of that is they also have to understand the intrinsic value and the purpose of the object. So it's okay if I pick up a Nerf ball and kind of squish it around because it's a Nerf ball, but if I go and pick up a Faberge egg and kind of test it out, that might be a problem. So it's not only understanding the physical nature and being able to have force sensors and touch sensors and actually soft skin, if you will, on the robot's side, but it's also understanding the meaning and value of objects in space. And even simple things like, what is the way of holding this object? If I hold my glass upright, that's great, but if I turn it upside down, I've defeated the purpose of the glass, and now all the water is on my foot. So part of what we have to actually help robots understand is what are things in the world, how do we learn about what their meanings are? And if you get back to the handover, in what ways, when we hand something over, should we hand it over to somebody else? Should I hand the glass to a person upright, or should I hand it to them upside down? And that will depend somewhat on whether the glass is full or empty, but it might also depend on where the handle is. And these are all sorts of questions. The world is big and complicated and there's a lot of things that a robot is going to have to learn about the world. Now, some of this we can do by sort of scraping the web and gathering pictures of objects, gathering pictures of people holding objects, going onto Amazon and looking things up and deciding what the price of that object is. There's things that robots can learn and probably will learn from what we've got out there already in the database that is the internet. And we are going to have to, as individual robot owners, with my home cleaning robot that's looking after my house, I'm going to have to be able to explain and teach that robot what's important to me. Again, these edge cases, just because the world thinks this is right - and maybe that works for 95% of it - there's still going to be this need for me as an individual to be able to tune and explain and communicate with my robot what my values are.

And this is where we come full circle to, or we meet up with artificial general intelligence. And a lot of AGI researchers think that we won't get that except through embodied learning - robots that learn about the world and feed that into artificial general intelligence to say, "Here's how the world works. This is what you need to know in order to have useful conversations with humans." And it goes in the other direction as well, that those robots can't interact with that world or even understand it without some kind of general intelligence, or *something* that helps them make sense of that world. Is it a Catch-22, or what's the path to that resolution?

Well, I think, really, robots are robots and humans are wonderful, complex, complicated beings that are creative and amazing. And robots are our creations. They are just what we built and they're imperfect. And they will continue to be imperfect, but they are getting much better. But I think we get back to this, how do we make robots better? We make robots better by teaching them. I think when you think about how do we approach general artificial intelligence, how do we learn? We learn some things by exploring the world and touching and feeling and all of these things. But we also learn from teachers. We also learn from our parents. We also learn from

people who tell us things that we didn't intuit on our own. And those things also build into our knowledge. And I think we can think about using that same expectation for robots, right? There are things, the 95%, that it will be able to pull out through reinforcement learning and machine learning and other tools that are amazing now. And the databases, the internet as a database will provide that. The interactions with the world and failing in the world, or even trying things out in sim worlds and learning in those sim worlds, that will be part of it. But fundamentally, we are the teachers of our robots. And so we need to work out how it's going to be effective for us to make our robots more effective. Because the thing that will determine whether there is a robot in every home, as I used to say or used to be one of the things I used to talk about, is whether that robot is useful to me, whether that robot provides benefit to me, whether that robot can learn from me and not be frustrating. And so it gets back to, are robots teachable by your average person?

I get this image of you one day being a kindergarten teacher in a classroom full of baby robots. Wouldn't that be something?

Yeah, probably not. But it's a fun image, isn't it?

Yeah. What do you think of the current crop of media-friendly robots like Boston Dynamics' Atlas, and the Optimus that Tesla is working on?

Yeah, I love watching Spot dance. Every time there's a Spot dance video, I'm there because it's pretty fantastic. But it's basically a kinematic copy. So that's pretty interesting. I am blown away by the capability that the engineers have in building these fantastic machines. It's so impressive. And wow, I think it shows us that the technology on the mechanical side, on the power side, on the dynamic side, the physics side: very, very impressive. Where is the gap? And I will say this gap is still when we come to interacting with people.

And I think that I'd like to see more convergence of that kind of work and what you're doing and what, say, Cynthia Breazeal, was doing with making expressive robots. Well; let's look ahead 10 years: ideal future. Researchers always answer that with, "Yeah, I'd like to have sustained funding." Let's assume that that's happened. Ten years from now, what do you think robots will look like in our world? What would you like that to be? What's feasible?

Yeah. There's a lot of things to unpack there because it does seem like we always get asked this question. The technology will evolve. The batteries will get better, the capabilities, the dynamics, the motors will get better. We will improve our ability to sense human behavior and communicate; those things will get better. The piece that I also think is really super critical that will unlock a lot is that we engage the public in the questions around what do we want these autonomous devices that have capabilities and will be occupying our society-- And we need to engage in that conversation now. Whether that's conversation about autonomous weapons systems, which is really an important question that we have to talk about, or questions about if I send my personal robot out to collect my lunch and it bumps into somebody else and they spill their coffee, who's at fault? What are the rules of engagement of robotics in society? Because a lot of the things that are stopping us even right now from having autonomous vehicles on roads is not the technology. It's actually the law and the insurance, what we are willing to accept.

Because autonomous driving has got to the point where it's statistically safer than people. But in a culture of individualism, we are not willing to accept those statistics. We want perfect. Because I, brilliant driver that I am, would've not made the mistake that the robot made. So we have to have a conversation about what we are thinking about, and as a collective, as well as individuals, what our expectations are for living with robots. And I think we really do need to get our governments, our industry associations, our engineering associations really engaged, and our community groups really engaged, in both understanding what is coming and what is here already and what that means for how we live. Because you see it, you see technology getting adopted, and suddenly it's here, and then we have to decide, well, how do we feel about our kids having cell phones in schools? Or how do we feel about facial recognition technology?

Or having ChatGPT in the schools? And that's happening, yes.

Exactly.

I get asked a lot about ethics in this field. We've gone the whole interview without mentioning Asimov's Three Laws of Robotics. I think we and the audience know what the limitations of that are. Personally, I think of ethics as being defensive programming for humans. Are there sort of ethical ground rules for robots? Anything like "first do no harm" kind of directives that you put in there to avoid accidents? I mean, even obviously in industrial safety where there are robots that can kill people, and in some cases have, they have some pretty strict rules around that sort of thing. Are we evolving the kind of basic safety that we need for these as we get more of them in our world?

This is what I'm alluding to when I talk about public conversations. Certainly, the IEEE has a special group working on ethics and robot ethics is a live research topic. And there are conferences like We Robot that bring lawyers and roboticists together to have these conversations. So these are live conversations, but it's not uniform internationally. Because the rules around robots are different in Europe and North America and in Asia; they are just different. So this is very loose, but what I would say is, in Europe, if something goes wrong with the robot, then the manufacturer is on the hook. In the US, Canada, it's more likely that the owner is on the hook. And in other countries, and I'm not going to name names or anything, it would be more likely that, "Oh, you shouldn't have put your finger in where the robot was," kind of buyer beware, you know. You're on your own watch there. So it's different, and that's societal values. That's how you decide in your society and then that turns into regulations and laws. And so, again, it is about our society, our laws, our rules, our policies getting ahead or trying to catch up actually with the technology that's already out there. And this will be really interesting. I just saw something in the news about in the EU, the programmer actually could be on the hook if their AI technology caused harm. And so, yeah, I think this actually goes back to, "Well, now that I know that, and I'm a programmer or I'm educating programmers, we're going to have to become very, very serious about how we educate people who are working in these fields to think about what are the repercussions of our technology."

I wonder whether as we send more robots out into the world that are bigger, heavier, and more chances of them interacting with people, there will be some sort of call for some kind of

minimum testing or programming that would get something like a UL or a BSI stamp on it that says, “This thing has been programmed to tell when there are humans around and not run into them,” or things like that.

So there are already robot standards that are being developed that I think are the rudimentary pieces of this. And they’ve been around - certainly for industrial robots, it’s been along for a while, but now we’re talking about robots that are in community, social robots, personal robots. And it’s been focused on safety and so that’s been the primary thing about when a human is in vicinity, what is the behavior of the robot versus when a human is not in vicinity. And so those are developing; but we’re going to get into much more complex situations very soon. Delivery robots moving down the streets in San Francisco, delivering stuff.

Yes. Amazing. Wow, it’s been a terrific interview. If people want to find out more about your work, what you’re doing, what you’re going to be doing, or even see if they can come work with you, where should they go?

So I’m new at UVic. I do have a UVic website that’s got my name and email and actually, all my papers are on Google Scholar so that’s the easiest way to find out what I’m up to because we publish everything and it’s all there. So yeah, have a look.

Okay, thank you. Well, so one last question. How would you like robots to be personally involved in your life in the future? What’s your dream for what they could do for you?

Oh my goodness. I think that if robots could make it possible, not only for me but for people to have longer independence to live well and have better, more enjoyable lives, whether that means that older people could live comfortably in their homes for longer and feel safe, I think that would be one of the really good outcomes. And we have an aging population, this is a big deal, and I really do think that there’s things that we could do just to make people’s lives better. So that would be really good to see, if some of the stuff that I’m working on does turn to make whatever age, making people’s lives better, that would be a good thing.

Fantastic. Well, thank you Elizabeth Croft for coming on the show.

My pleasure! And thank you for a wonderful interview. Thanks, Peter.

That’s the end of the interview. This was one of those times I think about airing video because Elizabeth was so animated about this discussion that the interview was exciting on a visual level as well as auditory. By the way, to explain a couple of abbreviations I used towards the end there: UL and BSI refer to the Underwriters Lab and British Standards Institute, which are US and British safety standards organizations respectively which test consumer products. And their stamp on something basically means it’s certified as safe and so you’ll have, for instance, electricians who won’t install anything without that stamp. In Canada it’s CSA, the Canadian Standards Association.

In today’s news ripped from the headlines about AI, Microsoft has an AI that can understand images well enough to solve visual puzzles. Their Kosmos-1 is a multimodal model that can pass visual IQ tests and understand natural language instructions. The researchers believe multimodal AI is an important step towards building artificial general intelligence. Their paper is titled “Language Is Not All You Need:

Aligning Perception with Language Models." Clearly a dig at the paper that introduced the large language transformer models like GPT, which was called "Attention is All You Need." They have examples where, for instance, someone is holding a card on which they've drawn a cartoon smile in front of a cat so that the smile is in front of the cat's mouth, and the question is, "Explain why this photo is funny?" and the answer it gives is "The cat is wearing a mask that gives the cat a smile." In another picture a little boy is sitting on the ground crying next to a scooter, and the question is, "Why did the little boy cry?" and the answer it gives is, "Because his scooter broke." In yet another there's an image of an analog clock and the question is "The time now is" and it gives the right answer. It also passed one of those pesky questions from the Raven IQ test that show a matrix of icons and ask you which one would go in the empty square.

This project came out of Microsoft without any involvement from Open AI. They call it a "multimodal large language model," or "MLLM," which may become a new popular acronym. I'll bet that we're talking about MLLMs a lot by the end of the year. Obviously it has a lot of GPT-like large language modeling in it to understand and answer the text questions. And it has been linked to the expected-to-be-imminent introduction of GPT-4, widely expected to be multimodal.

Next week, my guest will be Tigran Petrosyan, co-founder and CEO of SuperAnnotate, an AI-driven data annotation platform for data scientists and machine learning teams. 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