

AI and You

Transcript

Guest: Paolo Pirjanian

Episode 13

First Aired: Monday, September 14, 2020

Hi everyone, welcome to episode 13. Today we will complete the interview with Paolo Pirjanian, an expert in robots. He first got into that at NASA's Jet Propulsion Laboratory, then became the CEO of Evolution Robotics, which made robots that could navigate around homes, and clean the floors, and unsurprisingly, the company ended up being bought by iRobot, where he helped develop the essential technology of the Roomba. But he is now the CEO and founder of Embodied, Inc., which makes a robot called Moxie. Moxie, who you can see in a video on Embodied's web site, is designed to help children who are challenged in their social development come out of their shell and feel safer in having interactions with people around them.

When the child's parents present the child with Moxie, and this little robot wakes up, he says, "Hi, I'm Moxie. I need help understanding humans. Can you help me with that?" And this puts the child in the position of being the teacher, and you know, we learn best that which we teach. If you want to get really good at something start teaching it.

In the last episode we talked about Paolo's experiences that led him to get into robotics, and the technology of a social robot, what he calls their SocialX platform. It's a framework for making social technology like empathetic robots on.

The fact that that kind of technology, that I can even say the words "SocialX platform," "Framework for making social technology," is indicative of how much progress we are making in creating and making easier to create these kinds of robots and interactive devices. Every one of these advances makes it easier to create the next one that goes further than that.

In this episode, by the way, we mentioned something called ELIZA. ELIZA is a chatbot that was developed in the 60s by Joseph Weizenbaum as a computer experiment. Now, that chatbot had different personalities that were in the form of a subprogram that you could put into it that would shape the kinds of conversations that it would have. The one that it's most commonly known for is the DOCTOR script. Not Doctor Who, but doctor as in psychotherapist in this case, and this is what we usually mean when we talk about ELIZA, because any other personalities that it was ever given have been lost to history. But the DOCTOR one was such a winner that ELIZA is synonymous with it. And that caused this chatbot to assume the personality of a Rogerian psychotherapist. It would ask you how you were doing, it would encourage you to say more, it would refer back to something you had said earlier at some point in the conversation and say, "Earlier you mentioned that you were having trouble talking to your girlfriend, tell me more about that." Now, it was originally designed by Weizenbaum as an experiment to prove that much human communication was superficial and that this would be mistaken for some of those interactions in some way. That was achieved so successfully that he became rather horrified at how much people were investing emotionally in talking to this tool, which as you can imagine, being in the 1960s, it was very, very, very simple and so he discouraged further communication

with it. Of course, it's a very simple program, and it was spread around so much in later years and was so easy to implement, that a version of it was embedded in every instance of the Emacs text editor that has been around since probably the '80s. But it was remarkable in its time for just how realistic it was, because it managed to create a narrow context in which its expressions of empathy, as simple as they were, hit enough of a target that people would open up to it. And in a sense, that's what Moxie is doing; it's creating a context where empathy is established and then can be used to advance the interests of the child through carefully chosen interactions. Plus a lot of clever AI. We will find out a lot more about how Moxie works and what Moxie does in this episode. So without further ado. let's return to the interview with Paolo Pirjanian.

I'm thinking of history here, and I'm irresistibly reminded of Eliza, we'll put some context of that around this in the pre preamble to the show. That goes back to what the '70s? And the conversation with a very primitive chatbot that nevertheless was enough to have some people claiming therapeutic benefit to it because this was a chatbot that emulated a Vergarian psychotherapist and was smart enough to listen to what people said, raise what they said later on and ask questions, "Tell me more about what you're thinking." There's a story that Weizenbaum's secretary asked him to leave the room while she had a private conversation with it and that was the point at which he decided this thing has gone too far. And then, later on, we had Cynthia Breazeal at MIT develop Kismet, which I refer to as the Mrs. Potato Head of robotics, a very expressive robot that people actually got upset when it was turned off. So this seems like an outgrowth of those technologies in many ways. What other heritage is there in Moxie?

A lot, right? We are basically leveraging everything we know to date about computer science and computer vision and natural language understanding. So Eliza and Cynthia Breazeal's robots from MIT and so on were the very first attempts of doing this and we learned from those. As you mentioned, we learned empathy develops very quickly. Cynthia's robot was not even verbal because at the time having verbal interaction was terrible. The voice recognition and NLP were just at very infancy, so they didn't really work. So she avoided that by having a nonverbal robot that would make some sounds and so on, but not speak English. We saw that in Furby by the way, a very successful toy that sold more than 100 million units, building empathy and it was just making some noises. Sometimes it would learn to say a few words of English so you can emulate or sort of create this perception of intelligence and empathy, but those are experiences that last a few minutes, right? So, what we are doing is going way beyond that, by many, many, many orders of magnitude and for that to succeed, we are deploying all of the latest in computer vision, machine learning, natural language understanding and more. And this would have not been possible even five years ago. The advancements we have witnessed in the last few years, which are primarily fueled by a combination of deep neural networks, availability of high-performance computing at lower costs and availability of large amounts of labeled data, vision or audio and so on, are making this possible today and we are leveraging basically anything and everything. That's one of the reasons why this was a very ambitious project because most startups can maybe succeed in one of those things that we are doing. We are doing 50 different technologies brought together integrated into one user experience.

And speaking of advances in the last few weeks, do you look at the sort of things that GPT-3 is doing and think, "Oh, I'd like a piece of that"?

Absolutely, we are and by the way, so you mentioned that with chatbots you can maybe sustain a few seconds of consistent conversation. With GPT-3, you can actually go about 10 volleys - 10 back and forths - in a topic and it will still be internally consistent in a conversation and then you will [then] change topic and go to another one. For instance, with Moxie, when the child has a conversation with Moxie, the volleys that we expect are about four or five, and then [the] robot moves on to the next one. So four or five, we can sustain a very consistent, coherent conversation about pretty much anything, and then we move on to the next topic. That's the next pole we are going to so let's move on to the next thing and the next thing and the next thing. And by the way, these areas are improving exponentially. Like from things that used to take decades to do, we're seeing happening in quarters now, like three months, so yeah, it's evolving very fast.

It is amazing. One of the things I predicted for the near future in my book was, I said the convergence of the ability to interpret facial expressions and demeanor and the fact that smartphones have cameras plus teenagers spend a lot of time looking at their cameras, means it would be possible to detect whether they were planning on harming themselves, which is a risk with that age range. Do you get into the age range where self-harm is an issue? And do you look for it or are you limited to younger ages where that's not so much of a risk?

No, I think it is a risk. We are targeting five to 10-year-olds, developmentally 10 years old, right. So you can have maybe someone that is 13, 14 but developmentally they are maybe 10. And you would see amongst teenagers-- And we know, by the way, suicide rates in the last few years have been breaking records based on CDC reports, which is sad. We have chosen not to address this because we don't have enough context to be successful. It would be too much responsibility for us, so we leave that to the parents. Again, obviously, we are a tool to help with parenting, we are not going to replace the need for parenting and that's not our goal either. There are companies like Ginger.io and other digital therapy companies that are doing that for adults and those companies are actually being very successful. They basically have a software that runs on your cell phone and is able to interact and also measure anything they can measure from yourself from both movements and activity and so on, and can intervene when they feel that the subject is at risk and will intervene and try to get them off the ledge if you will, for those kinds of situations.

Right and I should clarify that I wasn't talking about Moxie intervening in that case but bringing it to the attention of someone. I don't want anyone thinking that we were suggesting that robots should be the tool of choice for assisting suicidal people.

No, that's a good clarification. And so that's a good question too, whether we would notify someone. For now, we have taken the position that we are not going to look for those things and detect and notify and one reason is that if that robot is perceived by the child to be a spying device, then you lose that trust that you need in order for them to be able to feel comfortable in a safe, non-judgmental way interacting with Moxie.

Right. So we've talked about establishing empathy and how important that is in Moxie's relationship with the child. And looking at devices like Google Home and Alexa, those have been deliberately engineered to not be anthropomorphic in any way. They are as uncreature-like as anything could get. Now here you are going in the other direction and if you're successful at creating empathy, what happens when it's time for Moxie to leave? Is there separation anxiety?

Yes, there is, actually there is. So in very early testing, we observed that you cannot just yank the robot out of the child's hands. They do bond with Moxie very quickly. And by the way, this is not uncommon, even with stuffed animals, right? Children can create attachment to objects that way because it's soothing and they like the object and so on, and specifically, with something that's animated and lifelike, I think that is much deeper. So I think it's best to leave it up to the child to decide when they sort of are tired or have outgrown Moxie, to leave. And that's why we have made it affordable. It's not a device that costs \$50,000 and you rent it for six months and then you have to give it back. Many families can afford to buy it and use it at their home. If you can afford to buy a smartphone like an iPhone, you should be able to afford to buy this if your child needs it. Yeah, if parents feel the child is too attached and they want to remove Moxie from them, they need to wean them off. They can't just do it in one shot. So we did that in our experiments early days where we only had one prototype. Obviously, we couldn't let every child go away with a robot after one test. We would have testing over a week, and a couple of days before the test ended, Moxie would start preparing them that "I need to go back to the lab with your input. We're going to make improvements to me and all the data, and then it was an easy transition.

Got it. From somewhere, I heard the statistics often repeated that about nine-tenths of communication is nonverbal. And that's certainly borne out if I look at the transcripts of the shows, even the most accurate and faithful ones do not make nearly as much sense as this recording, and that doesn't even have the video on it. So that's a vocal variety and intonation and all kinds of other things that you can't capture in text, pulling a workload there. And it's augmented with things like we can see each other at the moment and that is pulling even more of that workload. Now a lot of AI is focused on, as you mentioned earlier, NLP, Natural Language Processing, which is turning what I'm saying into text and then saying what does that text mean? All of that is devoid of that nonverbal content. So, if nonverbal content is supposed to be nine-tenths of communication, how much weight do you give it?

Well, we give it a lot of weight. This is what the name of the company is based on, the notion of embodiment because embodiment is neuroscientifically shown to be super important. Now, nonverbal language, the numbers vary depending on what source you go to but let's say it's anywhere from 30% to 90% of the communication is nonverbal cues. As I said, the name of the company was based on this notion of embodiment, and the neuroscientific evidence shows. By the way, the way it works is the following from neuroscience. The embodiment - so the fact that I can change the intonation of my voice, the fact that I have body language and can use my gestures and show my passion and excitement, and even my facial expressions and so on -

resonates with mirror neurons you have, so there are that neurons that are imitating neurons so they trigger when there is something that you can imitate. And [those] imitation neurons, the mirror neurons can help elicit the emotion that I'm exuding in you. So if you will, in laymen terms, it's a wireless way for us to communicate. And what has been shown with fMRI scans is that when you have interaction with something that's physically embodied and co-present, you can trigger these neurons at a much deeper and broader level than if you're interacting with something that's less embodied, let's say a character on a screen. This is the same way with you and I, right? We could be talking on the phone right now and my intonation of voice would be some embodiment but not as good as when we can see each other, you see my body language and my facial affect and so on. And the next best thing would be that we're actually physically in the same room where you would feel my excitement and passion about what I'm working on even more. So that's really the theory behind why embodiment matters.

So these mirror neurons are acting as the receiver to your nonverbal transmitter and echoing that inside me so that I pick up that content.

Exactly.

It's easy to see how we represent text, 26 letters will do; is there a lexicon for nonverbal content?

There is to some extent. If you look at sentiment text, we can extract the sentiment of what are the kinds of words that are being used. From facial expressions, there is actually a very nice, rigorous theoretical methodology that is actually looking at every single muscle you have in your face that fires, and the combination of these represent different emotional states. And of course, you can have very high-level states, very abstract states of emotion and just categorize them to, let's say, five, or you can go into much more detail and look at every single muscle that's firing. And actually, Moxie extracts those, and then from that we analyze the ones that are firing and can come up with some nuances of emotion that the child is expressing. And even measuring because some of the things for instance, for some children, we want to teach them about being able to express themselves too because there are some children that have neurodevelopmental challenges such as autism that are not very good at expressing their emotions. And Moxie has activities that will encourage the child to do that and it will measure how well the child is doing that. So yeah, there is a lexicon for this. And for body language, I am guessing yes, but I'm actually not too sure about that. But we are looking a lot at voice and we're looking a lot at face in the first version of Moxie.

Have your test subjects, the children that you've worked with so far, have they done anything with Moxie that you've found unexpected?

All the time. Children are definitely unpredictable and we knew that going in but then that makes it even harder. For instance, we do all this face recognition and voice recognition and body language and expressiveness by Moxie, and then we find that children want to be lying on the floor talking to Moxie where they're no longer in the field of view, or they're sitting in the middle of the dark, or they want to put a blanket on Moxie so Moxie's not cold and they want to hug

Moxie and put it in their bed, which are all nightmares for a technologist. So we see a lot of those things. But also, on the positive side, I've seen this a lot. Originally, we were actually helping families with kids on the spectrum. These are typically children that have a much harder time opening up to another person, even their own peers and even sometimes their own parents. I remember one of the very early prototypes, which was we did not have any of the AI yet. We had created a button board we called it where we could trigger certain responses from pushing the button. And so this was actually the first-ever test. We went to this family's home, the son was I think at the time seven years old, on the spectrum, very high functioning super-genius kid, but socially, not very strong. We put Moxie in front of him, he asks a question, we push a button, Moxie responds. And then he asks another question, we push another button, Moxie responds. And after a third interaction, the kid just completely opens up and talks to Moxie as if it's a friend he has known for a while, to the point that mom got a tear in her eye. And then I asked her "Why are you so emotional about it?" She said, "You know when he comes home from school how long it takes me to 'find him'?" And I put it in quotes "find him" to get on a wavelength where he would open up to mom that way. Like she would work very hard to do that and she was blown away that this child was able to open up to watch that quickly and we see this very often.

Well, I'm getting emotional listening to that because I'm seeing and hearing and deciphering the value that Moxie can bring here to these children. I'm vicariously experiencing the same with him. There are those mirror neurons firing from what you were saying there.

Yeah, no, I agree. This is as you pointed out, as one of your first questions, this is a completely super complex project. The reason that we are here doing it is not because it's easy, but it's because we are so passionate about helping these families, right? And that's what keeps us going against the odds. We are doing it and we have a team, every single one is so connected to this mission and working around the clock to get it done.

Let's talk about 10 years from now, you were talking about that earlier, and that vision sounded like a robot that could walk around, help people in the house or anywhere else and lift things. I was thinking of Boston Dynamics Atlas robot, which looks like it was designed by James Cameron. What would your version look like?

Yeah, a legged locomotion is of course a lot more flexible so you can deal with going up and down stairs and those kinds of things. And if that becomes affordable and safe to use in an in-home environment or an elderly care facility then that may be a better choice than wheels. So I'm agnostic to what mobility technology, but the functionality that we want to support is the ability to do assistive care, helping people that need some help. And there are many simple tests in a home that can be done. If you have the companionship already, with the technology we already have with SocialX on Moxie that can allow them having conversations as I mentioned, keeping them active and motivated and-- By the way, you know social isolation leads to lack of motivation, which leads to death? Social isolation is the largest cause to deterioration of health leading to death. So that is why we are focusing on that first, but then if you're in the home and can do other tasks to help the elderly or other vulnerable parts of our society, grab their arm and help them get out of the bed. Maybe support them to lean on you to walk towards the bathroom so they don't fall if they need that help. So tasks like that are just priceless for allowing someone

in need to be able to independently live in their homes. We are all headed in the same direction age-wise, right? At some point, we will get to a point where we will not be able to perform all of the tasks independently. And I can just imagine you can also relate to family members, your parents or grandparents that you have known in their healthy adult life and now they're getting to a point where they're becoming frail and they need help. And how much of a challenge is it for someone's dignity to need a stranger come help them in an intimate moment like going to the bathroom, right? So I want to be able to provide care with dignity and compassion, that is the goal and the technologies are there. Maybe if Boston Dynamics legged locomotion is not available at price points we can put in these robots at the time, we use wheels. It will be constrained but-- By the way, funny enough, when we get older, we try to make our environment safer from tripping hazards, and not too many stairs and all that. If there are stairs, usually you put the ramp on it and all that. So those environments that are safe for [the] elderly are also really good for robots. So it's a good match between those two.

And what you're describing there about assistive care for the elderly sounds exactly like the Japanese goal because of their demographics, meaning that they don't have enough younger people to look after all the elderly people. They've apparently been working on getting robot technology for assistive elder care for some time now. Where do you stand in respect to their efforts? And what sort of synergy is there?

Yeah, actually, in my previous company Evolution Robotics, I used to spend a lot of time in Japan, and Japan has been doing research and development in this area for at least three decades. And pretty much any company in Japan that you touch, that anyone would recognize the name of, they have a robotics division and have had it for the last 30 years. So they're not joking about this but these are challenging problems, right? And, yeah, there are going to be companies that may get to that goal earlier than others and maybe there are synergies that we focus on, the empathy and companionship while they focus on the physical side so I don't have to build that as well, so we combine the two. So Moxie becomes the face of those robots that allows them to have rapport and something pleasant to have in the home rather than an industrial machine in your home, right? No one would want to have an industrial-grade machine in their home. So it gives that softer touch to their solution and so I can see synergy there.

Thank you. We're getting near the end of our time here. Ten years out from now, what's one thing you'd like to be true in your vision of the world or what you've created in it at that point?

I think it's actually pretty simple. I'm focused on our first target, which is bringing Moxie to everyone's home to help the children. So if I'm successful in that goal, and having Moxie for every child that needs it, I would die with a big smile on my face. There are many steps to get there though. First, as we are launching this now longer term, I do want to make Moxie even more affordable although, as I mentioned, this is an order of magnitude or tours of magnitude cheaper than robots of similar capability. So it's affordable for most people to actually be able to buy it. But I know also there are families that cannot afford to spend \$1,500 on something like this and ideally it gets to a point where it's much cheaper than that, whether it's by subsidization through [an] insurance company or whatever, those are the steps I want to go through as well.

Thank you very much, Paolo. Where should people go to find out more about what you're doing or contribute or ask questions?

So embodied.com would be a good website or moxierobot.com, both of them go to the same place, and actually, you will see you can read more about Moxie. You can reach out to us to connect with us. I would also encourage you to if you're interested in look at some of the live demos I've done because the promotional video of course is professionally shot and edited and so on, this will be raw video. And you can even sign up to seeing a live demo of Moxie, we do that every other week.

Well, thank you. One more question just came into my mind because you said how you got into this robotics business and how taken you were with that. So if there's someone in say, high school, college, same kind of position now thinking, "Wow, I want to do more about this robotics thing." Where would you suggest that they start?

In high school, a lot of high schools now have robotic clubs. So there are robotic competitions either it's the first competition called First or the league LEGO League competitions. I would highly recommend you get into one of these robotic clubs. If not, if you don't have access to that, you can potentially find robotic clubs in your local community that you live in. There are robotic clubs in most areas now. If not, the other thing I would recommend is to buy a Lego Mindstorm set. It's a Lego set but it's a robotic kit. You can build and program your own robot, which will be a great introduction to robotics.

Wow. Fantastic. Paolo Pirjanian, thank you very much.

It's my pleasure, Peter. Thank you for having me.

What a lot to think about! I hope your mirror neurons are also firing as a result. We're looking here at the application of AI and robotics to make a real difference in that most human of needs and fields: child psychology. Not by threatening or subsuming any child psychologists' jobs, but by finding a unique opportunity, a unique opening, that they couldn't exploit themselves. Maybe before long Moxie's successors will be recognized as tools in the toolboxes of those psychologists.

Even more provocatively, Moxie teases the possibility that one day our AIs will be not just sapient but sentient. Which is to say, not just thinking, but feeling. Not just intelligent, but sensate, and emotional. If we're going to share the planet with beings as smart or smarter than us, shouldn't they have feelings? A controversial issue, bound to inflame feelings and opinions among many, but one we shouldn't and indeed can't back away from, because it's one of many fundamental philosophical questions barreling down the pike towards us.

So perhaps it's fortuitous that we will have a philosopher with us on next week's episode. We will be joined by Dr. Karina Vold, from the University of Toronto and the Leverhulme Centre for the Future of Intelligence, which is associated with the Centre for the Study of Existential Risk in Cambridge, England. That's one of several institutes that have sprung up around in the world in the last ten years or so, addressing the question of the survival of the human race: It's right there in their title. And the Centre for the Future of Intelligence is focused on exactly where intelligence and AI is going.

Karina has spoken about the ethics of AI at Cog-X and on the BBC, she has many papers and a chapter on existential risk of AI in a forthcoming Oxford handbook on digital ethics. That's on next week's episode of 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>