# AI and You

Transcript

## Guest: Dr. Ryan D’Arcy

## Episode 7

***(Part 2)***

First Aired: Monday, August 3, 2020

Hello everyone, welcome to part 2 of the interview with Dr. Ryan D’Arcy. I said in the last episode that we met when we were both speaking at TEDxBearCreekPark in British Columbia. Ryan is a neuroscientist, the President and Chief Scientific Officer of HealthTech Connex. If you go to their web site, you’ll see the headline “How’s Your Brain Doing Today?” which sums up Ryan’s mission: How IS your brain doing today? How would you know? How would you know whether it’s doing better than yesterday, or as well as it should? He invented the term brain vital signs to come up with a way of measuring the health of the brain the way we measure someone’s blood pressure or temperature or pulse.

Ryan is also a tenured professor at Simon Fraser University in Computer Science and Engineering Science, and at the University of British Columbia in the Djavad Mowafaghian Centre for Brain Health. If you look at the show notes you’ll see a picture of him next to a magnetic resonance imaging machine. He routinely looks inside brains; sometimes with the lid taken off… but he prefers to do it without opening people’s heads, hence the MRI. In fact he’s got some definite opinions about how we should be doing brain-machine interfaces without the head-opening step, and you’ll hear more about that in this episode.

I really like talking with people who are passionate about what they do, and Ryan is incredibly passionate about brains. He is working for a revolution in how we understand our brains and how we fix them when they’re broken.

I think you’ll see some stunning advances from Ryan and others in that field in the next few years. You see, the pandemic crisis that we’re in is focusing a massive proportion of the world’s scientific firepower on a single problem right now, and that’s going to accelerate our capabilities in medicine much faster than we’re used to. That kind of attention won’t just advance virology, or microbiology; it’ll spill out into neighboring fields. Mostly medicine, but with all the computing resources being used in that effort there will be a fusion that overlaps into computer science. That’ll put us right where Ryan is: The cutting edge of neuroscience, figuring out how this lump of jelly on top of our necks does so much and how to keep it at peak performance. Maybe we’ll figure out how to use computers as brain prosthetics to assist us. The way we use our smartphones now, only with direct brain connections so that we won’t be limited by how fast we can type or listen or speak, but only by how fast we can think. Who couldn’t use that? I for one could use some help in the face recognition department: it’s very hard for me to remember both the name and the face of someone new. Oftentimes I come away from an encounter realizing that I’ve remember their face but not their name, or vice-versa. And of course, I’m doing what I can to improve that with what I’ve got now, but – wow - wouldn’t it make a difference if I had some help from a computer, which is very good at remembering things, that could fill in the gaps? See someone and it recognizes them and reminds me who they were and where we met? How about you? Could you use something like that? What else would you use it for?

A headline this week was that researchers at the Moscow Institute of Physics and Technology were able to use an EEG to monitor the brainwaves of someone who was watching a video and from the brainwaves alone, use an AI to reconstruct – blurry - but recognizable frames of that video. That's huge. We can develop that technique to help people who have a stroke or are paralyzed to control devices like exoskeletons or wheelchairs.

Yeah, you can go in the other direction with this too and think about the negative implications of being able to tell what someone is thinking. That's one of the disruptive effects of AI that we'll need to grapple with in the years to come. It's not too soon to start thinking about how to deal with those.

But to get better brain machine interfaces we have to better understand the human brain. To learn more about that, here's part 2 of the interview with Ryan D'Arcy.

I've long been struck by how our concept of the mind has always been whatever our most advanced technology was at the time. So it goes back to, like engines and the Victorian era, and then telephone exchanges, telegraphs and now computers and neural networks and computers. It strikes me we have a few more generations of models to go. Our computers are architected along what's called a von Neumann architecture, where you've got a central processing unit that knows how to execute instructions, and then you've got memory where it finds the instructions and operates on data. And then you've got a place where you input questions and it outputs answers. The human brain doesn't look like that; it seems to look homogenous. I'm sure that's a naive interpretation. What sort of structure is there of the brain, and what sort of models does it suggest of its cognitive ability?

Yeah, it's a great question. I think the real global realization of this came around the time when IBM created the computer Watson and beat the Jeopardy champions. And at that time, most card-carrying neuroscientists like myself got entered into debates with computer scientists about the concept of, will computers outperform the human brain? And no one, really, I think, knows the answer. Time will tell as to what that that will unfold. But it is interesting to do a bit of a compare and contrast. The first thing to recognize is that I think we almost always get the comparison the wrong way around. So a lot of people try to map the brain with respect to the computer. But in AI with neural networking, the advances in deep neural network learning really did get led by modeling a computer after the brain. And so, like the brain, computers have central processors, and then they have distributed systems in which they transact information, and that's actually similar to how our neurons are structured. We have areas and regions that do some central processing. And then we have distributed networks in which that information can be not only sent but represented. And so in many senses, a neural Network and AI neural network are starting to converge. And the AI neural network is starting to learn from what evolution has put in place between our ears. And so I think that's really exciting, and it's definitely greatly impacting our world on a daily basis. Whether or not you see it in terms of autonomous driving or the ways that we can actually analyze the complex signals and information that comes from the brain and make sense of it, we're seeing those benefits in AI for sure.

Talking about how to make sense of those signals, I read something here that I copied an article that says Joseph Makin led a group from the University of California, San Francisco, to render a person's neural signals as English text while the person read a sentence aloud. Sometimes it produced gibberish, for instance, it translated brainwaves representing the woman is holding a broom into "the little is giggling, giggling." But much of its output was very close to the spoken words. "The ladder was used to rescue the cat and the man" came out as "Which ladder will be used to rescue the cat and the man?" So here's something that's being used to try and turn our thoughts into rendered language, and it really prompts the question for me of, what is a thought?

Well, thanks for asking the easy questions.

That's why they pay you the big bucks.

I think, first of all, that's an excellent example we called brain-computer interfaces or brain-machine interfaces or BCI or BMI but more BCI, and certainly anyone who watches things like Tesla and SpaceX knows about Neuralink with people like Elon Musk, and those are attempts that are very much focused on integrating our computers and our AI with our brain activity, and that's a very good demonstration of the potential. There's gonna be a whole lot between the potential that you can see in the laboratory and practical interfaces. And it's also going to not - sometimes it will be a bit of apples and oranges comparison, because computers are fundamentally different from our brains in multiple ways. One way is that our brains are, interestingly, this paradox between localized and distributed at the very same time. So in that experiment, to read that slice of a thought, they were going into a specific area of the brain. But that specific area did not allow them to see the entire distributed neural network, to integrate that particular sentiment with other thoughts. So I think that demonstrates that we've got a long way to go. They're really interesting steps forward, for sure, but we have quite a lot of catch up with Mother Nature to be able to understand how that then gets represented in a way that you and I would think of as a thought. And just to give you one recent science example from our group along those lines, when you look at brainwaves, and I mentioned earlier that it's almost the 100th anniversary of when brainwaves were first recorded in 1924, for the longest time, we would try to access the small brainwave signals and would move the blink activity. Every time you blink, you create electrical activity that was an artifact. And so what happens is that all the field was focused on this idea that these blinks were artifacts and bad, and we had to get to this small. It turns out that when you blink, in terms of a thought, it's like a hidden window into your conscious mind because there's brain activity and buried within that blank that relates to you environmentally scanning and becoming aware of your environment or re-registering your environment. So that's a thought, right? But it's not a thought that we ever thought, over 100 years' worth of research. So I think it's very difficult and very interesting to continue to look at what is a thought, what is consciousness, and we are still really in the infancy of us understanding that. And that's another reason why our understanding of our brains is just driving our AI and I think it will be interesting to see as AI climbs up if that helps us to better understand our brains. But right now, we all have to recognize that our understanding of our brains is so crude that we got a long way to go yet.

As you said, one of the harder questions. Lay people tend to think of thoughts as being these sentences bouncing around inside our brains and science fiction fans who read about telepathy think about that the same way as well. When we see uh, fictional shows about that happening, they represented it as words appearing in someone's head. But the reality must be much more complicated than that. For instance, I can tell when I'm speaking that the words I'm going to say are showing up in my brain ahead of my mouth; like I know this sentence 1/2 a second into it, and then my brain is just unbuffering everything else that I'm going to say. And sometimes the problem is that it gets three sentences ahead of me and then goes back and tries to edit the one that's on its way out of my mouth. So there's something going on there that is not necessarily connected to language or at least isn't happening at the same speed as vocalization. And do we have any idea of what those thoughts are? Can they be deconstructed without having to come out as vocalized words?

There's been some interesting research that's been able to analyze brain imaging data to be able to understand and pull out the equivalent of what was someone's thinking pattern. But again, if you think in the way of layers and we understand that we don't know how many layers of complexity are on the system, that was the very first layer, and a couple of straightforward examples would prove to you that from a thought experiment point of view that this is certainly a much deeper and complex and intricate system than that because, for example, almost I think all of us have experienced the concept of hearing music that pulls up a memory and that memory then pulls up a thought right? The same is true for smelling something that pulls up a memory and brings time. And then that brings a thought that creates something that you may or may not vocalize or do a volitional action over. So in those cases, you know right away that there are multiple ways by which our neural systems can bring thinking to bear, be it conscious, unconscious, automated or by control. And so it's really going to continue to be the case that with a system as complex as the human brain, we were not gonna be out of jobs quite yet in terms of computer AI telling us how it all works.

Going back to one of the first times that we met, you gave me some feedback on an early version of my talk where I started out by saying that the ability to perceive our motion through time was uniquely human and you corrected me and said, Actually, there are animals that can do that as well. So I took that out because it wasn't necessary for the talk. Now, some animals even have more neurons than the human brain, although they still don't have some of the capabilities that we have. But are there things about different animal brains that are optimized in ways that ours aren't, that provide research possibilities?

Oh, most certainly. The question is, what? What is the research possibility? There's been historically - the bulk of how we've understood the brain has been through studying non-human brains from the lowest level to the highest level. And it's absolutely the case that I think we are a little bit anthropocentric in our perspectives of the brain. I've often had to remind people that humans *are* animals. We sometimes like, categorize humans as separate to animals. But every brain is an information processor. But from a practical, functional purpose it's a supercomputer wired to do a specific job, which is to survive and evolve and follow the theories of natural selection in whatever environment and situation and context it finds. So the neural systems are always being optimized for that, and you would see, for example, just as one example, in non-human primates there, somato-sensory cortex in terms of their sensing and their touch is much greater and larger. Raccoons the same, so their brains have been rewired for abilities that allow them to survive and to thrive. And the same is true and we think that what's uniquely human is largely related to two things, actually. Number one, we have opposable thumbs, which has allowed us to engage our frontal lobes, our executive functioning to start to use tools to bring on the industrial age and the ages of technology and innovation. And number two - and they're not in rank order - we created communication systems that support the concept of language as we know it, which we haven't been able to necessarily discover yet that other animals have it in quite the same way. That's not to stay by any stretch of the imagination that there aren't other animals that have language systems that operate as *they* know it. And we just simply can't perceive them because of our perspective. But nonetheless, those have created, over evolution, a reorganization in terms of priorities and structure in our brains so that we can use the environment that way and that that is where we're all animals, fundamentally alike we're all built from neurons, fundamental neural circuits. And we've just specialized them for our unique needs.

And then there are the roles of feelings in people, and there's been great debate and controversy in scientific circles over what other animals experience feelings. But I do remember learning not that long ago, actually, after speaking so often about sentient life, thinking that that meant self-aware, conscious thinking that actually, no that's not what it means. Sentient means *feeling*. We are Homo SAPIENs, and it's *sapient* life, that is thinking, but *sentient* means that it's *feeling*. And there's no debate about how much human beings can feel and have emotions. What sort of role those play in our intelligence?

Well, I think there's an entire discipline of social neuroscience which has really been through, particularly through brain imaging advances, able to start to do some really interesting experiments in terms of us, the human equation. And it's, I'm sure that, watch your *Scientific American* brain articles and you'll see these - really, there's a host of them and there's so many it's difficult to give one example, but one example that I always that comes to mind for sure would be this concept of metascanning, that you can scan two brains at the same time and show that their activity, let's say between a mother and her baby or between people who have strong feelings for one another, that their brain waves actually synchronize uniquely. So we do know that there's a strong biological basis for these things. And what's interesting to me is this where this kind of pulls us full circle back to the whole concept of AI, computers and that sort of thing. Because if you look at how we came and evolved as a human species and then you make a comparison of how did we evolve computers, they're fundamentally different on a couple of key constructs and dimensions, and one would be certainly the concept that we evolved to have feelings and have sentience, whereas with computers necessarily that wasn't a practical engineering requirement. But the other is actually you can train. Let's play a simple experiment. We're going to train a computer. The most powerful analytic processor we have, the supercomputing and quantum computing added in and everything you can do that's on the planet, and you're gonna put it against the human brain and pattern recognition, and you're gonna show it, you're gonna train it a bunch of animals, right? So you're going to show pictures and pictures and pictures. And then you're gonna demonstrate that this your computer, your AI could perform better in terms of reaction time in processing new patterns of new animals. But what we forget to ask is the context. The human brain pattern recognizer, if some of those pictures were dangerous animals in the environment like a snake, the human brain will fundamentally analyze that situation differently than a computer. Right? So the computer was never built to have an evolutionary fear. So it was fundamentally wired and designed in a different way. So AI will, as close it gets will still operate on different operational parameters. Then you necessarily would to survive in the woods and not step on a poisonous snake. So I think we still have to reconcile a lot of those fundamental first principle issues when it comes to this, and certainly - to just bring it back - human feelings with respect to sentience is right up there as a first principle.

I'm continually struck by how much our cognitive systems are optimized for our survival, even to the extent that they then reveal that there are other things that are not so important to survival, that they get wrong. Optical illusions is what I'm thinking about. For instance, anything that's based on an abstract image can become an optical illusion because our brains have only got so much optical processing power. But it's all optimized for our survival in the real world. You mentioned Neuralink earlier, and a lot of our listeners are going to be familiar with Elon Musk's rationale for a starting a brain-computer interface company, which is that he foresees there being a time and future when AI evolves rapidly and would get bored communicating with us at the speed at which we talk or read. And being able to communicate with it at the speed at which we think could buy us some important time for being able to convince it of what's important to us. What are your thoughts about the current progress in Neuralink and other companies and, like, say, the Human Connectome project, of the current state of brain-computer interfaces? And where do you think that industry is going to go in the next 5 to 10 years?

Yeah, excellent topic, always love speaking about this. I look forward to the time I should I should ever be able to have a one on one chat with Elon Musk about this. This is the first time I've seen him get his first principles wrong. And so I think there will be some revisions along the way. And I'm not critical of what Neuralink is doing, but there are two first principles that I see that were not correct from my perspective. The first I’ve already mentioned, which is that the relationship of AI evolves under different operational parameters and circumstances than human brains. So assumptions that assume that AI will get bored of human brains may need to be rechecked on that front. But the other area of first principles is that if AI is evolving at the rate it does, then it will likely be - Neuralink took an invasive approach where they actually open your skull and get the sensors in to a specific area of your brain. And the reality that would be that the better bet is that a non-invasive BCI will benefit fast and quickly in terms of signal processing advances and AI. And so that, of course, will have the immediate advantage of being able to look at both localized and distributed activity in your brain. And so I would definitely put the bets on those plays. And there are some of those plays out in the world that are moving forward and have similarly sort of focused on that priority. In terms of where I think we will go, I think that the grand idea that computers and human brains can communicate more efficiently is an awesome one, and particularly it will be demonstrated in the practical benefits. Primarily, I would imagine things like performance-based human machine interfacing when you're flying as pilots or driving cars, Or astronauts. I think that that's historically been where the practical applications have started first, and it likely will continue to be there. Also, robotic surgeries as an example and assisted medical technologies that interface with human, I think in practical demonstrations like that is where we'll see it first, well we already have, and we'll see it continue to expand.

So if we turn AI around to IA we have intelligence augmentation. I, for one, would like to have a brain computer interface to Google so I can think a query and get the answer a lot faster without needing a device in my pocket. How far are we away from having something like that?

I would say we've seen lots of advances, and we're learning to communicate and rather than to recreate an imperfect interface, use interfaces that have been around for a long time and jump on those. So I think, uh, the smartphone is a perfect example. I can't spell in the spelling contests anymore, but I can spell with my fingers and so I rewired my brain as an interface to be able to communicate more efficiently that way. Same with voice recognition and speech to text and text to speech. I think advances have used AI largely and are creating why you have your Alexas and Google homes. So I think you'll probably still see that they'll be coming in on established channels into our brains rather than necessarily you'll be plugging yourself into something and just seamlessly be able to find something faster on the Internet. If I was a betting person, I wouldn't say that's not gonna happen, But the way that it happens will not look like what we think at the beginning, it will be a very different implementation.

Something like the next iteration of Google glasses that that we speak to and which projects images onto our retina, perhaps?

I kind of think this is where AI is gonna surprise us all by pulling in a whole lot of data. So whether or not it's Google glasses, it's looking at eye blinks, its recording brainwaves, basically, we're getting to a place where we can take all that complex information and throw it into AI and machine learning and extract value out of it. So I think it'll probably be a multimodal sort of phenomenon.

Well, this has been fascinating. I know as I've told you before, that I could go on with this conversation for hours and maybe we'll have the chance to do that again another time. But our listeners, and human beings, have only so much attention span. What's obvious to me from the moment we met is the passion and drive that you have for making a difference here. What would you like your mark on the world to be? How would you want others to describe your legacy when all is said and done?

That's simple. Our goal is to positively impact a billion or more brains, and to ensure that we can unleash the brain potential of humanity to solve our scariest problems. So that's what we're trying to do.

Well, that is a great way to sum that up. If people want to learn more about what you do or have something that they need to ask you or let you know how should they get in touch or find out more?

Oh, for sure, reach out to me. There's a number of ways you can you can reach out to me over the Internet, and I'm always happy to chat with people who are interested in what we're up to. And just like our connection, I think I'll end with this. We could calculate that the potential for brainpower should not just exist within our skulls. But it should actually be that you modulated my neural circuits and I've modulated yours. So for anyone who's interested, reach out and will rewire each other's brains.

Well, we've done a lot of rewiring in this session. I'd like to thank you so much for coming on this show.

That's the end of the interview, sorry for the rather abrupt ending there. I like how as much as we were talking about what could happen in the far future or certainly not next year we kept it grounded in also things that are going on right now and one of the things that Ryan is doing is working to get treatments and measurements of the brain into the hands of the doctors that you go and see for a checkup or for something more serious.

It's long been a frustration of mine that I keep reading about medical advances, new discoveries, new inventions, new treatments and they don't seem to show up. I go into my doctor's office and they're still using the same equipment as 50 years ago for the most part. Maybe they're taking my temperature with something in my ear instead of under my tongue but its still the same sphygmomanometer that measures blood pressure; they're still using a stethoscope; this is not my idea of a kind of progress that we should be making. As I said earlier, I think that this crisis will precipitate a lot of advances in medicine and so I'm hopeful that we will see some of that show up in the doctor's office by the end of it.

Ryan particularly feels that kind of frustration acutely with respect to the health and measurements of brains, and he's working on how that technology could be pushed out, again to your doctor's office. So maybe not far in the future you'll go in to your doctor and he'll scan your brain while you're there. And you can have Ryan to thank for that.

By the way, in the interview, when we were talking about animals that have more neurons than humans, I looked that up: The elephant has about three times as many neurons in its brain. According to Wikipedia, if you count a different way, which is the number of neurons or equivalents in the whole nervous system then some whales beat us.

OK, a headline item before we leave. I'd like to draw your attention to a recent publication with this title: "Using Sinusoidally-Modulated Noise as a Surrogate for Slow-Wave Sleep to Accomplish Stable Unsupervised Dictionary Learning in a Spike-Based Sparse Coding Model." Wow. Allow me to translate. This paper, that was presented at the CVPR Women in Computer Vision Workshop in Seattle in June was about optimizing the performance of so-called neuromorphic processors: chips that have been designed from scratch to emulate the human brain structure. So we're still on topic for this episode.

The really interesting thing about this is that they said that those chips need the equivalent of sleep in order to stay healthy. Now, this paper was published by, and the work done at, the Los Alamos National Laboratory, and here's a quote from their computer scientist Yijing Watkins:

"We were fascinated by the prospect of training a neuromorphic processor in a manner analogous to how humans and other biological systems learn from their environment during childhood development." She and her research team found that the network simulations became unstable after continuous periods of unsupervised learning. When they exposed the networks to states that are analogous to the waves that living brains experience during sleep, stability was restored.

"It was as though we were giving the neural networks the equivalent of a good night's rest," according to Watkins. Does this mean that there's more convergence between artificial intelligence and human neural structures coming up? I think that's a good bet.

In next week's episode, my guest will be Richard Foster-Fletcher, calling in from the United Kingdom. He hosts the *Boundless Podcast* - which I was the very first guest on, so turnabout is fair play. He's a keynote speaker on the topic of the Artificial Intelligence Roadmap To 2030, he's the founder of Neuralpath.io, and the chair of the Milton Keynes Artificial Intelligence group. I've participated in his podcast and on several of the events that he's hosted, and he's really a thoughtful and compassionate futurist with an expansive view of where we are going that's driven him to take on this futurist role.

So, one futurist to another - that's what the interview next week will be on the next 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.