Since you're here, I have a small favour to ask.

We need your financial support to continue to produce and improve the Conversations.

This show, as you know, does not chase the headlines. We following the trend lines, we reach out to experts you rarely hear from. Experts who go into detail, experts who shed light on facts that enlighten and inform you many times in contrast to a trending perspective.

Conversations is an independent program, we choose the topics and the guests. No one pays us to air the program and no one pays to be on the program.

We have one sponsor who helps us cover the cost of our closed captioning and subtitle costs.

If everyone who watches and likes Conversations That Matter helps support the production of the show, we'll be able to continue to produce the show for years.

For as little as $1 you can support Conversation through our Patreon Account https://goo.gl/ypXyDs – it only takes a minute.

Thank you.

Welcome to Conversations That Matter. This episode is brought to you by Odlum Brown, a client-focused investment firm that starts client relationships with straightforward conversations focused on you, your aspirations, and your investment priorities.

Odlum Brown has been a supporter of Conversations That Matter from the day we started the show. Their only condition was that we provide a non-biased conversation with people from all sides of all sorts of issues.

Of course, we couldn't produce this show without the support of Oh Boy Productions. If you're looking to produce a video or a podcast, I suggest you reach out to Oh Boy. They can help you produce it, and they can help you build your audience. And we also need your support.

Please pledge your dollar per show at patreon.com/conversationsthatmatter because those dollars add up and play an important role in helping us produce the show.

Now to this week's episode.

Humanoid artificial intelligence is coming, and it very well may come from Vancouver. The brains at work creating AI, human-like AI, live here. And the odds that they will succeed are high. They have an amazing track record.

One of those brains is the mastermind behind the development of quantum computing, a development that has manifested itself into a company known as D-Wave. At the core of the development of humanoid AI sits the existential question, what does it mean to be human. What motivates us? How do we decide right from wrong? And whose morals constitute the foundation of the programming of the machine that will self learn?

These are just a few of the questions that surround what many believe will be the last great human discovery. We invited Geordie Rose of Sanctuary AI to join us for a Conversation That Matters about artificial intelligence. Why, what, when, where, and how soon.

- Conversations That Matter is a partner program for the Centre for Dialogue at Simon Fraser University. The production of this program is made possible thanks to the support of the following, and viewers like you.

- Welcome.

- Thanks for having me.

- My pleasure.

- You are making a difference in not only the world of quantum computing but, as a result of that, moves towards artificial intelligence. But I was fascinated, and where I would like to start today is, I watched a talk with you where you said, well, before we can understand what artificial intelligence will be, we have to understand exactly what it means to be human. Why do we have to start there?

- Well, it depends what kind of artificial intelligence you're talking about.

- Well, I'm talking about humanoid.

- That's the critical thing. In the original days of AI, back when the term was coined and people started thinking about computers and what they might be able to do, AI meant that. It meant trying to build machines that thought like people. But that is not the way that people think of AI today.

- They think of it more as machine-based learning?

- Well, there's a bunch of different camps. And I shouldn't say that nobody thinks about that because some do. But the majority of people who work in the field now are applying what are fairly basic ideas to two types of activity. One are what you could call this feature extraction. So imagine you've got a picture. And you want to assign a label to it which will say there's a person in the picture or there isn't a person in the picture. That kind of thing is very difficult to do with conventional computer science ideas. Because how do you describe a person? You could say, well, they've got two eyes, and a nose, and so on. But at some point, those levels of description aren't very good. In practice they don't work very well. If you think about what you would say to describe a cat, and you were to list all of these characteristics, dogs also have nearly all of those that you'd list. There's something about the cat and the dog that look different to you and it may be hard to put it into words. So, the majority of people who work in AI build systems that essentially do that for different kinds of data and different kinds of labels. That is not intelligence. Calling it AI is a big stretch. So, sometimes what people do is they call it, they start appending some technical jargon like machine learning, or supervised learning, or something like this. But it's hard for the layperson to distinguish.

- It sounds more like nuanced analysis.

- Well, it's... It's labeling, that's all. Humans are very good at this. If you look around the room, you can see a lot of things that you identify. And if you think about what the word of that thing would be, it pops into your head. That's very difficult for machines. And mostly what people have done in the recent few years of AI is try to replicate something like how humans recognize objects, first in imagines and then in audio. For example, extracting text from speech, and some other weirder modalities. But it's mostly visual stuff and audio stuff. And there's a few other things, but that's basically all of it. So that's the first. The second thing which is more recent and a little bit more interesting but still not intelligence of any sort are what are called generative models. Generative model turns that process in reverse. Imagine I were to say the word cat to you and ask you to draw a picture of a cat. So you're gonna draw something. So that thing that you've drawn is a picture of a cat. It's not the word cat. So the generative AI attempts to go from symbols, like labels, to objects in the world, like a video, or images, or audio. And those generative things tend to be where all the computer art lies. So if you ask a computer, please generate me a painting in the style of Picasso, what you can do is feed all of the Picasso images into this. It will learn certain features that Picasso uses and then combine them in novel ways to generate a new kind of

Picasso. And that works quite well also. Most modern AI is in one of those two categories. And we have to go back to the '60s, '70s, and '80s, to get back to what you might consider to be the original roots of AI, which is how do you build machines that think like people, that reason and process information the way people do.

- I've gotta get you to hang on for a second while we take a quick commercial break. We'll be right back.

- Conversations That Matter is a not-for-profit program made possible thanks to the charitable support of the following, and from viewers like you. Please visit conversationsthatmatter.tv and help us to continue to produce this program.

- When it comes to artificial intelligence, you do talk about the fact that to be able to get to that point that you just made about where you want to go, you have to first be able to define what it means to be human. Why is that important?

- It matters a lot if what you're trying to do is build human-like intelligence. Because in order to do that, amongst other things, you have to define what you mean. Which means being very clear about what you'd expect a machine that had this property to be able to do or not do. Humans do a bunch of things. There's a lot known about how we do them in terms of neuroscience and the way that the brain works. If you're gonna build a machine that, say, can have this type of conversation, that machine needs to be very like us. It needs to understand the world the same way we do. And you can kind of see flashes of this in the nonhuman biological world. If you try to have a conversation with an octopus, it's not gonna go very well. It's not because octopi don't function perfectly well as creatures in their environments. And they communicate with each other, of course. They have to. But we can't communicate with them. And it's not as simple a reason as people tend to think. People tend to think that there's this thing called intelligence which is like a linear scale. And we're here, and octopuses are here or something. That's not right. So every creature in the world has an unbroken chain of successful evolution going back a billion years. It's ideally suited to its environment and niche. Maybe not ideally, but well enough to have survived.

- To survive and procreate.

- So, we can't communicate with those guys not because we're smarter. Well, not only that. But there's also a problem that the way that an octopus would understand the world is derived from its experience in the world. It lives underwater. It has eight arms covered with suckers. It eats different kinds of things and needs to know different kinds of things in order to survive. So all of those things you can think of as the thought patterns of an octopus, which are so alien to us that even if an octopus was as "smart" as we are, we would have a lot of trouble communicating. Because the things that it would think about, we don't have human words for. So what does it feel like.

- We don't have the experience either.

- That's exactly right. So, words for us are symbols of experiences in the world. A cup, a word, is a symbol that represents our sum total of experience with cups in the world. If we didn't have the ability to use these things or wanting to, we wouldn't form those notions. And if we were to communicate with an alien that was, say, made of gas and lived in a cloud, it wouldn't have any concept of what that thing was in a very deep way. You would have very difficult conversations that would ensue with them. So, the reason why I think it's so important to understand how human bodies and minds work is that when you're building a synthetic analog, you have to know what it is you're trying to build. It's a zero-order problem. What is it exactly that you're trying to replicate. And so, we have to understand things like this.

- When I think about the speed that the processing takes place, for me to understand in this conversation what you're talking about, my brain is moving incredibly fast. How are we going to be able to create humanoid-type artificial intelligence if we can't process the information at that speed? So therefore how important is quantum computing to be able to get to that stage?

- Well, I think once we understand how thought works, we'll be able to use something like off-the-shelf computing technology to embody it. I don't think we need to have big breakthroughs in hardware. Now, the but. Usually what people use hardware for is not actually running something after you know how it works. It's doing all the experiments to figure out what works and what doesn't in the first place. So there, we're vastly underpowered. If you think about a company like Google, or Facebook, or Amazon, they have access to the compute resources that the rest of us don't. And that's a significant advantage when you're doing R&D. But that's a different question as to whether or not you could run, once you've understood everything, on those things. And I think you could. I actually don't think that the way that humans work is necessarily computationally intensive. There's a lot of evidence that it isn't because of the very low power draw in our brains. So, our brains, while they draw a lot of power compared to the rest of our body, they're not nearly as power draws as even modern... Like, a desktop computer draws more power than your brain does. And there are reasons to believe that that is an indicator that the actual computational load on your brain isn't that high. Now, it's difficult to know for sure.

- Could that not possibly be an evolutionary development that the body has said, I've got to find efficiencies in the use of energy to be able to power this thing. Otherwise, I'm not gonna be able to survive.

- We have, I believe, the highest percentage of our energy goes to our brain of any creature on the planet.

- 25%, up to 25%, depending on how active it is.

- So, that's a lot if you think about the other things that you have to do during the day. Shuttling more power than that was probably not effective as a strategy evolutionarily. You get to a certain point where I think there's probably diminishing returns on powering more and more of this. So we have things that we do quite well, but I think those things are... They're probably not hard in the computational sense. So if you sat down and said, tried to prove Fermat's Last Theorem or something, like some very, a longstanding, pick a math problem that's been open for like 400, 500 years. You're gonna have to grind away with your brain for years trying to think through all of the different possibilities, and you're gonna spend an awful lot of time doing that. But most of the things that people do are not like that. So for example, knowing that this is a cup, that happens in a very short period of time. Likely because in your brain the signals that reach your brain are hitting very specific detectors that are specifically like the cup detector. So you can think of it almost like an analog process. Light flows through this network. It turns into electricity. Those electrical signals light up a certain detector and all of a sudden you know it's a cup. So that's not a lot of compute load, but it's an amazingly sophisticated computing system that's been built to remove the compute load from it.

- So it's been streamlined.

- Yeah, everything has been designed so that we can detect things that are important to us quickly using these built-in capabilities.

- This is our second break. We'll be back in a moment.

- Conversations That Matter is a not-for-profit program made possible thanks to the charitable support of the following, and from viewers like you. Please visit conversationsthatmatter.tv and help us to continue to produce this program.

- Okay, well, I gotta touch on a point that I saw you talk about, that technology ultimately will make capitalism redundant. Or maybe not redundant but impossible to support.

- Yeah, they're incompatible.

- What happens then? You've developed these machines, but you don't need me any longer.

- UBI, universal basic income. And the way to implement it is very clear. The concept is that the... Okay. So, the way that things work now, in the simplified economic model, is you have companies who pay people to do things. And then those companies create wealth, a part of which goes to the government as tax. And then, the people who have generated the wealth, part of their paycheck goes to the government also. And the government gets tax from both sides. So as technology grows in its capabilities, what happens is companies replace the labor of people with machines. So they stop paying humans. They pay machines instead, a lot less. And the profit grows, but that profit doesn't flow back to the people. Instead, it collects inside the company. So there's a very simple solution to this. Because those people are no longer getting money from the corporations, and the corporations are making more and more, the corporations need to pay more tax. That tax goes to the government that now pays the people. So it reverses the arrow. Instead of you paying tax to the government, in the future, what's gonna happen is the government is going to pay you. And the reason for this is very simple. It's that unless that happens, the people who the government serves are going to have no capital. And they're going to...

- Well, ultimately, riot against it, or rebel against it.

- And so, you're gonna have this very terrible situation where wealth collects inside companies and there are only a few people who own all of that wealth. Like there's three people who own half the United States. That's where technology and capitalism in its current model ends. So what we have to do is replace that with a very simple system, and it's not difficult, and it makes total sense, that companies have to pay what you can think of as an automation tax. The more of a corporation that's automated, the more money they pay to the government in tax. And then the government redistributes that money to everybody in the form of universal basic income. So this idea is really simple. It will work, and it will be the way that we go forward.

- Will it be a better world than the one that we're currently in right now? Because there's an awful lot of acrimony.

- The bottoming out of the middle class and people who are wont to work but the opportunity is not there. And even redefining what a job is. Like if you think about what people count as jobs now, those shouldn't be counted. A job for an adult is a wage that allows you to support yourself and your family. It's not working for $15 an hour. That's not a job. That's what you do when you're a teenager. A job is something where you make enough money so that you can have the sorts of things that you'd expect you'd have in a first world country. So, I think that the will it be better, yes. And the reason is that you break the unhealthy link between labor and being able to live. So right now, everybody has in the back of their head this fear that if I don't exchange the sweat of my brow for money, I can't feed myself and my family, I'm gonna lose my house. And that is a terrible situation for a technologically advanced civilization that doesn't need that to be the case to be in. We can automate out today most people's jobs. And the reason why I think that's a good thing is that we can pay for all of that, again, by taxing corporations, this kind of automation tax, and then redistributing that money back to the citizenry. It doesn't mean that if I own a big profitable company I'm not gonna be incredibly wealthy. Let's say I'm worth $100 billion because I own company X. I can give away half of that to the government and I won't notice a difference. Because I can only wear one suit at a time, and the difference between $50 billion and $100 billion is not material to me. I don't care who you are. You will not notice a difference no matter what you're doing. There's no project that you can imagine where you'd need to invest $100 billion in it.

- This is our last break. We'll be right back.

- Conversations That Matter is a not-for-profit program made possible thanks to the charitable support of the following, and from viewers like you. Please visit conversationsthatmatter.tv and help us to continue to produce this program.

- So, you envision a world where artificial intelligence and technology is gonna rewrite the social order.

- Absolutely, and for the better. It will be all of the drudgery of life is going to be done by machines that are better at those things than we are. And it will be almost like, okay, right now, you don't think about the... Let's say when you turn on the tap, water comes out, and it's kind of free. You don't really think about that. What if everything was like that?

- I can't help but go to a place where I start to think, okay, what is the meaning of life? Does my life have value? And for me, I am contributing to something by my very existence. And therefore I effect change. If I'm not engaged in the world of work, what then do I do to have that same kind of impact and, hopefully, drive meaning for my time here on Earth?

- That is an excellent question. I think, okay, so... I'm lucky enough that I don't have to work anymore. I've been very fortunate financially, I suppose, in the things that I've done. But I do anyway. And I would claim, actually, that a lot of what I do could be better done by machines now. But I would do it anyway. And I think that this idea that you wouldn't do something that you loved to do because there's somebody or something that's better at it than you is a really... It's a wrong idea. So, work doesn't have to be an endless grind. Like, there are different kinds of work. You can enjoy what you do, and you would do it even if you didn't have to. A lot of content creators, I think, are like that. I've been listening to a lot of writers talk about their craft. And this thing that comes up over and over again is that most of them have jobs that aren't writing, so they have to stuff the writing in in the morning, or at sometime. And they report doing it because they feel like they have to. They don't have a choice. So, a lot of people have these things, like something that you'd really want to do if you could, or spend more time on it. And the fact that you have to work a job doesn't change any of that. And I think what will happen in this post-scarcity economy that I'm describing is that, yes, there are going to be a lot of people who just sit on the couch playing video games. And then there are gonna be people who are not satisfied with that, and they're gonna want to do other things. And they'll do those too, but they don't have to.

- Are you talking quarter century, or less?

- Less. On the timescale of say the introduction of the driverless car. So who knows when that's gonna be. But it's probably within, you know, 10 years, they'll be deployed in scale. So that is not too far away and that will displace a lot of workers. This is coming immediately. I used to be concerned that governments wouldn't be able to react fast enough to this because they wouldn't be able to get their ducks in a row, and people would be looking at the wrong kinds of problems. Like walls, for example. Where the things that really mattered were not being paid attention to. But I actually don't believe that now. I think that, at least in the United States on the Democrat side, this understanding of the dangers of the automated economy, virtually every person on that side understands that this is an issue. And I think it's gonna actually be an important issue in the upcoming election in the United States.

- It's an interesting world that we're moving into.

- It is.

- And we designed it.

- Yeah, well, I mean, I don't believe in free will, so I'm not sure we had anything to do with it, but...

- Okay, that's an entirely different discussion for another day. Thank you very much for doing this.