

## [Transcript] This Week in Startups / Reverse-engineering autonomy in humanoid robots with Sanctuary AI CEO Geordie Rose | E1832

I want to just make this very clear that my perspective on AI and automation is that there's an upward spiral when you have more energy, you have more intelligence, you have more capability. These drive all the metrics of human flourishing up. They don't take. So when we think about the answer, when will you get lights out manufacturing? I think the answer is never because people will always find new things to do with the tools that we've built.

Even very powerful tools that can think and maybe are even self aware. These will only increase the number of jobs, the increase wages, but there'll be different kinds of jobs.

There'll be the sorts of things that maybe we can't even imagine now.

This week in startups is brought to you by InTouch CX. Looking for ways to make your startup more efficient?

InTouch CX has a groundbreaking suite of AI powered tools for end to end optimization to give your business the edge it needs to thrive.

Get started with your free consultation at [intouchcx.com](https://intouchcx.com) slash twist.

Fount, do you want access to the performance protocols that pro athletes and special ops use?

With Fount, an elite military operator supercharges your focus, sleep, recovery and longevity, all powered by your unique data.

Want a true edge in work and life? Go to [fount.bio](https://fount.bio) slash twist for \$500 off and

.techdomains has a new program called startups.tech where you can get your company featured on this week in startups.

Go to [startups.tech](https://startups.tech) slash jason to find out how.

Hey everybody, welcome to this week in startups. We've been focused a ton on AI this past year.

Of course, we talked about it over the last decade on the show, but things have heated up with language models and

you know, the very forgotten category of startups is of course robotics.

We see once in a while on the internet a trending video and the trending video tends to be when a Boston

dynamics robot's doing a backflip or we see maybe some surgery being done on a grape.

You've seen all these viral videos, but the idea of humans leaving the factory floor and going and doing things in the real world, well, we don't see many startups doing that.

We have one in our portfolio, Cafe X making coffees at SFOs airport right now.

And of course, our friends over at Tesla are making Optimus and there's another startup called Figure.

They're working on a humanoid robot. Sanctuary AI is today's guest.

They're another startup working on this problem and they're specifically focused on building robots with general intelligence.

What does this mean?

Well, it's not verticalized and they're not just making a cup of coffee.

But what if these robots could solve problems in the same way we do as biological creatures, as human beings?

And if that works, well, that's going to have the economic impact of humanity.

And it's going to go well beyond just the steam engine.

And we have the founder or I should say the co-founder and CEO of Sanctuary AI on the program.

**[Transcript] This Week in Startups / Reverse-engineering autonomy in humanoid robots with Sanctuary AI CEO Geordie Rose | E1832**

His name is Geordie Rose. Welcome to the program, Geordie.

Thanks for having me.

Great name. I am reminded of the Mark Knopfler lyric from the amazing song, Selling to Philadelphia, where he says, I am Jeremiah Dixon.

I am a Geordie boy. Do you understand the reference Geordie boy?

I do, yeah.

So let's talk a little bit about the company.

And I know you were founded in 2018.

So you've been working on this for a while.

You've raised close to 100 million bucks.

Where are you at with building this humanoid robot?

And I would love to see the latest.

I should start by saying that our approach to the problem and the reasons for us working on it are slightly different than most people who work in robotics.

For us, the motivation for doing it was a belief that human-like intelligence and more generally the intelligence of animals, which is kind of our model for what intelligence means, is very intimately tied to our presence in the world.

We have a body. We are a thing. We experience the world through our senses.

We develop understanding of it through interacting with it.

And then we act on it to achieve our goals.

All of those things are very difficult to do if you're not actually physically present in the world. So the starting point of this, which actually goes back more than a decade now through two different companies, was to explore this idea that intelligence, by which I mean general intelligence, emerges as a consequence of having to deal with the real world.

The real world, you never see the same thing twice.

You have to be able to generalize from your previous experiences to new experiences.

You have to be able to understand the common sense ways that the world is.

So we've been building software, which you could call general intelligence or AI, but it's also control systems for robots.

And we've always viewed the problem of artificial general intelligence through that lens is that, for us, a true general intelligence can be thought of as a control system for a robot that converts what it sees, hears, touches, feels about the world into actions that are intended to reach goals. So for us, the robot is somewhat means to an end. And because of that thesis, we focused almost exclusively on a very hard but very fundamental problem, which is the building and use of hands. So much of the humanoid robotics videos are performative.

They show robots doing things, but they're not valuable things. And for us, I think that the key value of doing this is to understand how an entity, a robot or a person, understands the world well enough to be able to manipulate it with its five-fingered opposable thumb hands. I believe that the hand played a big part in our technological evolution and also in the development of language, which are related things.

So that's what the thesis... How did that happen? How did the hand

play a role in language? I'm curious. Was it writing or the ability to hold pen?

Speaking. So how does a hand help you speak?

## [Transcript] This Week in Startups / Reverse-engineering autonomy in humanoid robots with Sanctuary AI CEO Geordie Rose | E1832

Yeah. So although this is speculative, there is a lot of evidence that the earliest spoken language was very strongly connected to the things that our hands do, like point, touch, grasp and so on. And some of the evidence for that is in neuroscience, where the part of your brain that controls the grasping and the use of the hands overlaps with your language center. These things are not disconnected. And when you actually try to build a system that touches, feels the world and can interact with it in the way we do, you see this explicitly, is that the cognition, you think of it as the domains of intelligence, many of them and maybe all of them are required in order to do something with your hands. It's a remarkable thing that a planning, reasoning, logic, all of these things are connected to the way that we interact with the world through our hands. That's fascinating. Like when you were saying that, I was thinking, so I put my hand on my chin. And then if I were to, if you and I were navigating the world, we were, you know, early sellers or somewhere, we might point towards the direction we want to go or I might put my hand on my chest to refer to myself or I might put my hand out and my palm up to refer to you in some sort of gracious way. Is that what we're referring to, this sort of instinctual thing that happens with our hands as we're talking? Yes, our view and the position that we've taken is that the hands and their use and the mind are interwoven in an inseparable way in people. So if you want to understand human-like intelligence, the kind of intelligence that we have, the hands are the appropriate starting point and that's why we focus so much on them. And I asked to see some things. I can actually show you one of the hands. Yeah, all right. So I see on the screen here, you've got, yeah, a very interesting looking hand, five digits, yeah, four fingers and a thumb and a palm. And it looks like something out of Terminator, but a little more elegant, in fact. Well, I think that I would not characterize it that way. I think that the way that we imagine this hand is that it's the best that the technologies that the global community knows how to build. It's the best that we can get to human hands. There's a lot of things in this hand that aren't immediately obvious just by looking at it. And those are mostly about the sensors. Our sense of touch is a very important thing for our intelligence and how we are in the world. We tend to take it for granted because it's always there. And when we look at screens and things, people are very visual and they think about the world in terms of seeing, which is fine. But there's an interesting observation that seeing is about the future. It's about planning, because the things that you see are away from you. So for example, if you look at a cup and you want to pick it up, the part of your brain that plans thinks about the future, but touch is a little different. Touch is an immediate thing that's in the now. Touch doesn't have foresight. It's all about the present moment. And when you make contact with the world, say you're seating in a chair or you're picking something up or you're turning a Rubik's Cube in your hand, the sense of now is intimately connected with the touch sense. And without it, you can't be who you are. So this is an important thing for building robots is that touch is not a second-class citizen if you're trying to build a system that behaves in things like we do. And so these hands are covered in very sophisticated touch sensors that allow them to feel the world, something like we do. Yeah. So when we're looking at each digit, I guess we have a couple of

knuckles.

And so the tip of your finger has one pad, then that middle knuckle, I guess has another pad, and then there's a longer pad in that third spot. So if you're looking at your own hand, you have those sort of three segments of a finger, they each have a pad on this robot. And the thumb obviously has the same configuration. And I guess when they touch each other, that's telling it something's hitting that pad. Is that correct?

It's more general than that. So the sensors in your hand are not just about a yes or no question about you're touching something. They're very rich. You can feel temperature. You can feel if something's sliding past your fingers, which is very important when you're trying to hold something or turn it in your hand. Imagine trying to put a key in a lock and turn it. A lot is going on in that thing in your brain. And a lot of it is driven by touch. If you didn't have the sense of touch, it would be very difficult to insert a key into a lock and turn it. Even something as simple as that. Because of resistance, right? You have a certain resistance on either side and the top or the bottom of your finger where it's touching the key. I'm imagining this as you're saying it. Yeah. The way that we do things in the world is not... We take it for granted. There's this thing called Morvex paradox, where the things that we take for granted and are easy are actually some of the hardest problems that there are. And the reason they're easy is that we've had a billion years of evolution to create a system that is fine-tuned to be able to deal with things like picking things up and putting things in things and things like this. But the reason why we have AI systems that can write at the level of GPT-4 or create images from scratch that are as beautiful as any human artist could draw, the wonders of the digital age, but we don't have a robot that can do laundry, is that doing laundry is a fundamentally much more difficult problem than any of the ones that modern AI has managed to master. Yeah, it's fascinating when you think about how complex this is. And that paradox you mentioned, that seems like a fascinating evolutionary moment. These systems are so complex that they must be automated. Because to actually, with cognition, to try to think, okay, I'm going to have to give some resistance as this key goes into it. And I'm going to have to feel it click a couple of times. And then I'm going to have to twist it left. But I'm going to need to put more pressure on the inside of my index finger versus my thumb. And if I put too much pressure, I'm going to break the key off in the locks. I mean, it's incredible when you think of all of that occurring. And it's occurring in just an automated fashion. It's just a chunk of a task to open the door. It's not even one test either. It's probably open the door, which is the key getting taken from your pocket, being put into lock, twisting, open the door, closing at the whole shebang. It's just abstracted into one instruction set, huh? Well, that's an interesting phrase. And I'm glad you mentioned it because that's the way that most serious embodied cognition efforts work is they have an idea of an instruction set, which is very similar to the way that processors work. I worked my first half of my career in building computer systems. And the marvels that we've built in the computing side are fundamentally based on a very non-trivial fact that every program that you could write boils down to the execution of roughly order 100 different tiny little programs that just happen in different orders. So every program that you can write on a computer is basically composed of only about 100 building blocks in modern processors. The reason you can do that is that processors have a natural way to turn the analog nature of the world into a digital character

that allows error correction, which is the fundamental reason why you can do all the things that we do today, that in the computer, there's this thing called a transistor, which is the basis of this digitization going from things that are just any number at all when you measure them like voltages and currents to something that's only a zero one in motion, that is the taking of actions in the physical world. You need to be able to find a way to do that same thing. And so what we've done is created this type of instruction set, which is a very small number of building blocks that you can compose in different orders to create massive complexity of tasks and potentially all of them. So if you think of a robot moving through the world or a person as a program, then you can imagine that any program could be written in say maybe just 100 things in different orders and that's what we're trying to do here is to figure out what those 100 things are and then use a technique called task planning, which is the idea that given a goal, let's say I ask the robot in natural language to do something, the robot can figure out how to sequence the things it knows how to do in order to achieve the goal and thereby achieve general intelligence. Because if I can ask the robot to do anything at all like in the human sphere and the robot can actually perform the task, then it would be fair, I think, to think of these things as reaching the goal of having general intelligence. By the way, I should mention that there's a concept that David Chalmers is a philosopher introduced called a philosophical zombie, where a system can have the appearance of being like us in the sense that it can do things like we do, but it doesn't have the first person conscious experience that we have. So there's a lot of mysteries about the relationship between being able to actually achieve goals, do things and whether that is related or different than the experience we have of being people, this thing that it feels like to be a thing. That's a deep, deep mystery. All right, listen, efficiency needs to be top of mind for every founder in 2023. Fundraising is drying up, so you need to extend your runway. And one great way to do that is automation, but it's hard to apply automation in your day-to-day operations, isn't it? So here's an amazing solution. Intouch CX provides easily integrated automation tools for customer support. You're wondering how it works? Well, let me tell you. Intouch CX provides automated and live chat, email and voice support. This eliminates unnecessary process and cost, and it will make you faster. Intouch CX will streamline your customer support process, cut back on repetitive and time-consuming tasks, and increase productivity by 30%. And it's going to simplify your business.

So refamp your workflows with Intouch CX. Intouch CX partners are experiencing 45% average cost savings in customer support ops so far. Find out how Intouch CX can improve your startup's efficiency. Get a free consultation with their automation experts and get started at [intouchcx.com-twist](https://intouchcx.com-twist). That's [intouchcx.com-twist](https://intouchcx.com-twist). Yeah, consciousness, the big C is for philosophers and religious people. What is consciousness? Is it just some illusion that we're having in this brain of ours, which is a collection of a bunch of subroutines as you're sort of alluding to here? Or is it this God molecule in our brains making us sentient and driving us to do things? I guess this is one of the exciting things about AI, is that we're in some way, or in your case, quite literally, trying to deconstruct and then reconstruct what is happening in cognition. But it has to start with, hey, pour me a glass of milk. So it has to know what milk is. Easy enough to do now with visual computing. But pour, okay, we know what that word means. It's moving some liquid from one place to another.



**[Transcript] This Week in Startups / Reverse-engineering autonomy in humanoid robots with Sanctuary AI CEO Geordie Rose | E1832**

And then we have to then, of course, make it do that accurately. So if we were breaking down a task like that, and you said there's about 100 things you're teaching, what are those little subroutines or those micro behaviors, or you used a term for it, what was the term used? We call them micro policies. So in the world of reinforcement learning, which a lot of this is grounded in, we came up through the reinforcement learning school of thinking about cognition. A policy is an action that you take from a current state. So it's a prescription for how you act, given the observation of the world that you have. So these micro policies are a collection of very specific types of behaviors, like say, for example, turning a key in a lock that we train individually in isolation from any other use. So the way that works is that we take the robot and we, a person who is teleoperating the robot, which is a process of a person controlling and being kind of immersed in the robot, receiving the senses of the robot and moving the robot through a rig, which is another type of robot that the person is strapped to. So the person moves the robot to accomplish the task, because the person knows what it means to pick up a key input in a lock and turn it. And we collect order hundreds of episodes, which are instances of solving that problem. And then we use that to seed a thing we call large behavior model. So large behavior model is much like a large language model, except the fundamental data is the data of experience. It's vision, audio, proprioception, which is the information from the servos and the robot, where it is and so on, how fast it's moving, and touch haptics. So if you can use the same idea where you take a bunch of data, and instead of predicting the next word or token in a text prompt to response, you take the past, which is the things that have happened to the robot, and you predict the future, which are the analog to the large language model predicting the next tokens. But in an interesting twist, the predictions of proprioception are predictions of where you will be, how you will move. And you can then send those predictions to the actual motors, and the motors can move. So that one of the most fundamental churning points in my professional career happened about 10, 12 years ago when I and some colleagues read Jeff Hawkins book called On Intelligence, which was really the first thing that I read that made sense about the potential model of human cognition. And central to that story was the idea that our brains are predictors, is that we imagine the future, and then we implement the imagining. So if I decide to pick up a cup, my brain is predicting how my motor's signals will fire, and then it sends those predictions to my muscles, and then I perform the task. So these large behavior models that we and others are working on now are of this sort, is that they predict the future based on the statistical properties of the data that they've looked at, and then they execute the tasks, and they work quite well. So for a human being, we're going to perform a task, we're going to pour this glass of milk, we will, in our minds, and we do this either consciously or maybe even subconsciously, okay, I'm going to pick up the glass, I'm going to pour the milk, I'm not going to spill it, it's going to pour in some kind of an arc, I'm going to watch it fill, I don't want to splash, I'm going to stop pouring at a certain point, and you kind of visualize this movie in your head, this potential future behavior, and so our minds are so powerful, they can actually essentially play a scenario, almost like a screenplay, like a little vignette, and then our muscles actually go play that routine, is that the concept here in terms of intelligence of what happens in our brains? Yeah, so an analogy would be, let's say you take a piano keyboard, and each of these little micro-polices that we're talking about

**[Transcript] This Week in Startups / Reverse-engineering autonomy in humanoid robots with Sanctuary AI CEO Geordie Rose | E1832**

are one of the keys on the keyboard. Your brain, your mind, because I want to make a distinction here is that I've come to believe very strongly recently that the mind is a creator of stories about the future. You, the awareness that you are, your conscious presence is not your mind, it's a different thing, you touched on it briefly, I don't have any proof of this, but I'm much more of the mind that the, there's a big mystery there about what it is to be the thing that is you, but it's not your mind, your mind is a machine just like your heart, and the job of your mind is to produce stories. So in the analogy to the piano, think of the mind as creating sheet music, and then the sheet music is automatically put on the piano and the keys are played and you hear a melody or a song. In the analogy, the song is the behavior, like for example, picking up a glass of milk and drinking it. All of the behaviors that we exhibit in this model are all different songs that are generated by pressing the keys in different orders and with different styles. So the mind is the creator of the sheet music that it does always, this is, brains are always doing this, and then it's played on your body, like a song. And, and our conscious perspective is sometimes not aware that these are separate things, you know, we have difficulty introspecting our minds and our behaviors for a variety of reasons. But I think that after working on this problem for a long time and seeing the, the synthetic analogues of us, how it actually works in robots, I think this is a good model is that your mind is a machine for creating sheet music that is immediately played on the, the instrument, which is you and your awareness is a separate thing that kind of watches this. And sometimes we get confused and we think we are the mind, but I don't think we are, I think we're something separate.

So there's the mind and then there's the machine. And this machine is going out, conceiving of these tasks, executing on them, playing the sheet music, running through the script. I use this, the analogy of a film, using the analogy of piano, but the script gets played, the sheet music gets played, it happens. But our awareness that I am a human, I am Jason, you're Jordy, we are on a podcast, we're having a conversation, I'm trying to understand what you're doing, you're trying to understand my questions. And then there's going to be 100 or 200,000 people who listen to this. And they're also going to try to, that's consciousness, the awareness of each other and ourselves and our place in the universe and that there is even a universe. These are two different things. But for some reason, we perceive these two different things that are occurring, the mechanical execution of tasks, through this very interesting project, process, which you're now recreating is different than consciousness. And consciousness, who knows when we're going to ever figure that out? Or if we can figure it out, this idea that we're aware that we are a living being, but we can figure out, at least at this point in time, it feels to you like we're going to figure out, and we're close to figuring out how our brains break down complex tasks and do them so elegantly. Am I, am I repeating that back to you correctly?

Yes, that's right. There's a spiritual leader, I suppose you could say, called Eckhart Tolle, who refers to the first thing, which is the not knowing that you are different than the plans that your mind makes as being unconscious is the phrase he uses. And I think that it's a natural state of people is to not be aware that the mechanical following of scripts, which is most of our behavior, and that's the sort of thing that you can, there's a shot to doing in a robot. So I'm fairly sure that we can build machines that can do all work, like all of it, that at least

## [Transcript] This Week in Startups / Reverse-engineering autonomy in humanoid robots with Sanctuary AI CEO Geordie Rose | E1832

does it's currently, you know, understood things like automotive manufacturing, logistics, bringing parcels to your home. I think that all of those things are within scope to do within say a decade, at least have the capability. So this idea of building a thing that appears as though it's intelligent and does all the things that you'd want, that's within reach. But the thing that I'm really kind of taken with is this other notion, you know, I used to be a theoretical physicist a long time ago, and I worked on foundational problems and quantum mechanics and general relativity. And I've always had, you know, at the base of who I am, I've always been interested in understanding how things work at some fundamental level. And it's always bothered me that all we ever experience of the world is this first person thing, the feeling of being you in the moment. But we don't understand that at all. And I think that this neglect of what is probably the most central direct experience we have of the world means that there is a discovery waiting to be made about the relationship between our experience and notions of space and time. And I think that this project is somehow, in some ways, aimed at that, is that, you know, it sort of starts from a weird spot because you see these robots, and there's mechanical hands, and then I'm talking about, you know, some fundamental relationship between the emergence of space and time and how we perceive it through our conscious perspective. And as he seems to be not related, but I actually think that they are very tightly related. You see these blue light glasses I'm wearing? I'm not wearing it for style, although they are very stylish. They've totally changed my life. Why? I started having headaches, right, and had eye strain. So I got these blue light blocking glasses that do a little magnification because I need readers. Yeah, I look nuts, but my eye strain's gone down, my headaches have gone away, and I'm sleeping better. Do you know how I got on this? I got on it because I now have a health coach. Who's my health coach? It's F-O-U-N-T. It's a health company that's created custom health and performance programs that are tailored to your body, obviously, also your goals, and they take into account your lifestyle. My coach is incredible. I text with him all the time. They did a blood work for me. They check out my wearable data, and we do weekly calls to see if I'm on track and getting the results I want. They also told me about some supplements I should be taking based on the blood work, and they do it at a fraction of the cost. We upgraded my diet. I'm doing a little more protein. We've optimized my sleep. That's great. I got the supplement packs. I feel great. I feel like I'm in control of my destiny. If you want to be like me and you're concerned about your health and you want to just try to do better, have some experts on your team. Build your own program. Go to [fount.bio-twist](https://fount.bio-twist). That's F-O-U-N-T.B-I-O-slash-twist. Get your free consultation. Mention twists. You get \$500 off your first month, and get your own personal health coach. Health is well, and if you're running a startup, if you're a CEO, if you're a capital allocator, take it seriously. I love this service. [Fount.bio-slash-twist](https://fount.bio-slash-twist).

Well, if we think about the experience of being human and our place in this universe that we're trying to figure out, performing the tasks, as you said earlier in our conversation, is how we navigate the world, and it's how we are actually doing this act of trying to figure out what it is to be human. And this all then starts to open up all kinds of possibilities, free will. When we pour that glass of milk, when we play that sheet music, where is our decision to do that occurring? What parts of it are automated? Which parts of it are just wrote and just get executed on? And so it does open up. And I agree with you. This is the question that we will always try to figure out. And this is why science fiction always winds up here, which is what does it mean



to be human, whether we're talking about Blade Runner or Prometheus in the Alien series and really Scott's take on it. So let's get back to reality here. When you're training the robot, you are not saying, hey, we're going to pick up a tennis ball here. If we're going to pick up a tennis ball, it has this size, therefore we're going to program it to pick it up. That's what people did with robots before. They very explicitly had it do some very narrow verticalized tasks. You're having a human being, like the guy who played Gollum, I guess, in Lord of the Rings, use gloves or something to send the instructions to the robot's actual physical actuator hands. And they're incredibly sensitive and have those pads on them. So we're teaching it, hey, I'm going to just pick up Andy Serkis. Yeah, he was a guy who played Gollum. We're going to actually just pick up the tennis ball. And then the AI that we train is going to know what happened. And that's what's going on here. Yeah, so I have got a, this might be helpful. Can you see the video? Ah, here we go. Yeah. Yeah, we got a video of a robot. Yeah. Yeah, so this is Phoenix. And what you're seeing is this process that we're talking about where there's a person in a suit, they have haptic gloves with force feedback so that they can feel the world they see through a heads-up display so they can, they feel like they're looking through the eyes of the robot. And they're connected to it, their own robot that when they move, moves the robot in an analogous fashion. So the, yeah, so this is what it looks like when you watch the robot side of teleoperation. You can see that these machines are capable of doing lots of different things. I mean, that might not be obvious from watching this, but the systems are nearly capable of doing anything that a person can do under this type of control as long as they don't have to move around the world. This is focused on the upper body stuff and the problem that I mentioned, which is the dexterous manipulation of the world. And you're seeing it there, you know, basically pick up an object and then scan it with a barcode as if you were working in an Amazon, let's say factory and shipping and packing boxes or even doing something as delicate as using a Ziploc bag, which we do unconsciously. We feel it, it feels like the Ziploc bag, yellow and blue may green and you just, you have that color system, but you also have the feeling of it. So humans do the tasks and then take me to what the software then does with the human having done the task. What does it do next in terms of building a model to then go do the next thing in the world? Yeah, so imagine you have a Reddit post, which is some sequence of words that someone says, I really love Diablo 2 because Amazon is my favorite character. So somebody's written something that sentence is the expression of a thought that a person has had into words. Now, when you train something like a GPT large language model, that sentence is used to help figure out the statistical likelihood of each word. Let's just keep it simple, following the preceding ones. And if you give this model enough words that people have written, the expression of their thoughts, then if I was to say type in a prompt, which is my favorite game is, then of all of the words that have ever been written to some approximation, there's a probability of what the next token will be given that prompt. And then the thing can unroll, which means I put the next word in and now I ask with all those four words, what's the fifth word? Okay, put the fifth word in. What's a six? And each time it's a probabilistic thing. So you roll a random number and you pick the thing that the random number says it should be. So with this type of model, these large behavior models, the data is a little different. It's the data of the sequence of successive nows. It's the time data from the person performing the task. So if I ask the robot to open a Ziploc bag, let's say that's the micro policy that we're going to train. So a person

**[Transcript] This Week in Startups / Reverse-engineering autonomy in humanoid robots with Sanctuary AI CEO Geordie Rose | E1832**

picks up the bag from the table through the robot, opens the bag, and maybe pulls it open a little bit. So that then becomes the analog of a sentence. It's a piece of data, which we're now going to use to train a model, where instead of predicting the next word, we're going to predict the next sequence of actions. And we unroll the same way we would a sentence. So every successive prediction becomes a movement pattern for the system. And in this case of the kinds of things we build, while it's similar in some sense, there are some very big technical difficulties in actually doing this that require the synthesis of many different kinds of artificial intelligence advances. For example, you could send the pixels from the camera in at every step to one of these models. But the pixels are not the thing that you really care about. What you really care about is where are the things and what are they, which is a much lower dimensional thing. So machine learning computer vision techniques have been developed that will take the camera feed and extract what you could think of as the semantic or important information about the scene. And those are typically the things that you put into these types of models. And that's not just true for vision, it's true for haptics and audio and proprioception. So on the audio side, the obvious thing is if a person speaking, you could use the actual audio waveform, but you could also use the text. And text is a much more compressed and high quality version of the data than the actual audio itself. So we tend to do text extraction from speech before we send them into these types of models as well. That's fascinating. So you can know with machine learning, hey, there's a bag in this scene and the bag is open. And but the bag is upside down. So it needs to be up. We should flip it around so these things don't fall out of the bag, etc. And so where are you at? Let's I think I understand what's happening here in terms of the language model analogy, and then just translating that into predicting the next best thing to do. And so where are you at in terms of training this in the real world? I assume that factories, and the example you give looks like a, you know, a packing and shipping, probably one of the most boring, monotonous, soul crushing jobs a human being could have. So why not give that to a robot? And sure, you could do it 24 hours a day or whatever the robots are going to be capable of. So where are you at in terms of taking this and actually having it at a fulfillment center, packing boxes and making sure that it scans them and puts the right objects into the box and then ships them onto the next person and being on you know, this distribution center floor. To be clear, the initial go to market is in automotive manufacturing. It's not in logistics and retail. We focus almost exclusively on that with one exception. In automotive manufacturing, if you take a look at a video that, say, Toyota makes of their factory floor, automotive manufacturing is one of the most automated systems there are in any industry. But if you watch what actually happens, there are hundreds or thousands of people in automotive manufacturing facilities all the time. The question is why? Why are they being automated? So when you look at what they're doing, there's kind of two categories of answer. One is it may be beyond the bounds of science. We may not know how to do the thing that they're doing, but there's another answer, which is that often people are used to connect machines. So let's say I have a machine for stamping a part and I have a machine for putting, making the part in the first place. So moving the part from the one machine to the other machine is a very difficult process that involves all of these things that we talk about. You need to be able to know what a thing is, where it is, localize it, use your hands to

**[Transcript] This Week in Startups / Reverse-engineering autonomy in humanoid robots with Sanctuary AI CEO Geordie Rose | E1832**

pick it up, sometimes out of a cluttered mess, put it somewhere, which often requires putting something on a jig, which is a difficult thing. You need to be able to move around and so on. So a lot of the work that's done in automotive manufacturing specifically is a combination of different solve problems that have never been put together in a way that you could make economic.

And one of the key factors of these general purpose machines that we and others are building is that this is exactly the kind of thing that's required in order to actually do this for real, is that if I was to spend all my time and energy building a machine that did one of the things that somebody in this factory floor does, it would be very difficult to build a business. But if you could build a machine that could do, say, 50 of the kinds of things, now we're talking. So our initial use cases are nearly all of this sort. They're automotive manufacturing, they're the connector problems where you're moving between machines with parts or things of the material. And even the things that aren't automotive manufacturing that we've looked at all share the same feature. For example, in warehouses, which was my last business built robots for e-commerce distribution centers, there's a problem called induction. So induction is the problem of taking things that usually come off trucks, big pallets and just stuff. Imagine all of the things that you could buy on Amazon coming into a warehouse. And then taking them from their point of delivery and then getting them wherever they should be in the system, on a shelf, in a box, whatever. So induction is another kind of problem that's related to this, where you're dealing with system things that you need to manipulate with your hands, opening boxes, closing boxes, putting things in boxes, taking things out of boxes and so on. And so that's another category of things that's related. But nearly everything we're doing now is helping automotive manufacturers dramatically improve the efficiency and productivity of their workforce. We're back with another Pitch It to J-Cal. This is this segment brought to you by our friends at Dot Tech Domains. Dot Tech Domains are giving twist listeners a chance to show off their startup on this weekend's startups. So go to [startups.tech.json](https://startups.tech.json), that's [startups.t-e-c-h slash Jason to apply](https://startups.t-e-c-h slash Jason to apply). There's only one rule. You need to have a Dot Tech domain name to get featured. This week, I received a great pitch from Label Drive, which you can find at [LabelDrive.tech](https://LabelDrive.tech). Label Drive helps other companies manage their AI data. And they've built a tool for collecting and labeling data that's especially focused on identifying and catarising objects to save your time, save your money and build better products. And as we all know, that's crucial for AI training. So I want you to go right now to [LabelDrive.tech](https://LabelDrive.tech). And if you're interested in getting featured on this weekend's startups with your new Dot Tech domain name, I want you to go to [startups.t-e-c-h slash Jason](https://startups.t-e-c-h slash Jason) and apply today. That's [startups.t-e-c-h slash Jason](https://startups.t-e-c-h slash Jason) and fill out the form to apply. If this works, when do you think you'll have the ability to have the robot find those 50 different things to do? Let's say you nail that and feels like you're well on your way. The first question, when do you get that solved and in factories and just doing it day in and day out? The plan that we've got takes us from where we are now to the full automation of important tasks by which, I mean, there are markets for, say, billion dollars of annually occurring revenue for us. So let's say that's a kind of the thing that we want to target. We enter into agreements with our customers where the first step is that we mock up their situation in our facility here in Vancouver. Think of it as a digital twin or not a digital twin, a real-world twin. There's also a digital twin, by the way, but the real-world twin. And then the processes that they pay us

## [Transcript] This Week in Startups / Reverse-engineering autonomy in humanoid robots with Sanctuary AI CEO Geordie Rose | E1832

to show that we can automate using this type of thing, the kinds of tasks that they want as a first step. So there's a period of roughly two and a half years that we see where we go from where we are today to being able to really do something for real in the lab of the sort that you could then scale. So that's the first step. When we start scaling is likely around the middle of 2026, where you're going to start to see the increasing number of these types of machines actually deployed inside automotive manufacturing plants contributing to the productivity of the plant. So this is the plan of record. Now, I've done quantum computing and all sorts of things where it's very difficult to predict how things will go. So in something like this, you have a plan, things could go faster, they could go slower, it's unclear, but that's what we're aiming at. I think you'll start to see the beginnings of large-scale deployments of these machines somewhere in 2026. And so 2026, you start seeing the deployment of these. And then when do you think factories start to remove humans? I guess they call that the lights out moment. You don't even have to install lights in a space. I know it's funny, but when do you think you have that lights out moment and factories don't need to have humans in it? So I want to make a point about this. There is a myth that AI and automation reduces labor. It's not true. Throughout history, there have been a series of moral panics where the next big technology thing is believed to do something terrible to employment. It's never happened. Every single time there's been a new thing introduced. And I think that the central reason for us thinking this is that it's the lump of labor fallacy, the idea that there's a fixed amount of work. And if you give the work to the robots, there's nothing left, that's simply wrong. The way that it actually works in practice is that when you give, say like you give a bunch of labor, like I want to build 80 million cars. So that's a fixed amount. Let's say we could do that all with robots. The amount of work that's available for the general human population expands as a consequence of that. It doesn't shrink. So I want to just make this very clear that my perspective on AI and automation is that there's an upward spiral. When you have more energy, you have more intelligence, you have more capability. These drive all the metrics of human flourishing up. They don't take. So when we think about the answer, when will you get lights out manufacturing? I think the answer is never because people will always find new things to do with the tools that we've built, even very powerful tools that can think and maybe even self-aware. These will only increase the number of jobs, the increase wages, but there'll be different kinds of jobs. There'll be the sorts of things that maybe we can't even imagine now that are made possible by these things. Like look at the internet 20 years ago, podcasting. This is a great example. Yeah. Now we have an entire podcasting history. We have people who take pictures or there's an incredible company called Songfinch. What they do is you go there and you tell it you want to make a song for your mom or your dad. They pair you with an artist and you pay them \$200 and they'll write a song about your mom for her birthday. That's very cool. Well, I mean, there's humans out there and I guess these used to be bards or court jesters or whatever who would do these kind of tasks as well, but we find things and you're just thinking about your robot and, oh, well, we have this new problem, forest fires. How are we going to clean up the, how are we going to rake up as our former president joked of, how are we going to rake up all that debris under the trees there in the mountains in California? If somebody had 10 of these robots able to do a task, they might say, oh, you know,

## [Transcript] This Week in Startups / Reverse-engineering autonomy in humanoid robots with Sanctuary AI CEO Geordie Rose | E1832

I have an interesting idea. Maybe we could clean up and do some deforestation with them and they will eventually, in your mind, a decade from now or two decades from now, not just be in factories. They'll be in our lives. They'll be side by side with us solving problems in the real world. Yeah, that's the ultimate vision here so they can leave the factories. Yeah, I think of them as being a kind of thing like the automotive industry where at some point they'll be ubiquitous and parts of, and our entire civilization will be built in synergy with this new thing like we did with cars, you know, roads and so on. By the way, I wanted to mention that this happened, this business of the job upgrade happened to me when I was starting school. There were no quantum computers at all, except maybe theoretically. And we started a company to try to build one. This is an example where the, we probably hired about over time, I don't know, maybe 300, 400 people who had PhDs in physics in that company. And this is D-Wave, yeah. This is D-Wave, yeah. That was a new kind of job that was created as a consequence of a revolutionary new idea. So this is the sort of thing that always happens with innovation is that, and I'm kind of emphasizing this a little bit because we're at a very weird time right now where there's an attempt to do regulatory capture in artificial intelligence. It's a very dangerous idea. This idea of de-acceleration or stagnation or holding back, which are connected to ideas of the old ideas that were rooted in communism. These are very dangerous social ideas that I think it's important that we don't stay silent about. People like me who have very strongly held beliefs that technology is the solution to maybe all of our problems, not only the ones we create, but also the ones that might emerge as a consequence of our natural habitat, global warming or meteors or whatever. The idea that the better we can get at creating new things, the better we are all is a very important policy idea that I don't think is being communicated effectively by the community of people who build technologies. There are a group of people who want to have the government and they've specifically gone to Washington and said, hey, please regulate us. And there happened to be the people who were at the, maybe at the forefront or some amongst the people at the forefront. And building a bunch of regulation into this would benefit the people who have the lead today, as opposed to say open source people or folks who are coming up. Is that the thinking of what their motivation is? Because this is a group of technologists who are on the cutting edge. Why would they go to Washington and want to have a bunch of non-technical politicians slow things down? What do you think their motivation is? I think the best answer to this is somebody that you had on Bill Gurley. Yeah. So his take on this I really resonated with. It was one of the most, you know, sometimes you watch something and you're like, I'm disagreeing with everything this guy's saying right now. I think I would, if people are interested in this subject and they haven't seen that, I would most definitely recommend it. Bill Gurley all-in talk. Yeah, we'll put on the show notes for everybody. Yeah. But the regulatory capture is what they're going for. It calcifies the winners as the winners. It builds up a moat for them. And this could be just cataclysmic for humans, right? We need this technology to solve problems. Yeah. And I think that's the point is that the solving of problems comes from innovation and growth. And the forces of stagnation, the people who are pushing for not that are very strong right now. And I think it's dangerous because I think that my view of this is that civilizations metrics like how well people are doing



**[Transcript] This Week in Startups / Reverse-engineering autonomy in humanoid robots with Sanctuary AI CEO Geordie Rose | E1832**

are very strongly correlated with growth. And there's an idea that we have to slow everything down, which is, I think, a dangerous idea. I think that what would happen if we were to implement policies that were restrictive is the same thing that happened. I'll use an example with nuclear power. A lot of the problems that we face today in the global warming sense and catastrophic potential futures that we might be looking at are connected very strongly to the precautionary principle, which in the nuclear industry was, well, we don't want to build nuclear power plants because we're afraid of nuclear bombs, which is ridiculous, by the way, because they're not the same thing at all.

They're not the same thing. Yeah.

And maybe a reactor melted down once or twice.

You don't count all the deaths that happened in all these other industries, which were massively higher. If we had not done that with nuclear instead embraced it, we would not be where we are today. And so there's examples of this fear, which is a rational fear that can become policy that it could be very dangerous here, because I think that these technologies we're talking about, which is the AI robotics to a certain extent, but not just those things more generally, we should take the attitude that the upward spiral is the objective. We want more energy. We want everybody on the planet to come to the energy consumption of us. We don't want to reduce energy consumption. We want to increase it. And then we want to increase everybody by another 1,000 times. And we need to be able to find ways that technology can enable that and then enable solving the problems that might come of it from second order effects like global warming. These things are all solved by innovation and technology. Innovation and change is not the enemy. It's our friend. It's a necessary part. And it's connected to who we are as people. People are explorers. We're adventurers. We want novelty. We want to go to places that no one's ever gone, either literally or figuratively. And that is the essence of the human spirit to me. And we want to be advocating for that as technologists and leaders in our fields. Yeah. And it's so paradoxical. I remember when I was a kid, all these great musicians who I loved, Bob Dylan and et cetera, did the No Nukes concerts. And we really were indoctrinated into this fear of nuclear. And the second order effect is that we burned more coal and we burned more oil and we heated up the planet. And now we're trying to solve the problem. And the solution was there in the 70s. And then sometime in the 80s, we decided, hey, let's stop doing this. And now 80s, 90s, 2000s, 2010s, we're sitting here four decades later. And finally, people are starting to realize 40 years later, oh, you know what? Maybe that was a mistake. Should we start building these again? And now we've got to convince everybody that we went on a 50 year side quest that made no sense. And it's incredibly frustrating. And, you know, it's, yeah, to some of our friends, people have been on the pod, Sam Altman, Reed Hoffman, Mustafa, like, they, I think they're misguided here. We, we could have conversations about this, right? I mean, there's nothing wrong with having a conversation. Hey, how do you make nuclear safer? Hey, could these robots, I mean, it sounds farcical, but could the robots escape and then do bad things in the world? Sure, we could have this conversation. But that doesn't mean that we need to have a bunch of regulators come in and say, oh, somebody in Washington is going to approve your language model and your code. And that doesn't make much sense to me. That seems like they're doing regulatory capture, I agree 100%. And you know, the other thing I realized about what we were saying, Jordy, is there's something about solving

## [Transcript] This Week in Startups / Reverse-engineering autonomy in humanoid robots with Sanctuary AI CEO Geordie Rose | E1832

problems, I realized in this conversation, that when, and when we were talking about jobs, and there's a sensitivity to that, with good reason, we, you know, we want to pay things and a large amount of jobs could go away quickly. And there could be displacement, of course. But when the mind

and consciousness is left alone, our minds are designed in a very interesting way to think and find the next problem to solve. There's something fundamental about human consciousness and this brain and, you know, Darwin and evolution that our species survived, dominated, and evolved

with something inherent in our code, which is understand the world and find the next problem to solve. Does that resonate with you? Oh, yeah. I mean, everybody who's laid awake at night and they can't get their mind to stop spinning through all the negative scenarios that could happen. Everybody, I think, experiences this. You're exactly right, is that this tool that we've got, this beautiful mind that does all these wonderful things, it creates the worst nightmare as possible about what will happen as a consequence of it working well. And so with technology, our mind spins up all of these horror science fantasy ideas, we turn them into movies like Terminator or Black Mirror. None of that is real. I think there's a very important, powerful message here, is that the terrible stories our minds tell us when you lie awake at night about your personal life is the same process that generates fear about the outcomes of change. So when we do something new, we innovate, we discover something about the world, there's a natural tendency that all of us have to imagine what might go wrong. Yeah, so I would advocate for being aware of that, is that it's a story your mind is telling you. The Terminator thing is not true. It's not real. It's never going to happen. It's just a story that somebody made up that resonates with our basal, base nature, fears and concerns about the future and so on. But it's not real. What's real is very different. Yeah, in our minds, there was a reason this obviously existed, the person who worried, hey, I wonder if these berries are poisonous or not, or I wonder if there's something dangerous in that body of water, maybe I should be cautious. Yeah, a little bit of caution, thoughtfulness, probably extended life and people who were reckless probably had shorter lives. And so, yeah, the gene pool probably evolved this way. But you must be aware of how catastrophizing it is. I mean, people can get really wound up. We see this with social media presenting us with so much bad news in the world. Our brains are not designed to process that, are they? No, and this is an example of how technology can have unintended consequences that are negative.

Social media hijacks this propensity that we have to tribalize, to fear, to other people. Other people is being different. What part of this idea of thinking of this conscious perspective that you have is separate from your brain, carries with it another idea that we're all connected. We all have this thing, we all share in it. The analogy that Eckhart Tolle uses is that there's an ocean and we're ripples on the ocean, but this ocean is the same for all of us. This idea is a powerful one when you're trying to think about why you're reacting in a certain way to certain things. The social media stuff is an amplifier of the negative aspects of how we function as people. But that doesn't mean that we shouldn't have done it. I think this is the point. Like you said before, we want to talk about it. We want to have a frank discussion about it, but the solution to these things doesn't come from shutting things down. It comes from having this discussion and making good clear-minded decisions about how to build, not how to break.

**[Transcript] This Week in Startups / Reverse-engineering autonomy in humanoid robots with Sanctuary AI CEO Geordie Rose | E1832**

One of the great paradoxes of all of this might be we build up this AI and we get to some general intelligence. It might tell us, it's a non-zero chance, it might explain things to us about our own consciousness, why we're here, and what consciousness is that we ourselves could not come to the answer. We may unlock some mysteries that explain our own existence in a way. That is, just to me, would be a wonderful gift of accelerating this. What if this machine, what if this artificial intelligence can be more objective about us and can teach us something? That would be a pretty mind-blowing outcome, I sure would.

All right. Listen, continue success with this from just working on quantum computing and now to robotics and figuring out how to make the sequences play. It's going to be very interesting to watch your progress and listen, accelerate it all. Let's go. I'm assuming you're hiring and this must be one of the most fascinating places to work in the world. If people are interested in learning more or maybe applying for a position to build this out and accelerate human intelligence and augment it so beautifully, we're going to find out more.

So, I and one of the other founders of the company, Dr. Suzanne Gildert, have a podcast called the Sanctuary Ground Truth Podcast. That's a place that you could look. We also, at our website, [sanctuary.ai](https://sanctuary.ai), there is a careers page. We are hiring and growing quite quickly, and there are positions for all sorts of different kinds of people. We mostly hire technical people, of course, but there are some other things. And if anybody's interested, please watch the Ground Truth Podcast and go to the website and check us out.

Amazing. All right. And we'll see you all next time on This Week in Starter. Bye-bye.