AI and You

Transcript

Guest: Anil Seth, part 2

Episode 138

First Aired: Monday, February 6, 2023

Hello, and welcome to episode 138! Today we will conclude the interview with **Anil Seth**, whose TED talk, "Your brain hallucinates your conscious reality" has over 13 million views. Anil is Professor of Cognitive and Computational Neuroscience at the University of Sussex in the UK, has a PhD in AI, and authored the bestselling book, *Being You: A New Science of Consciousness*, which was a 2021 Book of the Year for *The Guardian*, *The Economist*, *The New Statesman*, *Bloomberg Business*, and a 2021 Science Book of the Year for *The Guardian* and *The Financial Times*. He has published over 200 research papers and has even rapped about consciousness with Baba Brinkman.

Last week we talked about how Anil got into this field, the wakefulness aspect to consciousness, Nagle's definition of consciousness as what it feels like to be something, the Garland Test vs the Turing Test, and… philosophical zombies.

One day this may come to an actual court case, and you'll hear me refer to Martine Rothblatt in the interview. In 2003, Martine, a lawyer and entrepreneur and executive produced of the documentary film of Ray Kurzweil's book *The Singularity is Near,* took a hypothetical lawsuit to mock trial. The case was titled BINA48 vs Exabit Corporation. Exabit was a fictional corporation that had invented an AI called BINA48, which declared itself to be conscious and through Rothblatt, filed for injunctive relief to prevent Exabit from turning it off. I'll remind you this was a mock trial; the purpose of it was to determine what would happen if such a case went to court. We may not be far away from such an event, given that Blake Lemoine asserted that Google's LaMDA had asked for an attorney and said it did not want to be switched off. So it may be instructive to look to the mock trial. I don't know if you can cite a mock trial as precedent. The outcome was that the jury voted 5 to 1 in favor of BINA48, but the judge set the decision aside because he thought the court had no authority to act without the legislature. In 2005, the California Supreme Court denied the appeal. It gets complicated after that. Anyway, it was an excellent device for focusing on the issues, and so I referenced it in the interview.

It's time to get back into the interview with Anil Seth.

So, I want to get into how to test for consciousness versus intelligence. But this is all surrounding this idea of is it going to be possible to tell the difference from the outside? I want to ask; if there's no way to tell whether something is a philosophical zombie from the outside, going to the theoretical definition of a philosophical zombie from the outside, going by the theoretical definition. Going by definition, I couldn't tell if you were a philosophical zombie. How do I know that *I'm* not a philosophical zombie?

Well, then you're really reaching sort of a far corner of this philosophical adventure. I mean, you might not be conscious in the way you think you are. This is true. So, there are some people who would say that we have some deeply-set intuitions about what consciousness entails that don't

even apply to us. Like if we think there's some sort of special consciousness stuff - qualia, as philosophers say - in the brain, then there are people like Daniel Dennett, who will say, "No we're mistaken about that. All we are is a bundle of dispositions to react to things in particular ways and this is why we imbue mental states with this peculiar consciousness stuff, that doesn't really exist." Now, I don't go all the way there; I think consciousness is real. And in fact, I tend to think that it's the only thing we can be really sure of, in this sort of, twist on Descartes and other people who said this, that the fact of consciousness is the only thing that I can be sure of; everything else is a little bit up for grabs. Doesn't mean I'm conscious in the way I might *a priori* think I am. But the basic fact of being aware is a starting point, and require a lot of evidence to convince me otherwise. But you're right, that it's just looking at the behavior of system from the outside is not very satisfactory, and in practice is not what people would do anyway. I mean, we would look at what's going on the inside too. And here there are a number of theories of consciousness, which make different claims about the necessary and sufficient properties for a system on the inside to be conscious. And one of the big fault lines in this debate, which is relevant to AI, is this thorny question of substrate dependence or substrate independence. And it's a bit jargony, again, but the basic idea is very simple. Does consciousness depend on the stuff something is made off? Now, some people are very confident that it doesn't. But consciousness could arise in a system made out of tin cans wired up the right way. If the system was constructed with sufficient complexity, in the right way, with the right kind of causal interactions, and so on. This is the perspective of functionalism and philosophy there, that it's the substrate doesn't matter, it's just what it does that matters. What it does on the inside and on the outside. The other position is that the substrate *does* matter and that consciousness is the kind of thing that only certain systems that are made of a certain kind of stuff can have. And I mean, we know that we're carbon-based life forms, and we're living systems, the only conscious systems for which there's any degree of consensus, that they're conscious are all living systems. There's no full consensus, we still argue about whether a lot of other living systems are conscious or not. Now, this really matters because if consciousness *is* substrate dependent, if it depends on some way on our nature is living organisms, then however persuasive a system made out of silicon is that's sort of implementing functions but with some very different underlying mechanism, it won't be conscious. And what's the answer to this? How do we know? Well, I don't think we do know. I'm suspicious of people who make very strong claims either way. There is as far as I can see, no knockdown reason to believe either that consciousness does depend on the stuff the system is made of, or that it doesn't. I think most a lot of possibly the majority of people that I know in the field lean towards the substrate doesn't matter, because they can't think of a reason why it should. But I actually tend to lean the other way and it's only a tendency. It's not that there's a knockdown argument for this. But I tend to think the burden of proof is to show that it doesn't matter, because all the systems that we know of that are conscious are living systems. There is good theoretical reason to tie consciousness to a nature as living systems and this is, of course, the argument that I make in the book that we are conscious it because of our nature as living creatures, and there's also the very tricky question of, where does the substrate start? If you're look in a computer, these sort of standard computers, it's pretty clear you've got your hardware and you got your software. But if you look inside a brain, and you ask, well, where does the mind where start and the wetware stop? It's actually a very

difficult question to answer. Every time a neuron fires, the structure of the brain changes. individual neurons are sort of little adaptive organisms, in some sense, they're trying to preserve their own balance of activity. And this drive to stay alive that we organisms have goes all the way right down to individual cells, and even within cells. So, given these tight links between consciousness and life, and the difficulty of thinking where the substrate starts, it's tempting to think that it's not so much consciousness as a matter of information processing of some sort - which is a term people slap around with I think, quite disreputable facility, it's like, "what do you mean by inflammation and by processing?" it might be life that sort of breathes fire into the equations of consciousness.

And there's this pervasive thought it seems almost a primitive subconscious belief that our mind is the software and our brain is the hardware. And I only bring that up to give it any sort of credence, because we've said that consciousness is something that we have this intuitive feeling for, and that seems to be a powerful intuition. What you've done there is connect some deep aspects of the hardware, like the microcode of our cells with a higher-level psychology, like survival instinct, and made that connection, so says one of those, that feeling, that drive comes from things that are encoded at a cellular level. And if a machine could be conscious, then we would have substrate independence, because software can be moved from one machine to another.

Yeah, if we're talking about machines as we currently have them.

If we're talking about the same kind of machines, and whereas this is much more debatable when it comes to humans, but is it possible that a machine could still be conscious even though human consciousness *couldn't* be substrate independent?

Right. "I don't know" is the answer to that. I mean, again, my intuition is that all consciousness is related fundamentally, to our nature as living creatures. But this is not a view that I hold with a high degree of credence or confidence because I think the jury is out. I mean, we just don't have a consensus theory about the sufficient mechanisms for consciousness, I just think it is unwise to assume that the substrate doesn't matter because for the reasons we've already explained, it it could indeed be that there's a middle ground, which is that consciousness that is illuminated by the kinds of emotions and valence that human and animal consciousness is characterized by things feel good or bad, that might be fundamentally tied to life. But there may be other kinds of consciousness that turn out to be less or not substrate dependent, that have different characters. So, perhaps, there would be a form of machine consciousness, which would be a kind of awareness that is totally devoid of any emotional content; kind of *Star Trek* Commander Data type of thing. Again, who's to say? It's just from what we do know and certainly the way I've been thinking, emotional experiences can be very, I think, usefully understood as a rising from the brain's predictions about how well the body is doing at staying alive. Now that is something then that at minimum, we can only understand the nature of emotional experience through the lens of understanding our nature's living systems and perhaps a stronger claim applies to that we can't have those experiences unless we are living systems. And the rebuttal to this is often that, "Well, you can just simulate all these things". You can have simulated emotions. And yes, of

course you can. But this doesn't always mean that you instantiate the phenomenon in question. There are some things for which if you simulate them, well, that's enough. If you simulate chess playing, you actually are chess playing. One of my colleagues likes to say you're actually playing the history of chess, that's what most chess computers do. But there are other things for which if you simulate them, it remains just a simulation. And there are very non-controversial examples of this. If you simulate a weather system, as people are very good at doing now, nobody even thinks for a minute that it actually gets wet and windy inside the computer that's simulating a storm. It's purely a simulation. So, this isn't to say that consciousness has to be one or the other, it's just to say that we don't really know which it will be and to make strong assumptions either way is dangerous. Now, I would be very worried if consciousness turned out to be substrate independent, because I think the, the possibility of easily building a conscious machine is something that is very ethically problematic indeed, because as soon as we build something that has conscious experiences, we have ethical obligations towards it. And in the case of a machine, this would be very hard to tune, because we would have no particular idea what those experiences might be like, or perhaps even whether they were going on at all.

And we'll get to that, but just wanted to go back to the simulation, that a simulation of weather is not getting wet, and the counter example is that a simulation of a chess game is a chess game and that says that a chess game is actually software, it's actually data. It is substrate independent by definition that you could represent it as the description, the list of all the moves that were made and is not the same as moving a particular piece on a particular board and so that the chess game there is a software. Does that illuminate that distinction?

Yeah, no, I think you're absolutely right and I think there's another perhaps useful perspective on that, is back to this distinction between consciousness and intelligence. Here, I think we can have much clearer intuitions that intelligence in general, is very likely substrate independent. We're already getting there. We haven't got to general AI yet, the general capability and soon as we get there, of course, we'll just go right by it. But it seems that we've enough existence proofs of systems that display forms of narrow AI, that intelligence in general doesn't seem to be something that is substrate dependent. Equivalently, there's no really strong reason to think that it ought to be in the same way that I think that there are some arguments why consciousness should be substrate dependent. But back to your other distinction, it may be that the particular human way of being intelligent, still is substrate dependent in some way, it might still require the kind of emotional input the kind interactions that we have via our body and environment that give human beings the particular intelligence that we have. But intelligence in general, I think can be substrate independent, but maybe not consciousness.

Well, then let's go for the brass ring here and imagine some point in the future a situation like LaMDA, Blake Lemoine, one, it's a lot more contentious now, it's in a courtroom and this has actually been played out in a mock trial some years ago by Martine Rothblatt, if anyone wants to look up the history of that. But let's say there's something that an AI that one group of people is asserting is conscious and another group isn't and it has come to a court and you are called as an expert witness for either side. But what is your argument? Dr. Seth, you are on the stand, how should we tell whether this thing is conscious?

Well, I firstly I try and have the people who designed the machine arrested for negligent use of technology because I do think we shouldn't be pursuing that goal of even trying to build conscious machines, it's just not a cool or useful thing to do, because they might have experiences and that would be ethically catastrophic. But alright, so we're there, we're in the courtroom. If this was happening right now it would be a very difficult position. What I would say would be, okay, it is not a question of how similar the system is to you or me on which we should base this decision. It's not a question about how intelligent the system seems to you or me on which we should base this decision, we need to base it on our best understanding of the necessary and sufficient mechanisms for consciousness. And this is an evolving set of criteria. There are currently a number of theories. So, I would want to know of the designers of the system, how does their AI stack up against all the different theories that we have? Each of them makes certain sort of predictions of the signatures of the markers that we might associate with human or animal consciousness, both in terms of their behavior, but also in terms of the mechanism, in terms of the stuff inside. And so you'd want to collect all those data points, and see on how many different theories would the claim stand up, that this system is conscious? And then what would you do with the answer? Now, that's a decision that escapes that particular courtroom, and that's a decision for society as a whole? You know, what are the criteria that we settle on as meaningful? And this is why I think these kinds of conversations are very important, because we've seen how much mess can be made with LaMDA and those kinds of those kinds of hubristic media disasters, people are very attracted to this idea. And again, science fiction has treated this very well like the science fiction series *Westworld;* done this brilliantly by examining the effects on human psychology, that inhabiting a world where there are AIs that give the strong appearance of being conscious, are present and we're encouraged to treat them as if they're not. This turns people into sort of instant psychopaths, it's not a place you want to be. So, the decision is much broader. It's not just like, "Let's find the criterion and go yes or no", there are consequences to the criteria that we apply for how society will function in general. So, it's a case of taking what we know; there will be gray areas, and I don't think it's going to be in the near term, that we get a satisfactory resolution to the substrate dependence question and this is still this is still in my mind very much up in the air. But if I were the expert witness, and the only thing holding me back was the fact that the substrate is not the same. Like if all the internal other mechanisms were lined up with the most compelling theories of the time. No, then even for me, I would have to say, "Okay, look, I don't really know whether it matters or not. So, we have to obey the Precautionary Principle here and be on the side of caution." But I would just not want society to get into that position.

We may not have that choice.

That's right.

At what point would this decision be elevated above prior decisions regarding animal rights? Would it be the degree of intelligence that - I know we've said that then it's not related to consciousness, but then what would you say that would tell the court that they shouldn't just refer to some previous decision regarding animal welfare and animal rights?

Well, firstly, intelligence is related to consciousness. It's just not the same thing. So, and it's related in in a way that's relevant to this conversation because one aspect of intelligence is mental time travel, for instance: the ability to imagine possible futures and remember the past and it's because of our intelligence that we are capable of having experiences like regret rather than just being sad, because we can imagine alternative counterfactual possibilities and usually do. So an animal welfare decision - this is a little out of my wheelhouse, but my understanding is they would ideally don't he made on the capacity for suffering, and one's level of intelligence changes the capacity for suffering. And I think this does play out, for instance in how animals are housed in farms and in zoos and things like that, you know, caters to a little bit about what not just a natural habitat would be, but what the cognitive capabilities of the creature are that might lead them to suffer in certain environments rather than others. But it's very different because when we're faced with this hypothetical courtroom situation, one scenario where this is perhaps most likely to reach the courtroom is where the AI engages in language with us - I mean, this is something that even AI today is getting quite good at, as we've discussed. And it's going to play on our anthropomorphic tendencies very strongly. And that's something that there's always this tension, thinking making attributions of consciousness beyond the human between our tendency to attribute properties of the human consciousness and intelligence to other systems, anthropomorphism; and anthropocentrism, this tendency to think that we're special, and have properties that that are reserved for us. And navigating that tension is really the hardest thing and I think it does play out differently for animals versus machines, because machines are likely to play on our cognitive empathies much more than our more basic emotional and physiological resonances and then there are other systems I mean, this courtroom could become quite busy because we've got animals, we've got like bumblebees coming in in courtroom one, we've got amoeba down the corridor, we've obviously we've dealt with all the mammals, they're fine, they're aware, they're conscious, we've got whole schools of fish coming in, and courtroom three around the corner and then in some other aspects. And in other rooms in this courthouse, we've got brain organoids, which, which are these emerging neuro technologies, which are sort of lab grown mini brains. Now, here's a really peculiar thing, because the current organoids are not very smart, but they're made out of the same stuff. So, the whole ambiguity or about substrate independence is out of the window. So, we've got now sort of more complex organoids rolling around in there, now maybe aquariums or something with cameras attached. And then probably the most contentious courtroom of all, we have human fetuses and we have the question of okay, at what point does consciousness happen? Is it the moment of birth? Is it before? Or bizarrely, it might even be slightly after?

Right. Right.

All of these things, just, I think, really spotlight why a science of consciousness, even if it doesn't have all the answers, is really essential. A voice needs to be in the room for all of these kinds of debates. Not to determine exactly what happens. But as you say, I think it's a really nice context to be an expert witness to give the best-informed perspective, based on what we know.

And to be asking those questions. And there are people doing research now, as you mentioned, where they're binding human neurons to chips, to make this kind of hybrid intelligence that

uses the neurons. So, as you say, that would definitely blur the line about the substrate independence. You wrote in the book about talking to GPT-3, and finding the edges of it. In the last two months, OpenAI has come out with ChatGPT, which is a leap ahead of what they had before. Have you interacted with it? And would you revise your opinion in any way?

I have interacted with it and I have not revised my opinion in any way at all, it seems to be very good at writing haikus. I mean, it's fun to play with, you can get it to write a haiku about something and then you can ask it to do it again, but more sad and it does it again, more sadly, it's really impressive. I mean, I actually think this is important because to the extent that this argument about artificial consciousness is taking center stage, it's really distracting attention from the progress that these systems are making: ChatGPT compared to GPT-3, and compared to these systems as they were five years ago, I mean, the progress is phenomenal and we shouldn't lose sight of that just by continuing to make drama about whether it's conscious or not. But yeah, ChatGPT was fun to play with. But it was still very easy to trip up. I found it quite unsatisfying, actually, because it sorts of combined a certain fluency with a total lack of substance. You know, It seemed really boring. Apart from, it could do party tricks. But I think it's a big problem for education, because you can probably get it to write a passable essay about anything at all, even a possible abstract for a scientific paper, which isn't, you know, saying too much about the quality of most papers that get published. But no, I found it again, not sorts compelling that there was a mind a conscious mind or even an intelligent mind behind it. I think part of it is that OpenAI has put certain guardrails there that there are certain questions that that you might ask as someone interested in consciousness that it's steering clear of. So, I think wisely, so that it doesn't get caught up in the same kind of hubris that that LaMDA did. But no, the haikus that it wrote were, yeah, alright. Well, I think I asked it to write one about snorkeling with manatees in Mexico, because I just been snorkeling with manatees in Mexico. It certainly did a better job of writing a haiku about that than I could have done.

I've asked it to write haikus as well.

Why do people do that? And that seems to be the thing that everyone's trying to do. Because again, it because it exemplifies the party trickiness of it. It's not it's not a serious conversational partner.

There have also been conversations, it has a much larger context and previous models. Again, not imputing it's crossed some philosophical line yet, but it is certainly grabbed attention in a way that others haven't. And I think I also have the experience of it as kind of boring. But I think this is also because I know what I'm talking to and I know it's not going to originate - well at least I don't believe it's going to originate ideas that haven't existed on the internet somewhere already. But I think in a sense, that is actually unfair; that if I didn't know that I was talking to a machine, I might have a more satisfying conversation. I don't know quite how to make that experiment happen, though.

I think you're right to point that out. I think you're absolutely right to point that out and that gets to this sort of the four different questions we were talking about; that that for different humans, it will take a different amount of evidence for that person to believe that the system that is either

conscious or intelligent, independently of whether it actually is. The inverse Turing Test and inverse Garland Test are going to be relative to the particular human on the other end.

This is fascinating. I really appreciated the tension you drew between our tendency to anthropomorphize and find human qualities in everything from dogs to Tamagotchi, and at the same time, be anthropocentric and say nothing but humans can have these qualities that we want to reserve to ourselves, and I think that's sounds like an entire book idea there. We are running out of time, I wanted to ask one penultimate question. That is, there are people that are highly invested in the concept of substrate independence; Ray Kurzweil, for instance, made no secret about his goal to be eventually uploaded to a machine. Do you think these people have any hope?

Oh, I think, unfortunately, very little hope. I mean, I think this is again, it's kind of problematic, really, it adds to this sort of slightly distasteful, ethical atmosphere in these conversations about machine consciousness because they are all tied up with this sort of transhumanist uploading culture as well. Like you know, it's very tempting if you're working in this area, maybe to think of yourself as pivotal to this moment in human transition and that what's to be gained is not merely being some kind of God where you create a conscious organism, but you gain a prize of immortality too because you can upload yourself and live forever in pristine circuits of some future neuro computer. And here the personal motivations get so mixed up with any kind of they really occlude, I think, any philosophical clarity, and actually any realism about the systems that we have that we're interacting with. Yeah, even if consciousness can be substrate independent, the idea that it's a plausible near future technology, I think it's just nuts. I think that in that area, I think there's a lot more to be worried about, about life extension technology, I think those are much more plausible and of course, all the billionaires are investing tons of money in those, and that's going to really, I mean, we've got so much inequality already. But the great leveler, that everybody dies within plus or minus x years; [if] that goes out the window, then we're in real trouble.

Thank you. So, here you are running a lab researching consciousness. What would you, if you dream for that to have achieved ten years from now?

Sustained funding would be nice. That's always a dream. But in terms of achievement, I think to have to be realistic about it. To be realistic, I would like to have made substantial progress in thinking of some of the questions that I have - not necessarily about the big question of how does consciousness happen at all, but I would like to be clear about the different theories that we have. Okay, we know, for instance, that I didn't know global workspace theory, it seemed good at the time. But now we know that's not a very good explanation. This, by the way, is a theory that associates consciousness with global broadcast information in the brain. It's a very nice theory, but we have a set of theories, I would like to see those defined narrowed down; the theories that we have to be more specific, carry more explanatory power. I would love it if we were in a position to be much more confident in answering the kinds of questions that we've been discussing today. To be much clearer about what the criteria are for consciousness to exist in a system, and what the criteria are for us to attribute consciousness to a system.

Do you think at that time, we could get more unification or bridge between the philosophy of consciousness and the neuroscience? Draw the line between your definitions and brain activity?

I think that's essential. I think one of the attractive things about this area - one of the things that attracted me certainly - always was the internet links between philosophy and neuroscience and psychology and computer science. It's a fundamentally interdisciplinary area, you're going to need these different perspectives to be not just working on the same problem, but working together on the problem, in order to make progress. So, yeah, I want to see that deepen and continue, and I'm optimistic about that. There's been some brilliant philosophers who've inspired me over the years like Dan Dennett and Patricia Churchland, and Thomas Metzinger and now there are new generations of philosophers who also are trained in maths and trained in biology and trained in neuroscience. And it's when people that sort of grew up with this multidisciplinary training that I think the potential is really exciting. The other thing, by the way, that I would like to see in ten years is consciousness science making real difference in psychiatry. It's already making a difference in neurology, where some of the tests of consciousness in humans are being used, in some cases to inform decisions about diagnosis, prognosis in patients who appear to have lost consciousness entirely. But then if we look at the whole realm of mental health, and psychiatry, psychiatric disorders generally present as disturbances in conscious experience. Things come to our consciousness differently, sometimes in behavior as well and explicit behavior. But fundamentally, it's about emotional experiences being distressing, it's about hallucinations and delusions and things like that. And most psychiatric medicines - so pharmaceutical approaches - tend to just be about addressing the symptoms, rather than treating the causes. We don't have an antibiotic or an antiviral for psychiatry, we just have things that we have things like aspirin or paracetamol, to turn down symptoms. So, an understanding of how the brain generates particular kinds of conscious perceptions in situations - this is something we're working on in my lab - what are the what are the precise neural mechanisms that that explain different kinds of hallucinations that can come about in different conditions, Parkinson's disease, schizophrenia, psychedelics, and so on. Yeah, drilling down to that level of mechanism, I think gives us new ways to think about treatment, and interventions, and really do something for mental health and psychiatry. I think that is very possible. That's something I'd love to see happen.

Fascinating and let's hope the funding shows up, then. Anil, thank you so much. We have run out of time, there will be a link to your TED Talk and *Being You* the book in the show notes and the transcript. And what would you like to tell people who would like to know more about what you're doing, where they could find that maybe if they want to come and work with you, or other things that are relevant to people wanting to find out more about you and your work?

Thanks for the opportunity. Well, of course, the easiest place to start is with the TED Talk and the book *Being You* and the other thing that I'm working on right now, that would be wonderful, if your listeners would math's engage with it's a thing called the Perception Census, you can find it on my webpage, which is AnilSeth.com and it's a big online citizen science study of perceptual diversity, one of the implications of my understanding of consciousness is based on the brain's

predictions is that we all inhabit a slightly different subjective world. Even if we're faced with the same objective world, you know, our brains bring a lot to the party and we all have different brains. But we know very little about inner diversity, because it's just private and subjective. You know, we can see that we're different skin colors, and heights and body shapes, and so on. But unless inner differences get sufficiently dramatic, that we start behaving differently, and we slap labels on neurodivergent conditions, we'll never know, like is the blue sky that I see the same blue, the blue that you see? And so, this project is setting out to measure this hidden landscape of perceptual diversity and we've had more than 20,000 people take part already from 100 countries, but I'm really hoping to more than double that. So they're fun, illuminating informative brain teasers and interactive illusions, and so on and, it's called the Perception Census. You can find it on https://www.anilseth.com/. That would be wonderful and the last thing I'll say is that actually, a lot of my lab is focused on the productive interactions between AI and consciousness research, decides all these questions about whether machines can be conscious or not, there's a huge opportunity to use techniques from machine learning and AI as models of how brains accomplish tasks in perception, and in action in cognition. To address this, what I call the Real Problem, to try and map between mechanism and phenomenology and so that's a real focus for me going forward is, is so, developing or finding these productive synergies between cutting edge AI research using generative modeling and things like this and but using them as computational implementations of theories about the about conscious perception. So, that's something our lab is working on and if that's of interest to people listening to this podcast, yeah, keep in touch about that. That would be great.

Oh, fantastic. Thank you. We've answered many questions and left even more open for future research and debate. But thank you so much, Anil Seth for coming on *AI and You*.

Thank you, Peter. It's been a real pleasure.

That's the end of the interview. I hope you've enjoyed this fascinating romp through the field of consciousness with a top international expert. I was particularly interested in what we discussed about substrate independence and the kinds of holes that are being poked at that in arguments about human consciousness. It reminded me of how we have nerve cells in our intestines that are connected to our brains, so we are doing part of our thinking with our digestive system. Think about that the next time you say, "My gut is telling me…" If you enjoyed the episode, please take a moment to give us five stars and a review, it makes a huge difference to how many people get to find out about the show.

In today's news ripped from the headlines about AI, the question-and-answer website Quora launched a platform called Poe, standing for Platform for Open Exploration, an iOS app that that lets people interact with a number of different AI agents, including, of course, ChatGPT. A spokeman said it "is designed to be the best way for someone to get an instant answer to any question they have, using natural conversation." So you can see here that the large language models are really coming of age in their ability to be of general use.

The Perception Census that Anil referenced is research into "perceptual diversity." Just as we all differ on the outside, we all differ on the inside too, in how we perceive the world. The study involves a series of fun, interactive illusions, brain teasers, and games playing with color, illusion, sound, and more. Participation is easy - all you need is your own computer - and by joining in you'll help advance the

research, and learn more about your own powers of perception too. There's a link to that and to Anil's book, *Being You*, in the show notes and transcript.

Next week, my guest will be Risto Uuk, a policy researcher at the Future of Life Institute and expert on the European Union AI Act, which we will be talking about. That's next week, on *AI and You.*

Until then, remember: no matter how much computers learn how to do, it's how we come together as *humans* that matters.

http://aiandyou.net