



MIND & LIFE

Mind & Life Podcast Transcript

Anil Seth - How Our Minds Predict Our Reality

Original Air Date: May 13, 2020

Retrieved from: <https://podcast.mindandlife.org/anil-seth/>

Opening Quote – Anil Seth (00:00:03): I do see this close alignment between the insights that we're getting from cognitive neuroscience and Buddhism, to the extent that things are not necessarily the way they seem, there is an impermanence to our experiences, there's an impermanence to the self. And recognizing that impermanence does open a space for change.

Intro – Wendy Hasenkamp (00:00:31): Welcome to Mind & Life. I'm Wendy Hasenkamp. This week I'm speaking with Anil Seth. Anil is a cognitive scientist who studies how consciousness works. As you'll hear, he started his career being interested in how we perceive the world around us, but that line of research has taken him far beyond the realm of basic perception, touching on the nature of our emotions, our sense of self, and even the question of free will.

(00:00:58) I caught up with Anil at last year's Mind & Life Summer Research Institute. This is a really unique event. It's a week-long gathering, held each June in upstate New York — part academic conference, part meditation retreat. There's about 150 participants, and we all stay together at the Garrison Institute, overlooking the Hudson River. And as part of the contemplative aspect of that program, portions of the event are held in silence, and you'll hear us reference that dynamic at one point in the conversation.

(00:01:28) The perspective that Anil's research has helped advance centers on the idea that our brains, instead of merely sensing the world and responding to it, are also very much involved in prediction and inference. And that has some surprising implications for pretty much all of our subjective experience.

(00:01:47) In our conversation, we discuss these views of the brain and mind, how they impact our sense of self and how we relate to others, possibilities for change, the overlap of these ideas with some key concepts in Buddhist theory, and we also touch on the central role of the body in all these processes and how that relates to our experience of emotion. If you're interested in learning more about Anil's work, we've linked to some additional resources in the show notes for this episode, including his Ted Talk, which has been viewed over nine million times. So I really hope you enjoy this conversation. It's my pleasure to bring you Anil Seth.

Wendy Hasenkamp (00:02:30): Okay. I'm here with Anil Seth. Thank you so much for joining us today on the podcast.

Anil Seth ([00:02:34](#)): Thank you.

Wendy Hasenkamp ([00:02:35](#)): I would love to begin just hearing a little bit about the roots of your interest in consciousness and how we perceive the world.

Anil Seth ([00:02:42](#)): I think like many people, I've been interested in these big questions like consciousness for as long as I can remember. I remember arguing with my friends in the pub in England, or before pub, while at school. And it's just been one of those animating questions. And for some reason, and somehow I've just been fortunate enough to still be allowed to be interested in it, and work on it for a career.

([00:03:07](#)) At the time I was a student in the mid '90s, it wasn't really an obvious thing that you could embark on a scientific research career to study. It was a topic in philosophy, maybe a topic in religious studies, even. And a fringe topic in psychology and neuroscience at best. But it still seemed to me the most interesting phenomenon. How does it happen that we have conscious experiences? That there is something it is like to be me, or to be you, or to be any conscious creature?

([00:03:40](#)) There was no particular good answer. And I remember thinking to start with, that maybe physics is the right way to understand this. It's a big mystery. It's a big gap in the whole structure of our understanding of nature. And so, I started out studying physics. And after some time, I kind of reached my limit at the ability to solve difficult equations, but also it just became evident to me that actually there are real opportunities in studying psychology and neuroscience to address this question. And some of the explanations of consciousness that were coming up from the physics direction were not making any sense. So these were sort of explanations like, "Well, consciousness is mysterious and quantum mechanics is mysterious, so they have to be linked." And this was kind of dissatisfying, to put it mildly. It was really rather off-putting. So it was with some relief that I kind of transitioned into psychology and neuroscience. And then things changed, I think very generally across the landscape of the brain and cognitive sciences, and consciousness became — it's not exactly the most legitimate and well-established branch, but you can teach it, you can research it. And most of all, you can occasionally even get grant funding to study it.

([00:05:01](#)) So basically that's what I did. I think one important step I wanted to mention was my postdoc in the United States, in San Diego, and that was in the early 2000s. And at that time, it was still pretty uncommon. So coming to the States was a big transitional moment for me. My boss was Gerald Edelman, who was one of these figures he'd won the Nobel Prize. He was allowed to study consciousness. Francis Crick was just down the road. So San Diego at that time was wonderful. It was wonderful in many ways, but it was one of the few places where consciousness was properly on the menu. Stayed there six years, and then went back to Sussex University where I am now, still, and set up a group there.

Wendy Hasenkamp ([00:05:40](#)): Yeah. So you describe the brain as a "predictive machine." Can you say more about that?

Anil Seth ([00:05:45](#)): Yeah. There's so many metaphors that people bring to bear when they try to understand the brain. It's an enormously complex object, it defies any kind of intuition. It's 90-odd billion neurons, 1000 times more connections. We don't know the fine wiring diagram of the brain. We just know it's important for a lot of things. People have always reached for kind of dominant technology as the metaphor — a system of complex plumbing, a computer, and more recently, maybe the internet.

(00:06:17) I've always found the computer metaphor a bit off-putting, a bit misleading, it de-emphasizes consciousness. It de-emphasizes the role of the body, and the brain just is not a computer. Computers you can use to simulate all kinds of things, but the brain isn't a computer.

(00:06:37) The way I started to think about it — and it certainly was not original idea, as I later realized, it's built on a very long tradition in philosophy and in psychology — is to think from a starting point more of, what is the problem that the brain is trying to solve? It's like a functional perspective on it. And if we start with perception, so one of the main jobs of the brain is to perceive both the body and the world so that the body can act, the organism can act within the world. And the problem with perception, it's tempting to think in a lot of classical models of perception, think of it as this kind of bottom up process, where the sensory data comes in through the senses and is then read out by the self, which sits somewhere inside the skull.

(00:07:31) But that's a pretty unsatisfying view because the world that we perceive is only indirectly related to what, if anything, is actually out there in the world. There are no colors actually out there in the world, colors are something that's constructed by the brain. So I started to think... There was a literature already on perception as a predictive act. And I don't mean predicting the future. I mean, prediction in the sense of making a best guess about what caused sensory inputs. So prediction in this sense is more like filling in missing data or making an inference. And we get all the sensory data, but it doesn't come with labels about where it's from, or even what kind of sensory data it is. But we end up perceiving this world full of structured objects with particular properties.

(00:08:29) So the idea is that perception really depends on these top-down predictions, that the brain is always bringing to bear, to give shape and form to the sensory data that comes in. Now, this idea goes way back to Plato, even, in the shadows on the cave and people thinking those are real. And then there's an Arabian scholar in about the 11th century who started talking about perception inference. And then in the 19th century, the German scientist Herman Von Helmholtz was the first to really articulate this idea of perception as inference.

(00:09:04) And now this is an idea with quite substantial momentum. People call it predictive coding, predictive processing, the Bayesian brain. But the shared idea is that perceptual content, what we perceive, isn't a readout of sensory data. It's a top-down best guess, a projection, a hypothesis, a prediction about the causes of sensory signals. And that's what we perceive. It's the brain's best guess of what's out there.

Wendy Hasenkamp (00:09:33): It's interesting that this idea has been around for so long, in fact, in different ways. It seems really to be taking off now in cognitive science in a way that I think maybe it hasn't held sway before. Do you have a sense of why now people seem to be jumping on board with this idea?

Anil Seth (00:09:51): It's always hard to see why ideas gain a certain resonance at a particular time. I mean, part of it, I think, is that the various disciplines that constitute cognitive neuroscience have reached a stage where these ideas seem sensible. So for instance, we can start to look at direction of information flow in the brain. We're getting a little bit tired of cognitive neuroscience that just says it's this area or that area that's responsible for this process or that process. I think people want to know more about, well, what are the mechanisms? What's the circuitry? How does it work? Not just where is it?

(00:10:35) And also, I think there's been an increasing interest in the phenomenology of this. So if we look at things like psychosis or other forms of hallucination, it makes sense to think of these phenomena in terms of predictions, perceptual predictions that are not restrained by the sensory data, that

overwhelm the sensory data in some measurable way. And so I think progress can be made. In fact, there's a phrase which I've used quite a lot. I didn't make this phrase up at all. I heard it from one of my mentors, Chris Frith, in London. He describes perception as a "controlled hallucination."

Wendy Hasenkamp ([00:11:13](#)): Can you describe that?

Anil Seth ([00:11:14](#)): Well, I think it's just this idea that we typically think of hallucination as a false perception — when you see something that nobody else next to you sees, we call it a hallucination. And the assumption is that it's a totally different category from normal perception, which is revealing the world as it is ideally. But I think this phrase picks out a continuity between normal perception and hallucination.

([00:11:43](#)) They both result in us perceiving something, and they both depend on this interplay, this dance between the brain's predictions about what's there and the sensory data, which is giving feedback about these predictions, prediction errors, we'd call them computationally. And when people have hallucinations, what's happening, or the idea of what's happening is that the prediction errors, the sensory data, aren't updating the predictions in a way that happens for the rest of us.

([00:12:14](#)) So people's perceptual content becomes unmoored from the causes in the world, and people report perceiving things that other people don't. So hallucination is, if you like, a kind of uncontrolled perception. And then by the same token, perception is a controlled hallucination. And it's controlled not by the individual. I'm not saying like, "I am in control of my perceptions, or my hallucination," as like me (Anil Seth) decides to perceive things as they are. No, the control is coming from the sensory data. So my perceptual predictions are being controlled, being reined in by prediction errors that are reported from my eyes, my ears, my nose, my other senses.

Wendy Hasenkamp ([00:12:57](#)): Yeah. So what enables the brain or the mind to generate these predictions? Like where are they coming from?

Anil Seth ([00:13:05](#)): That's where I think a lot of work needs to be done. It's a nice idea. But one of the challenges is that it fits with a lot of evidence, but there's very little direct evidence that discriminates this idea from other ideas. To actually characterize, let's say, a visual sensory input as a prediction error, as reporting the difference between what the brain is expecting and getting, it's hard to come by that kind of evidence. So there's a big program of research that needs to be done. In my lab, where we're doing some of that, to try to lock down what's happening in the brain when predictions are being made and being updated.

([00:13:47](#)) Where do they come from in the first place? That's another interesting question. And again, it's a slightly dissatisfying answer, because there's going to be... it's a combination of factors. So you can imagine that certain perceptual predictions — we can call them, another word for them would be "priors." So, prior beliefs about the way the world is. And these aren't beliefs you're necessarily aware of having. It could be a belief your visual system has that light generally comes from above. Now this could be a belief encoded in your visual system through evolution, through development. It's just baked in there, pretty much anything that we think brains do — if you ask what gave rise to that, it's going to be a mixture of evolutionary factors, developmental factors, lifetime learning, and moment-to-moment experience.

Wendy Hasenkamp ([00:14:34](#)): So how do you see these ideas about the brain and the mind as implications for how we think about ourselves?

Anil Seth (00:14:42): I think this is where things get really interesting. And this was where this line of thought took me in a direction I wasn't necessarily expecting to go. I'd started out mainly interested in perception of the world and how that happened. And in that view, the self is the thing that does the perceiving, and then makes some decisions and performs some actions. And so that's the, sort of, typical way we think about it. That there's a world out there. We sense it. We read out perceptions in this outside-in direction, the self does the perceiving, and what we perceive is related to the world. If you think of perception as prediction, then it becomes natural, at least to me, to think that the self is not the thing that does the perceiving, the self is a perception.

Wendy Hasenkamp (00:15:30): How do you get to that?

Anil Seth (00:15:33): I think you get to it by thinking, what do experiences of self consist in? And then you realize that experiences of self are just forms of perception. And there are many different forms of perception that constitute our experience of being a self. There's the experience, the perception of, what in the world is my body compared to things that aren't part of my body. There's the experience of having a first-person perspective on the world. There are experiences of intending to do things, of what typically or colloquially is called "free will" experiences. Then there is experiences of memories, of a sense of person identity that extends over time.

(00:16:16) And all of these aspects of self, I think can be described as perceptions. And we have known for a long time, from both case studies in neurology and experiments, that these different aspects of self can fragment. People can lose their autobiographical memory. Part of their self is altered, but much of it remains. People's experience of what is their body or what is not their body can be altered in disease or in the lab, which shows that part of ourself is also constructed on the fly. It's not to be taken for granted. The temptation is just to take the self for granted — that I perceive myself this way, and therefore, there's nothing to explain. But I think there's a lot to explain.

(00:17:03) And it's particularly relevant when we think about psychiatric conditions, which can often be well described as altered perception of self. And if we can start to understand how we perceive the self according to the same principles that we perceive the world, it gives us a view into what does it mean to be a self. But it also gives us an avenue into thinking about how to understand psychiatric conditions, unusual aberrant perceptions, and so on.

Wendy Hasenkamp (00:17:34): And how about implications then for changing our perceptions, our beliefs, our sense of self, even? Does this view point to ways of creating change?

Anil Seth (00:17:47): I think it opens space for it. One thing about change, which I've been recently thinking about, which I find potentially fascinating, is again just taking some parallels from what we know about how perception works in the world, and then seeing if that gives any insight into how we perceive being a self.

(00:18:08) So there's this phenomenon of change blindness, which has been well studied in psychology. One expression of change blindness is if a scene changes very slowly, we typically don't perceive the change. Like the background color of a screen can change very slowly, and if somebody's focusing on something else, they might not notice that half their visual field has, or at least the input has, changed color. There's an interesting question there — what do they perceive? Is it that they perceive what they started out with, or has their perception changed, but they haven't noticed the change? I think it's the latter. These experiments show that the change of your perception is not the same thing as perception of change. Change itself is something you perceive, like you perceive color, like you perceive sound, like you perceive shape.

(00:18:53) And then if you think about how that applies to the self, one of the characteristics of how people experience being a self (unless they're in some sort of psychological crisis) is as stable. We perceive ourselves as unchanging over time, or rather we don't perceive ourselves as changing. So I think we have a kind of self-change blindness, which is quite adaptive. It's quite useful for us to perceive ourselves as not changing, because it gives us a stable anchor for our psychological identity. In fact, we'll come to it later, probably, but the point of perceiving things can often be to regulate or control them.

(00:19:29) And so I think we perceive ourselves the way we do in order to maintain ourselves as stable physiological entities. And so, that itself entails that we perceive ourselves as more stable over time than we actually are. Because if you think about it, we are... I'm not the same person I was 10 years ago or 20 years ago or will be-

Wendy Hasenkamp (00:19:48): Or even yesterday.

Anil Seth (00:19:50): Or even yesterday. But I don't ever notice that change. And I think I'm systematically wired to not notice the change in myself. And of course, that can break down in certain situations. And that's sort of one way into how these ideas can shed interesting angles on psychiatric conditions, mental illness, or other forms of... I'm focusing on those because they're more distressing, but any change in one's conscious experience, which would be described as unusual, I think it can be addressed this way.

Wendy Hasenkamp (00:20:24): So a couple of things coming up for me there. The way you described it — as the normal mode of operation is to perceive ourselves as very stable and unchanging over time — makes me think of foundational principles in Buddhism that suggest that that very tendency is, first of all an illusion, which aligns exactly with your work, but also one of the foundational sources of all of our suffering. Because we don't perceive the reality of the ways that we change and the ways that reality is interdependent on all the contextual factors and things that play in. Have you gone into Buddhist theory at all as it relates to your work?

Anil Seth (00:21:03): I've always been interested in it, and I admit that it's not something I've studied explicitly much. I read some of the stuff when I was young, and I've always been fascinated by the apparent overlap between (at least my limited understanding of) the Buddhist doctrine and Buddhist philosophy and Buddhist practice, and what modern neuroscience is teaching us.

(00:21:32) And I think you pick up a couple of the points that stood out to me as well. The idea of perception being a construction. We tend to experience our perception as being real. That's why it's so counterintuitive to think of it as a best guess or as a top-down prediction, because I really experienced the table as being red. And perception has to work that way, because if it didn't, I would be constantly wasting time thinking, is it really red? It works by giving me the appearance of reality.

(00:22:06) And in the same way, my experience of a stable self works because it gives me this experience that the self really is stable, and that's often useful. But of course, it's also, I think, as you just explained, it's the root of a lot of problems and potential suffering too. So I do see this close alignment between the insights that we're getting from cognitive neuroscience and Buddhism, to the extent that things are not necessarily the way they seem, things can change.

(00:22:37) There is an impermanence to our experiences, to the world. There's an impermanence to the self. And recognizing that impermanence, whether it's in a sort of first-person confrontation with your own experience through meditation, or whether it's through scientific insights, which I think gets you to

a similar place, but by a different direction — not exactly the same place, I'd never claim that. But that does open a space for change. It's the first step, I think. Once you realize that your perception is a construction, once you realize that the self can change and does change, that at least gives you the potential for allowing change to happen, and for inviting ways to make the changes, or to enable or to permit other factors to allow you to change.

Wendy Hasenkamp (00:23:34): So I could imagine on one side, these ideas can be quite liberating and can open up a lot of space, like you were just saying, for possibility. On the other hand, I could also see cases where it could feel quite destabilizing, like — the world is not as it appears, my self isn't either. What does all that mean about reality? So have you, yourself dealt with some of those struggles? Have people given you pushback or issues as you talk about these ideas?

Anil Seth (00:24:06): A bit of both. I think these thoughts about perception have helped me, if anything. I've had my own share of distressing mental experiences. I've had episodes of depression over the last 10 years or so, not too bad, but not pleasant. Really unpleasant, these things are. I think during those times, it has helped, to the extent that I've been able to focus on some of these ideas of perception as construction, and try to recognize that — what's going on in my brain and body to deliver this perception?

(00:24:44) Now this, in a way it's like impromptu meditation. Instead of getting caught up in the thought, noticing it. But my strategy for noticing it is to think about, how does that happen? What's going on in my brain? But it's the same, strikes me as the same process, just putting a little distance between the experience and your perspective on it. So I think it's been very helpful personally for me. It's not a magic bullet or anything. I wouldn't recommend it as a replacement for therapy.

(00:25:14) How these ideas impacts other people is... I think there is a lot of individual variance about that. I think most people find them very positive ways of thinking. But of course, it all depends on context and setting, who you are and what situation you're in when you start to think about things this way.

(00:25:35) This is a Ted Talk a couple of years ago. And of course, then you get a massive exposure compared to pretty much anything else. And so you will get a variety of people who haven't benefited, or not immediately anyway. Part of the problem was that I think the title given to the talk was, "You Hallucinate Your Conscious Reality."

Wendy Hasenkamp (00:25:55): Right! It's a very famous talk.

Anil Seth (00:25:57): And of course, the title is the only thing you don't get to choose.

Wendy Hasenkamp (00:26:00): Oh, you didn't get to choose that!

Anil Seth (00:26:01): Yeah... Any kind of media stuff, right? You never choose your own headline. Right. So I didn't choose that. And I think it's a catchy title, but there's one way in which it can be misinterpreted, which is this idea that nothing is real. And that it's all up for grabs. And so, sometimes, probably once a week I get an email from somebody saying like, "If you think nothing is real, go and jump in front of a train and see how that works out."

(00:26:27) And it's like, "No, that's not what I'm saying." Trains exist. Whether you're looking at them or not. It's the way in which a train appears to you in your conscious perception that's a construction. I mean, Locke the philosopher goes back to this, where there's a distinction between things that exist

independently of our observation and things that don't. Things have solidity independently of whether we look at them or not. Things like money and color, are different kinds of things. They don't have an independent existence.

(00:26:56) So I'm not saying that everything is up for grabs. And this is the key idea about controlled hallucinations that most of the time, our perceptual systems are engaged in this dance between perceptual predictions and sensory signals. But this doesn't mean that an ideal perceptual system would actually reveal things as they are. I think there are two reasons for this. First is, the level at which objective reality is manifest. Whether it's like quantum foam or I don't know what's out there... Who knows, clouds of probability... who knows? There's no reason we should expect our perceptual systems to ultimately have access to that objective reality. And Kant talks about this as the noumenon — this objective reality forever hidden behind a sensory veil.

(00:27:49) So there's that sense. And there's another sense in which, why would evolution have delivered as a perceptual system that's maximally accurate? In a sense, evolution cares about our survival, our reproduction, and it's going to shape our perceptual systems with that as the main criteria. So colors, for instance, they're not an accurate representation of how things are. Colors are an invention of the brain that help us in the service of our behavior. So I think there are probably many examples of this — in psychology, we call these positive biases and things like that — where it's better for us to almost deliberately (well, our brain is doing this deliberately) misperceive things.

Wendy Hasenkamp (00:28:35): Because it's more functional for our survival.

Anil Seth (00:28:38): It's more functional. So back to this thing about self, if we misperceive ourselves as being more stable than we actually are, then we're going to regulate ourselves towards that stability, almost like a self-fulfilling belief, but at the level of basic mechanisms of perception rather than positive thinking.

Wendy Hasenkamp (00:28:55): Yeah. Okay. Can you talk a little more about the importance of regulation of the body in this whole process?

Anil Seth (00:29:02): Yeah. So this is where I think I started to put things together a bit more. Because the starting point for these ideas was — yes, perception of the world outside is a prediction, is a process of predictive inference. Then that applies to the self too. So the way I experienced the self is a process of prediction about the causes of self-related sensory signals.

(00:29:24) And then the next step was to think, well, what's perception for in both of these cases, and why is it that I perceive being a self in a very phenomenologically different way than I perceive the world around me? Like the world around me, especially through vision, is... There's objects in a structured space with spaces between them. But I don't perceive the interior of my body this way. I don't perceive my internal organs in their spatial configurations with colors or shapes. I perceive this basic sense of being an embodied creature. And then I might have moods and emotions, and amorphous senses of agency and will that don't have spatial locations or shapes. Why is that?

Wendy Hasenkamp (00:30:07): It's all very vague, yeah.

Anil Seth (00:30:07): It becomes vague in that sense, but specific in another sense. A very strong emotion is a very specific feeling, just not one that's specific in the sense of color. (Metaphorically, yeah, you can [say] anger is red and so on. But it's not red in the same way that a red thing is red.)

(00:30:25) So I started to think then, well, another way to think about perception is for control. And then there's a whole other literature, a whole different literature from 20th century cybernetics, and one of the early progenitors of AI that's got lost a little bit. But they talked about the importance of feedback for control, and the importance of having predictive models for control. And this could be as simple as something like controlling a guided missile or something, you give it feedback. Or a thermostat and a central heating system. If you want to control something, it really helps to be able to have a predictive model of what it is you're controlling. Because then you can even deploy anticipatory control, you can prevent something from deviating even before it has deviated.

(00:31:14) And if you think about perception of the body, ultimately you want to keep the body alive. So there are things like blood pressure, heart rate, heart rate variability, blood salinity. There are lots of physiological variables that have to maintain themselves within pretty tight ranges for you to stay alive. That's what perception of the body is for, it's not for figuring out what's there, it's for controlling and regulating the "internal milieu," as physiologists call it, the bodily state. And so, that's why I think that perception of the body, predictive perception of body, starts with physiological regulation. That's why when we experience the body, we don't experience objects with shapes. We experience, how well is this physiological regulation of the body going? Is it going well, or is it going badly? Is it likely to go well or badly in the future?

(00:32:08) And for me, this was the sort of final step in a series of inversions. We start out by thinking, "Okay, I perceive the world as it is, and it comes in. Then, the self does the perceiving." So those are the first two. No — we don't perceive the world as it is, it's a construction. And the self is itself a perception. And the third thing is that, actually all of this — all the mechanisms of perception, whether they're of the self or of the world — they all stem from a basic imperative to keep the body alive. So the way in which we perceive the outside world is built through evolution, through development, on mechanisms that are shaped by the imperative to regulate the body.

(00:32:49) And for me, this highlighted another connection with the contemplative tradition, which is this connection between mind and life. And people have argued about this forever in philosophy, what's the connection between mind and life? Thinking of the brain as a computer leads you to divorce the two. The key thing about a computer is it's substrate independent. You can make a computer out of empty Coke cans and pieces of string if you want.

(00:33:10) But this way of thinking just revealed this really strong connection between life and mind. I'm not saying that only living things can have minds. I don't know. But what it does suggest to me is the shape of our minds, or the minds of any living creature, can only be understood in light of their being alive. And this actually joined up with one of my mentors' inspirations, Karl Friston, who also started from a position like this, instead of starting from the position of how we perceive the world around us, like, what does it take for a system to persist over time and be a system? And from that, he derived mathematically all the mechanisms of predictive processing and predictive coding, which I found very, very satisfying. Just because it worked from... You can go from either direction. You can think, well, how do we perceive the world? And then end up figuring out what perception of the body is all about and for. Or you can start from, well, what does it take for a system to persist over time and to regulate itself against perturbations and fluctuations? And then go all the way out.

Wendy Hasenkamp (00:34:25): Right. So I'm thinking of what you said about mind in relationship to life. I was going to say, it sounds like then you wouldn't be on board with anything in the panpsychism world, where consciousness could be a property of matter in some ways. But you made a distinction there. So are you still allowing space for that, or do you feel like consciousness is bound up in life, living systems?

Anil Seth (00:34:51): The honest answer is, I don't know what kind of thing could be conscious. I don't find any of the available positions completely satisfying on this. And I find strong assumptions about them off-putting. So if people, for instance, just assume that computers can be conscious, it's just a matter of having them sophisticated and powerful. I think that's a misled assumption based on the idea that consciousness is closely associated with intelligent systems. I think that that's just wrong and that's part of our arrogance as a species. But that's a popular position in philosophy — functionalism. It doesn't matter what the thing is made out of, it just matters what it does. It's not entirely the same as computational theory of consciousness, but it's related.

(00:35:35) And then you have panpsychism, which is this idea that consciousness is, in some sense — and it depends on the variety of panpsychism that you choose — everywhere. It could be a variety of panpsychism where you say, consciousness is a fundamental property of the universe, much like mass, energy, or charge. It could be. I'm not sure what the utility of thinking that is, though, because even people who strongly believe that don't claim that there is something it is like to be an electron in anything like the sense that there is something it is like to be you or me. So, what testable proposition does it give you? And what problem does it resolve?

(00:36:13) It's a very coherent philosophical position. I was having a dialogue with a philosopher called Philip Goff, who's written a good book about this. And I like to see it sensibly articulated because it's easy to make fun of an idea like panpsychism and say, "Yeah, you just think consciousness is spread everywhere like some jam all over the universe." And some people have said things like that. But the sophisticated panpsychism would not say anything like that.

(00:36:40) And then you have theories like Giulio Tononi's Integrated information theory of consciousness, a very sophisticated theory of consciousness. And it's, what I would say, semi-panpsychist. It's saying that anything can be conscious if it has the appropriate internal structure, but not everything is conscious because not everything does have that appropriate internal structure. So from a theory like that, the right kind of computer could be conscious, but not in virtue of being a computer. And not everything living... There are certain systems that you can say, "Well, that's doing something, but it's not conscious."

(00:37:18) The problem with panpsychism I find is that, I guess I've been learning about myself that I'm attracted to pragmatism in science, and the philosophy of science. And unless it gives me something to do — and it doesn't have to be a particular experiment, it can be a line of thought to follow that leads to a new intuition about why things are the way they are — then I just get a bit bored.

Wendy Hasenkamp (00:37:43): So to come back to what you were saying about bodily regulation, is there a way that that links in with our experience of emotion?

Anil Seth (00:37:50): There is indeed, I think, a very close relationship between emotion and these ideas of predictive regulation. In fact, that was the first, I guess, contribution that I published in this area that I thought, "Oh, that makes sense." And it was just this idea that — okay, if our perception of the outside world is this process of prediction, this Bayesian best guessing, this controlled hallucination, well, then the same should apply to our sense of the body from within. And at Sussex, I learned a lot from my colleague Hugo Critchley, who's a professor of psychiatry there, but a leading expert on interoception — all the sensory data coming from deep within the body, reporting things like blood pressure levels and heartbeat, all these things.

Wendy Hasenkamp (00:38:36): So, helping with this regulation.

Anil Seth (00:38:37): Yeah, exactly. And so it occurred to me, well, why not apply the same principles to interoception? So our perception of interoceptive signals is again, a kind of Bayesian best guess, a controlled hallucination, but one that's geared towards regulation and control. And then the question is, "Well, what does the relevant best guess feel like, in terms of its perceptual content?" And the obvious thing was that, well, that's what an emotion is. An emotion is the best guess of the causes of interoceptive sensory signals. And it's a bit more sophisticated than that. Then you think, well, it's actually, it's the perception of how well regulation of those interoceptive signals is going. But just to put it most simply, when we have an emotion, what that is, is a best guess of what caused the current barrage of signals from within the body.

(00:39:32) And this is just another restatement of a very old idea in the psychology of emotion, which is back to William James and Carl Lange, that emotions are perceptions of changes in bodily states. And so all I was really doing there is just re-articulating that insight in the modern language of predictive processing. And I should say, because I wrote about this, I think in 2011 and then later, and Lisa [Feldman] Barrett, who's a psychologist of emotion here in the US, was having, I think, similar ideas around the same time. And we only recently met to talk about them a couple of years ago, but it was kind of nice to see that convergence thinking. And she wrote with much more specificity and detail about the anatomy than I was able to do, and I think... Well, I'm more convinced having read her stuff now, [thinking] you know, yeah, it could actually work.

Wendy Hasenkamp (00:40:27): Yeah. So I guess it's the same inversion that you're doing of our normal intuitions. I would imagine that many people would view emotions as something that's happening in the world. And then they're reacting to that thing.

Anil Seth (00:40:40): That's a good question. I actually don't know really what the, we call it the "folk idea" of emotion is. I think that would be quite cool to figure out — what do people actually, if you asked them, generally, what do people think emotions are? I think there's probably quite a wide diversity of them, but I think it would be unlikely that most people would naturally think that emotions are predictive perceptions of physiological regulatory, relevant changes in bodily state. Yeah. I think that would not probably come up as people's number one choice of what an emotion actually is.

Wendy Hasenkamp (00:41:21): How does your work and your research on the mind, and these controlled hallucinations and predictions, make you think about the idea of free will? It seems like it could be viewed as a quite deterministic system. And so how does that land for you?

Anil Seth (00:41:39): Free will. We always have to come to free will.

Wendy Hasenkamp (00:41:42): It always ends up with free will.

Anil Seth (00:41:42): We always end up there. It's a very challenging and frustrating thing to think about, I think, but also very important. Because for a lot of people, if you start to talk to them about, where is the essence of themselves, what's really the most important thing about who they are? You can tell them that their perception of their body can be changed, give them a rubber hand or whatever. It's okay. You can have out of body experience. It's okay too. But if you start to say, well, free will also a scientific explanation, and it's not some sort of ability to autonomously determine your actions, free from any prior cause, people get a bit worried. It seems to be, that's an idea that people are very reluctant to let go of. But what are they actually letting go of?

(00:42:33) And the more I think about free will, I start thinking actually less about free will as a problem itself, but more about, why is it such a sticky problem? Why do people find it so difficult to relinquish the

impression that they are an uncaused cause? Because if you look in the dictionary about how free will is defined, on the one hand, it's defined as a voluntary action or decision, which is kind of almost trivial, but true.

[\(00:43:06\)](#) And on the other definition that's usually provided, it's as something that happens without a prior cause, a mysterious human ability to make things happen that are not previously caused, which picks out what people's intuitions are often like. "I did this, I didn't have to do this. I could have done something else. I did it of my own free will. There was no prior cause apart from me."

[\(00:43:30\)](#) But if you take that at face value, it's what I call "spooky free will." It's completely inconsistent with everything we know from physics. And it doesn't even matter whether you buy into determinism or not. This is a big red herring as far as I'm concerned. Because sprinkling a little bit of chance doesn't give you the autonomy that people feel they have when they express their free will. Random behavior is not the expression of free will.

[\(00:43:57\)](#) So, I think it's a total red herring. I think philosophically I'm usually described as a compatibilist. I can think of a free will in a way that's compatible with the universe being deterministic, but actually I don't care whether the universe is deterministic or not.

Wendy Hasenkamp [\(00:44:11\)](#): Okay, so that's not the interesting question.

Anil Seth [\(00:44:12\)](#): I don't think it's actually relevant. I just go back to try to apply the same principles, ways of thinking, that started with this idea of — how do we perceive the world? Because the experience of free will is another perception. I just have... The content of this perception is that I performed an action I didn't have to perform and I could have done otherwise. That's the content of that perception in the same way that the content of the visual perception has things like colors, shapes, other stuff in it.

[\(00:44:42\)](#) And that I think is the key. Actually that's the reason people get so wound up about free will. Because the content of the perception of free will is precisely that the experience does cause the action. That is the content of the perception of free will. Whereas the perception of something else, like a table or a chair or an arm or a leg, doesn't contain that kind of metaphysically problematic step. Like when I perceive something as being red, I'm not perceiving that my experience of redness is actually causing anything to happen. So it's easier to accept that, "Yeah, that's the construction of the brain, it's a perception. And regardless of the actual relationship between consciousness and matter, I can still tell a story about why I perceive the world the way I perceive it in my conscious experience."

[\(00:45:32\)](#) But the content of an experience of free will doesn't easily allow you to do that, because it is precisely that my experience is having a causal role, and it could have been otherwise. That's the problem, but that's also the secret. Because if I think that's — alright, so why am I perceiving things that way? What is it about what's happening in an expression of a voluntary action that I experienced as free will? Why do I experience it that way? Well, it's not that it could have been otherwise, given the state of the universe at the time, take it, sprinkle a bit of random noise, whatever you want, that doesn't help. It could not have been otherwise. I did what I did, could not have been otherwise. But perhaps next time a similar situation comes around, maybe I will do something different because the universe won't be in exactly the same state. If you replay the tape, you'll get the same thing. But everything we do, everything that happens changes the brain.

[\(00:46:27\)](#) So I think — and this is a bit aligned with Patrick Haggard in London — I think that when we have an experience of free will, it's divided into two parts. There's an experience of intention, the urge to do something, and agency, the experience that I did that. I push a button, the light came on. I did that.

We have these experiences when the brain makes an inference that the cause of an action came from within.

Wendy Hasenkamp ([00:46:53](#)): From the self?

Anil Seth ([00:46:54](#)): From the self. From somehow... Not from any external imposition. I wasn't coerced to, I wasn't hypnotized to, I wasn't forced to, I wasn't brain stimulated to. The brain can make an inference about, what were the causes of a particular action. And if that inference is that, well, the causal chain that resulted in the action was relatively internal and was aligned with beliefs and desires of the self, then I will experience that action as freely willed. And part of the phenomenology of that is that I could have done otherwise. Because to have that experience, if the causes are relatively internal, then there is a sense in which the next time you might do otherwise, because you're not being so constrained by the external situation.

([00:47:45](#)) So again, it's just a case of starting with the parallel, starting with the idea that, what happens if I think of an experience of free will in the same way that I think of an experience of color? And then just following the thread from there. And that's... Then I think you can navigate around these red herrings of determinism. And also people again, I think, want free will to be a particular way, because they feel that if they think of it a different way, either they'll dissolve into some psychological puddle on the floor, which is not going to happen, or that somehow all our moral and ethical frameworks will be up for grabs then there will be anarchy, which also won't happen.

Wendy Hasenkamp ([00:48:23](#)): Yeah. So if I'm understanding you correctly, then you're suggesting that the experience of agency and free will and choice, is yet another construction in the case of any "voluntary action." Is that what you mean?

Anil Seth ([00:48:38](#)): That's right. I think... This is basically following me all the way down the rabbit hole. I think everything is a perception. Everything that we experienced can be understood as a perception of one sort or another. And that applies to colors, shapes, tables, chairs, cats, coffee cups, bodies, brains, free will too. So there is a sense in which it's an illusion. The sense in which really is an illusion is the spooky sense. This idea that there's some "soul force" that is still unconstrained by material reality, that can come in, swoop down, intervene, and change the course of physical events in some useful way.

([00:49:16](#)) That just doesn't make any sense to me. Things can't be that way, unless you go back to dualism and say, "Yeah, there's a soul, and it interacts with the brain through the pineal gland or something like that." But if you're not a dualist, then it really doesn't make sense to think that way. And so we have to come up with another way of thinking that does justice to the phenomenology of free will, to the function of voluntary behavior. And I think this is a way that does it.

([00:49:49](#)) And it's a combination of views from different people, I think. I'm still trying to figure it out for myself. I'm writing this book at the moment, and so far it's been just a kind of troll through the ideas we've been discussing. And then of course, I also left the free will chapter till the end! And I've tried to write it three times already now, and it's like, "I have to try and get this story straight." Because it's sort of... it is slippery. We do have such strong intuitions that free will has a kind of reality about it that is somehow deeper than, let's say color. But I just don't think, in the end, that that stands up.

Wendy Hasenkamp ([00:50:26](#)): That makes a lot of sense. I think, when I go through these gyrations around trying to figure out free will questions... Yeah, I'm inclined to think that it doesn't exist, or as you say, it's a construction, which I think is a great way of putting it. But it sure feels like we have it.

Anil Seth (00:50:39): Yeah. I don't think... It's not that the experience of free will suddenly goes away when you realize, or have a thought about, that (if I'm currently believing myself, the story that I'm telling myself about free will), it doesn't go away. I'm still feeling voluntary actions. I still feel this sense of like, "I could have done otherwise, I could have had something else."

(00:50:59) We were talking earlier about what the effect of coming to these ideas scientifically, has for one's personal psychological life. And the effects are subtle. And it's the same... I still see colors. It's not that understanding more about the psychological basis and neural basis of color perception doesn't make you suddenly stop seeing colors. I'm going to continue to feel free will. The hope is that as with all these other ways of thinking about self and world, that if I can understand free will for what it is, and not this spooky soul stuff, then that will make me more at ease with my voluntary choices.

Wendy Hasenkamp (00:51:41): Yeah. It seems like much in the way we were discussing how the way we construct our perceptions, as well as the way we construct the self, are very functional and very useful... Would you say it's also very useful to think that we have agency?

Anil Seth (00:51:55): Absolutely. Yeah, yeah. It's really helpful because, if you take the brain's perspective on it, it's very useful for the brain to distinguish those things that happen, that the brain could have not controlled, had no control over — if somebody just drops a brick on your head, there's no point in experiencing agency over that — and events for which the organism could have had some control. Because we are complex pieces of flesh and blood, and we can do many things. And so I think experiences of agency clearly have a functional role in making that distinction. So the brain can pay attention and learn from those actions that were relatively internally caused, from actions that were relatively externally caused.

(00:52:42) And psychologically, we kind of notice, it sort of helps to feel agency over some things rather than other things. And of course, there's a... One thing I haven't thought much about is, there's a big cultural dependency to this, to the extent that we attribute agency sometimes extrapersonally. It's a relatively individualist, Western thing to take responsibility for everything that we do. God used to take some of that weight off our shoulders, right?

Wendy Hasenkamp (00:53:08): Right. And I appreciate your naming the particular stickiness of this issue, and why it's so challenging for people to let go of. And I wonder if it's also related to the self. From Buddhism, the idea of clinging, and a "self," and the idea of this stable fixed identity over time... It's not necessarily that that itself is the problem, but it's the way we cling to it. And then the behaviors and things that arise from that clinging. So it seems also, maybe, why we find the idea that free will is bit of an illusion so challenging is, similarly, because it challenges again, that sense of self.

Anil Seth (00:53:49): Yeah. I think these lines of inquiry are revealing for the biases that we have about what is important about us. So where are they most stickiest for an individual, or for a bunch of people? So yeah, I think I've learned a lot about what's important to me as a self, through thinking, well, how hard is this intuition for me to be challenged in this way?

Wendy Hasenkamp (00:54:15): Right. So as you've talked about how these ideas have implications for our sense of self, I'm wondering about similar implications for the way we perceive of others — particularly in different groups, stereotypes and things like that, where we may be deploying predictions, or having constructions about others. So, how do you think this relates to possibilities for interacting with others?

Anil Seth (00:54:41): Yeah. I think this is a great question, and it's not something I've worked much on, either on my own or in the lab. But it's one of those things, being at a meeting in the Mind & Life Institute's really fascinating for this, because this is a topic of continual discussion.

(00:54:55) And I think there are two relevant points here. The first is that my perception of another's mental state is also, I think, constructed by the same principles. I am always making a best guess about your mental state, and it's being continually updated by the conversation that we're now having. And it's just a deeper inference to get to somebody else's mental state. It's not just, what light is being reflected from the surface of your face. I have to make a series of further inferences, my brain has to make further inferences about that. And of course my inference about your mental state will partly depend on how I think you're perceiving my mental state. So there's a recursiveness, which makes things even more challenging.

(00:55:38) But ultimately, the idea is the same, which of course means that I can now open the possibility that my perception of your mental state is not the way it is, is not necessarily accurate to whatever the normative criteria might be. It's also a construction. And that we're co-constructing mutual perceptions of each other's mental states to facilitate, in this case, a dialogue. So that's one idea, and that's just thinking of social perception through the same lens. And of course, if we have people then who might come from different backgrounds, we can start to think, well, how would their priors influence how they perceive my mental state? And how does that enter into this recursive cycle, by which we infer mental states?

(00:56:26) But there's this other aspect, I think, which is — any interpersonal interaction is framed by the context in which it takes place. And we tend to assume, because we perceive things as real, that we're both experiencing the same set and setting. You know, we're both in this room here, and I assume you see the table the same way I do. Probably you do, but probably not exactly the same. And it's because we naturally experience the contents of our perception as being real, it becomes harder to appreciate that somebody else might not experience things exactly the same way. And of course, they'll be experiencing things as real, too.

(00:57:06) And there's a trade-off here. I think on the one hand, we have to experience ourselves as inhabiting a shared perceptual reality — it's a kind of ground rules for having a dialogue. We assume we're starting from the same space. But if we overestimate that, then we're not going to understand how people can have different views, or different perceptions. So I think there's a push and pull dynamic going on here.

(00:57:30) And I remember thinking... I was reading the other day about a dinner at the Vatican, where the rebel leader from South Sudan sat with the current president of South Sudan, and the Pope, and a couple of other people. And they ate in silence. And I was wondering, why did they eat in silence? Haven't they got a lot to talk about? Shouldn't they get going?

(00:57:49) And there may have been many reasons why they did this, but it struck me one reason might be that, if they just sit there in a room and eat, then that provides a shared perceptual reality that they're likely to have more in common. Once they start talking, they're going to start going off, their individual perceptual worlds are going to start diverging, are going to be harder to keep them in register. So I just thought maybe whether that's strategically planned, or whether it's happening at all, I don't know, but I thought maybe that's a nice thought. Maybe give people a shared perceptual reality by preventing the talking for a little while. It gives you a stable platform to build on.

Wendy Hasenkamp (00:58:29): That's nice. It makes me think of what happens on a silent retreat as well. And there's a sense of community that can happen just from being in silence.

Anil Seth (00:58:37): Yeah, no, that's right. I mean I must admit, I find that the silence is unusual. I'm not used to it. But I really love it at breakfast. Because I just hate people talking to me at breakfast anyway. So this is like, paradise — when you've got all these people... It's breakfast... Silence. It's wonderful.

Wendy Hasenkamp (00:58:51): It's very unusual. But yeah. The last thing I would just love to get your input on or advice... I think that you're one of the most skillful communicators and translators of this kind of complicated science for public audience and-

Anil Seth (00:59:05):

Thank you.

Wendy Hasenkamp (00:59:05): ... and non-academics. And so I'm just wondering your experience in that world as a communicator, and if you have advice to give to other academics about how they can make their work more accessible.

Anil Seth (00:59:21): That's a very good question. And I guess the first thing that comes to mind about that is that it's really not that everybody needs to do this. I think it's often becoming a bit like this expectation that everybody in science needs to also get up on stage and shout about what they do. And some people should do that, but not everybody has to do that. So it's perfectly fine if that's really not your cup of tea.

(00:59:46) I do it, I started doing it in pubs and small places in Brighton, and I just really enjoyed it. And I think it's fortunate because I was talking about consciousness, and everybody's interested in consciousness. And so, it's an easy sell. You can get people to come and engage quite easily. So, that's fortunate. But I've done progressively more of this... I'm never sure how to describe it. It's not communication. It's not this kind of one-way dissemination of knowledge from the all-powerful scientists to the ignorant public — not at all. It's much more bi-directional and it's outreach, engagement, all these ideas. It is becoming, I think, more important in general, and rightly so. Not only because we get funded by governments usually, so there's a responsibility, but because I think it's part of the dialogue that we, as society, should be having. The voices of scientists should be more prominent within a lot of the current political debates in society, there's something useful to say. Even the scientific method itself is not as well understood generally as it could be. And I think there's actually... Public engagement about the scientific method would be more valuable than anything else right now.

(01:01:03) So the advice is really, if you enjoy it, do it. And it doesn't depend on you being an extrovert, I don't think. I mean, it's a distinctive thing. I think some extroverts hate it. Some introverts love it. In a way, it's a chance to tell people about what you're excited about. And the other thing is, it does take a lot of work. It's like anything, it's a skill. It takes practice. I'm still learning. I've been lucky to have a lot of opportunities to get informally trained, I've not had formal training, but informal feedback. And you just learn through repeated partial successes and partial failures, what lands and what doesn't land.

(01:01:47) And the other final thing I'd say about this is that I've also learned an awful lot. Because one of the great benefits of a wider engagement about your ideas is you come into contact with people you wouldn't otherwise come into contact with, who will talk to you about their work, or talk to you about how their work interacts with what you've been saying. And you will find yourself thinking about your own work in a new way as a result.

(01:02:16) And the final, final thing I'll say, is that it's nice to have impact. It's great... If you write a paper, that's brilliant, and if it's a fundamentally important paper, that's amazingly brilliant. But when you can crystallize your ideas into a form, and you can get them out to a larger variety of people who can digest them, then that can be extremely rewarding and satisfying.

Wendy Hasenkamp (01:02:46): Great. I think that's a wonderful place to leave it. So, thank you so much for taking your time today and joining us.

Anil Seth (01:02:53): Thank you. It's been great to chat.

Outro – Wendy Hasenkamp (01:02:58): *This episode was edited and produced by me and Phil Walker. Music on the show is from Blue Dot Sessions and Universal. Show notes and resources for this and other episodes can be found at podcast.mindandlife.org. If you enjoyed this episode, please rate and review us on iTunes and share it with a friend. If something in this conversation sparked insight for you, we'd love to know about it. You can send an email or voice memo to podcast@mindandlife.org. Mind & Life is a production of the Mind & Life Institute. Visit us at mindandlife.org, where you can learn more about how we bridge science and contemplative wisdom to foster insight and inspire action. There, you can also support our work, including this podcast.*