How Humility and Curiosity Can Help Neuroscience Mature

https://mindmatters.ai/podcast/ep368

Robert J. Marks:

Greetings and welcome to Mind Matters News. I'm your more than a meat computer host, Robert J. Marks. And co-hosting and doing the heavy lifting today is the always effervescent, bagpipe playing, barbecue master chef, Brian Krause. Brian, how you doing?

Brian Krause:

I'm doing great, Bob. It's good to be here.

Robert J. Marks:

Okay. You're still playing the bagpipes and doing barbecue master chef stuff?

Brian Krause:

That's true. It's true. I also like to bake sourdough bread. I think that we should get that on the list.

Robert J. Marks:

Okay. What should I call you? Sourdough master or something like that?

Brian Krause:

That sounds good. Yeah.

Robert J. Marks:

Okay. We have a fascinating topic today, and it's about neuroscience. Modern neuroscience has made some truly amazing advances. Everybody's heard of Neuralink, which is the brain machine interface that Elon Musk has going. But there's other fascinating stuff, too, that I've learned about, such as optogenetics and calcium imaging that they use on the brain. And it turns out now we can record, map, and even manipulate brain activity in ways that once seemed impossible, but there's a catch. We still don't really understand how the brain actually works. We can see what it does, but we can't yet explain how thoughts, emotions, and perceptions emerge from all that neural activity. This gap in understanding has led to a kind of overconfidence, and sometimes even over-claiming by scientists, especially when people say things like, "We are computers made out of meat or our brain is all there is." In truth, our guest today makes the case that neuroscience is still a young science compared to fields like physics and chemistry, because it lacks the strong theories that can predict and explain what's going on.

For now, it can't settle big questions about consciousness or the mind-body problem. And our guest today says that the wisest path is for neuroscience to stay humble, be honest, but stay curiou.s and keep building bridges between, say science, and philosophy, and things we observe. With Angus Menuge, my co-host today, Brian Krause and I have edited a book. And it's called Minding the Brain, that addresses whether the mind does more than the brain. A fascinating chapter in this book is called On the Limitations of Cutting-Edge Neuroscience. The author is Dr. Joseph Green. Joseph Green, which by the way is a pseudonym. He's an assistant professor. He studies computation, and biological, and neural systems. He was a postdoctoral researcher in systems neuroscience. He has a bachelor's and master's degrees in physics, as well as a PhD in neuroscience. Welcome, Dr. Green. How are you?

Joseph Green:

Thank you for the kind introduction. Very well.

Robert J. Marks:

Okay. With that, I'm winded, so I'm going to turn the questioning over to Brian. Brian?

Brian Krause:

Okay. Thanks, Bob. So, Joseph, your chapter is fascinating. And as we've mentioned, the title of it is The Limits of Cutting-Edge Neuroscience. And so, maybe a good way to get into this is to talk a little bit about big picture, those two elements, the cutting-edge aspect of neuroscience, where are we at with the science, and the limits. And maybe one way into this is early in the chapter you contrast the state of neuroscience, big picture, with other more what you call mature sciences, like physics and chemistry. It might be helpful to start there. What do you mean by that? What makes a mature science? And why would you say where neuroscience is on that maturity curve?

Joseph Green:

Yeah. This often sounds like an odd claim, the fact that you can rank sciences according to different traits. But it's also true that different sciences have a different flavor and a different way to go into it. And some sciences, like mathematics or physics, or chemistry are known as being more theoretical, or having a different kind of flavor to it than other sciences that are more in the humanities or in between. Neuroscience sits in between the spectrum of sciences, from more quantitative sciences to more descriptive sciences, because it goes about a system which is maybe the most complicated system that we will ever experience, which is our own brain. You often hear the quote that we have more neurons in the brain than stars in the sky. And I don't even know if that is true, but it's certainly true in that it's very, very difficult to analyze the brain.

And certainly, our science, our way of going about it is still not as mature as the way we go about other research questions, like the atoms in the matter from a chemistry or physical point of view, or theorems in mathematics. So, neuroscience, as I see it, sits a bit in a transition period, hopefully, where it's moving from a more descriptive side of things to a more quantitative description of the brain. But we don't have yet a theory that is leading predictions or generating a clear quantitative understanding of how our thoughts are generated or specific mechanisms in the brain. We're still trying to craft it. We don't know yet what is that mathematical magic formalism that will capture all of these.

Brian Krause:

Okay. I think I'm tracking. So, you use the term it's more descriptive rather than mathematical. So, is that because we're in a state where a lot of the experiments are basically collecting findings and facts, but we don't know yet enough how to build a comparative to a classical model of physics and making predictions?

Joseph Green:

I think of it as when electromagnetism was first discovered. You found that there were these electromagnetic waves and something magical was happening. But we were still far maybe from a real understanding in physics of a real theory that was later developed by Maxwell and others. And back then, you see something, a phenomenon. You want to describe it, but it takes maybe 50 to 100 years to really come up with a math that captures it. And I think in the brain we might be in a similar position, but with respect to a more complicated system, when we don't know if that year is far away, 50 or a 100 years, but we're certain that we still don't have it.

Brian Krause:

And this is the area that you work in mostly, the theoretical side of the neuroscience, trying to come up with these models?

Joseph Green:

Yeah. Part of my job is indeed to model the brain. I see my job as making sure that we can keep going with this understanding of generating models or theories of different brain functions.

Brian Krause:

Yeah. Okay. That makes sense. Now, to this phrase, cutting-edge neuroscience. Can you talk to us about some of these recent neuroscience breakthroughs that you described in the chapter? There's several that sound pretty darn cool. And I'm sure it's probably evolved a bit since the chapter was published.

Joseph Green:

Yeah. Neuroscience is a fantastic field. And one of the thing that makes it fantastic is that it is extremely technologically driven. So, when we say that it's cutting-edge, it's the fact that understanding builds on really new technologies that come up every five to 10 years and revolutionize our understanding of our abilities to record things from the brain, but also our understanding with it. So, for example, in the last say two decades, there are several technologies that came out. And I think one that has been revolutionizing is calcium imaging. This is a specific technology where essentially you modify, genetically, animals, so usually mice. And this genetically modified animals, essentially whenever the neurons take some calcium, they emit some light magically.

You can record, you can put a camera on top of the brain of these little animals. And whenever a neuron activates, you would see some light that is emitted from this neuron. So, you can see videos online of this tiny blobs that are collected over a field of view of an animal's brain. And there you can record maybe 5,000 neurons or 10,000 neurons simultaneously.

Brian Krause:

So, you can measure this while an animal is moving around?

Joseph Green:

You can measure this while an animal is freely behaving. There are different technologies that allow for it. And this is one of the leading technologies to record neural activity in mice or animals in general.

Brian Krause:

So, that must generate a huge amount of data to then you have to figure out how to make sense of, I imagine.

Joseph Green:

Yeah. We have an infinite amount of data these days. Another technology that I think has been revolutionizing, the recent landscape of neuroscience is called optogenetics. And this is the first causal technology, I would say. And what this does essentially is similar to what I described before. You start with a genetically modified animal. Genetic modifications are amazing these days. They can do things that I don't understand, honestly. But in this specific case, what you can do here, essentially you can genetically code an animal in such a way that if you shed light on specific areas of the brain or on specific neurons, this neuron would be active or deactivated. So, you can activate or deactivate specific neurons as much as you want, almost. And that's a dream come true, because you can then understand what those neurons are coding. So, if you are seeing something and then you would deactivate a lot of neurons in the cortex, do you stop seeing? So, this is a real causative power on the brain.

Brian Krause:

So, just like at an individual neuron level, you could turn them on or turn them off, basically?

Joseph Green:

Yeah. Yes.

Brian Krause:

Any examples of the kinds of things that has allowed people to learn?

Joseph Green:

I think, yeah. I mean, one aspect of contemporary science that is coming up is about decision-making and whether or not we can affect the decisions that are being taken by mice, where these techniques are usually deployed in simple decision-making tasks. For example, the animal needs to decide to go right or left. Can you influence the decision to go right or left and modify that to some degree? There are several other deployment cases. In the sensory cortices, we often try to understand whether or not what the animal perceives, either sees or hears, is modified by the fact that we activate or deactivate some of the neurons.

Brian Krause:

That's cool. There's something else, another category, I think they call it brain-machine interface. Is this still an active area of research? I remember reading some things in the news about people maybe with a missing arm or something be able to control a prosthetic robotic arm with capturing signals from the brain. Would you consider this a category in the same vein?

Joseph Green:

It's amazing the fact that we can and even attempt something like this. These are the kind of stuff that draw me into the neuroscience field at first. And they're really progressing fast.

Brian Krause:

They're still progressing fast, yeah.

Robert J. Marks:

So, I have a question, Joseph, about detecting brain signals. They got to go into the brain, so they got to cut a flap out of your head, go in and put the electronics on your head. Is there any way that we could actually do this from the outside? I know that you can actually buy stuff that measures your brain signals from the outside, but I think the noise is just going to be too terrible and you're not going to be able to do anything about it. Do you have any thoughts on that?

Joseph Green:

There are different technologies with which you can record brain signal. And there is also division between doing it, obviously on animals or humans. So, when it comes to humans, one of the most widespread technologies that many people experience is when you go to the hospital, you might get an EEG. And EEG is something that records the activity of the brain, but from the outside. So, they would put some kind of mask on top of your head and then record neural signal, electrical signal from the outside. The problem is that the skull stops most of that signal from coming through. And there is a similar technology that is called ECoG, where essentially for some people that have epilepsy or other sort of brain problems, they would remove a big part of the skull and then do the same but directly on top of the brain.

And in that case, the signal is much, much better, obviously because there is no skull to stop the signal. And yet another technique is when you actually go inside the brain with probes. And in that case there are big probes, long probes, but very small in their section that can be inserted in the brain. And these are used oftentimes for epileptic surgeries. If whenever neurosurgeons want to detect the center of epileptic seizures, they would go inside with some of these probes, and detect, and wait for seizures in some of the patients. And when they collect enough seizures, find the place, essentially, that generates the seizures, and then they're able to go inside and surgically remove or cauterize part of the neural circuit that generates the seizures. So, these are the leading technologies or fMRI as well, which is yet another one.

Robert J. Marks:

Okay. So, I have a follow-up question. I know that the brain is partitioned into different sections that do different things. How do they determine where to put those probes? Are they just hit or miss? Or is this something which is determined beforehand?

Joseph Green:

In the case of epileptic patients, for example, which I think is one of the most interesting cases, you would start with an EEG, like what you capture at the beginning. You would put a mask on top of the skull, record the neural, the electrical signal from the outside, wait for some seizures. And then from there, you could already localize the seizure at the level of the area. Is it in the frontal lobe? Is it in the hippocampus? In which area it is? Once you know roughly where it is, then you can put inside the electrodes and localize it at the level of surgical procedure. So, it's kind of a layered technique.

Robert J. Marks:

Okay. Thank you.

Brian Krause:

It's interesting. It seems like just intuitively like there's a real mismatch in terms of the granularity of the number of things you have measuring. And then the hundreds of billions of neurons inside the brain that might have meaningful signals, that they could make any sense out of that signal enough to actually move an arm, as you would want.

Joseph Green:

It's true. There are many, many incredible things about the brain. One of that, I think, Mountcastle, one of the first famous most neuroscientists, I think he used to call the general relativity principle of the brain, the fact that wherever you look in the brain, in the cortex specifically, the architecture of the cortex is kind of the same, but it encodes many different things. So, you can encode with the same neural architecture vision of whatever we see or with the same neural architecture, like the others that we perceive or the taste. And one of the specific things of this architecture is that is extremely robust. Even if you actually perturb or damage it, you can still... There are incredible cases. You can still sense or function perfectly. So, whenever you have a surgery, if they were to even cut out a good part of your brain, you might still function extremely well. It's really surprising.

Brian Krause:

That's interesting. Wow. So, that's a good list to give our audience a sampling of some of the cool things happening in neuroscience that's pushing the edge of technology and understanding. But what about the other side, the limits of our understanding? You have several interesting examples in the chapter that highlight where... First, blush, you might think, hey, we've really got a handle with all this fancy tech on what's going on in the brain, but maybe that's not the whole story.

Joseph Green:

I would say that there are many different... The brain is a very complicated system and we want to understand it. And we have started understanding and our understanding keeps proceeding at the speed of light in many ways. But yet, the amount of things that we don't know is massive. While our theory develops and it takes major leaps with the different technological cycles that come about, there are still basic questions that we do know how to answer. Recently, I think one of the question that I stumbled upon is the fact that a lot of the cells in the brain are called astrocytes and they're not neurons.

Robert J. Marks:

Are called what? Could you say that again?

Joseph Green:

In the brain, there are many different cells and the usual main actors, the leading actors are taught to be neurons. The neurons are these cells that send information to one another by means of electrical signaling. But there are many other cells in the brain and some of the most common are called astrocytes.

Robert J. Marks:

Astrocytes? Okay. Thank you.

Joseph Green:

Astrocytes. This is a cell that is dominant in the brain. And we still know very little of astrocytes, simply because they have a different kind of signaling and it's harder to record from them. So, one of the axiom, I would say, because it's not really... Their function has been recently more and more studied, with new studies that came up even very important prestigious journals in the last few years, suggesting that astrocytes can encode in a way as much information as neurons, but in a different way. But the bottom line is that we still don't know what astrocytes do or do not.

And we're talking about a lot of the cells in the brain. So, I think this is an interesting example, because it shows that while the vast majority of the neuroscientific community we're talking, the 98% doesn't study astrocytes. I would say a good 50% of them doesn't even care of their existence. It's still true that there are these other major players in the brain that are under-understood or not even considered. And I think that's an example of a bit our understanding. We do understand how neurons communicate to each other, encode information, pass information to each other to a very high degree, but we're still neglecting a lot of what's going on. Another example I would say that came to my table recently is the fact that there are different communities in neuroscience going about different sorts of questions from different perspective.

And some more come from the human studies, human technologies like we just mentioned, and some from the other kind of animal science, where you can do more perturbative studies like optogenetic that I mentioned before. And these two communities have a different way about trying to understand the brain. For example, in the human community, there is a lot of interest in what we call modes or states. What is the brain doing as a whole? Are there waves that travel through the brain from one part to the other? What are these waves about? Do they change with sleep, when you rest, with anesthetics? And this a specific kind of science of the brain, it's about global mode.

But in animal research, it's the opposite. You can go down to a single area, and really pinpoint it and perturb it. And then you look at really the details of those say 1,000 neurons or 100 neurons that encode a signal in a very local way. So, there is this distinction of scales from 100 neurons to millions, and millions and different perspectives that come about it. And we are still very far from reconciling what is happening in between these scales.

Brian Krause:

Speaking of the animal studies, there's an example that I thought was fascinating in your chapter where you talked about the C. elegans. This is a small little nematode worm with, what was it, only like 302 neurons?

Joseph Green:

Yes. It's about 300 plus neurons.

Brian Krause:

But it's the case, I guess, if I understand right, that we can observe these little worms wiggling, and doing their behaviors, and we can measure the signals on all 302 neurons, just get a complete system, unlike the brain where we're taking these samplings of hundreds of billions.

Joseph Green:

Yes, it's true. We can record all of the neurons.

Brian Krause:

And yet is it right that even with a full recording of all of the neurons, we still don't really have a mathematical model that explains the dynamics and the behavior; is that right?

Joseph Green:

Yes. I would say that the claim is a bit more nuanced, but I would say that overall it's still very true that by knowing all of the connectomics, because C. elegans is famous because it was the first animal for which they map the entire connectomics. Which means the connectomics is a word that describes the fact that we know how neurons are wired to one another. So, we know how all of these neurons are wired to one another in the brain, in the animal's brain. And then we can record the neural activity and we know the underlying connectomics. And still, we struggle to come up with a real... this mathematical description of how the animal is computing the information that is important and relevant for its behaviors. And this can be mating or moving and collecting food.

Brian Krause:

Do you have any sense of what's missing? I mean, I don't want to make it seem like an easy thing. But I'm like, come on, guys, there's 302 neurons. Why can't we describe...?

Joseph Green:

Partly is what I was saying before, that the neurons are not the entire description of the system. So, the neurons are believed or have been axiomatically taken, I would say, to be the full description, but there is a lot going on. There is a regulation of genes, the regulations of proteins. There is much more machinery than just the neurons that is really playing a huge role. And there are important studies coming about even there on how the genetic and the replenishing of different parts of the body or simply biological material that is not neural and affect the behavior in an important way.

Brian Krause:

Interesting. Interesting. So, in the same vein, but maybe coming at it from a slightly different perspective, there was also a paper that you described where, if I'm understanding this right, that this is sort of a test to see how well the analytic tools that they're using to decode and reconstruct these dynamics from the data worked on a system that was really well understood. And in this example, it was a microprocessor, like some kind of a engineered chip where we designed it. We know how this thing works, but then they would take measurements that were analogous to the kinds of measurements that we take of the brain and see, okay, can we actually reverse engineer the behavior of this chip? And we found that it didn't really work all that well.

Joseph Green:

Yes. It's a very interesting study. I believe it was a study by the group of Konrad Körding, that ask a very simple but yet complicated question. That is, if we were to approach a chip as if it was a brain, so that we know everything of its architecture, we could record pretty much everything. And we know what it's meant to be doing. We know it's function. Would we understand from the outside what it's actually doing, how it is designed? And the simple answer is that a chip is extremely complicated. And even if you've been around, and you poke here and there, you would have a very, very hard time understanding how it has been put together. And so, the paper that they published was essentially this sarcastic, I find very beautiful and interesting work, saying that essentially what we're doing in the brain, it could be similar.

This system is designed in a beautiful way. It's fantastic. It's crazy cool. And yet we poke it around and we pretend to understand something, but still very far from it. So, the emergence of machine learning has been this humbling, I would say predation for neuroscientists and theoretical neuroscientists like myself, in the fact that all of a sudden we have these neural networks that perform very complicated tasks.

Brian Krause:

You're talking about the stuff that's powering ChatGPT and these familiar systems?

Joseph Green:

Yeah, these days. Yes. Before 10 years ago was doing face recognition and then we saw it learn to play Go and so on. And these machinery really is everywhere now for good reasons, because it's extremely powerful. And we have access to the entire system. It's still a neural network. It has neurons. It does specific computations. We know the data set on which it's trained. We know the learning rule on which is trained, and yet we don't understand it as much as we should, or at least our understanding is piling up at the rate that is lower than what we expected.

Robert J. Marks:

Well, if I could comment, it looks like just probing the chip and trying to figure out how it works without any domain expertise, if you will, is monumentally difficult. However, the designer of the chip knows exactly how it works. And because we know that there's a designer of the chip, we know that there is a mathematics, there is an algorithm that goes into the workings of the chip. I don't even know if we're sure about that with the mind and with the brain. I'm not even sure that we know that a mathematical model exists. Maybe it doesn't. Maybe it does.

Joseph Green:

Yeah, it's an important question.

Robert J. Marks:

Yeah. Okay.

Brian Krause:

Even I suppose for like a nematode?

Joseph Green:

I have many colleagues that believe that the math exists and many others that are not sure whether or not the math for capturing the brain will ever be out there. My opinion, I would say went both sides, because it's unclear whether or not this math has yet been invented or whether or not it's there to some degree. I tend to believe that it's there and one day we'll find it, but I am not 100% sure.

Brian Krause:

Yeah. I'm thinking back to this topic of the prosthetic arm control. It seems that even though we don't have this mathematical understanding, we can still mix some use out of the signals. We can still read intentionality enough off of the signals.

Joseph Green:

I think the important distinction is that I wouldn't consider mathematical understanding in a teleological way. It's not a mathematical understanding. It's the final understanding that if you miss, you can't work with the system. You can work with the system even without full mathematical understanding. We can engineer it without knowing everything of whether or not there are equations that capture every single transitions in the system or process, or mechanism. And this is a beautiful example. We can read out activity of an arm from the brain in a very, very efficient way, in a very precise way, even if we don't understand everything of the brain. So, it works. It's not that we need the math for everything.

Brian Krause:

Yeah. Although, I imagine if you had the mathematical description, that would allow for a lot more careful, fine-grained engineering interactions and things like this.

Joseph Green:

Absolutely. Yes.

Brian Krause:

So, maybe to close out this section of our discussion, where we're talking about the limitations and the advances of neuroscience today, what would you say, can you imagine what it would take for neuroscience to move from this more descriptive phase to get into a more mature science? I'm asking you to just speculate here, I know. But any thoughts on what that path is going to look like? What the research program is going to look like?

Joseph Green:

It is hard to look into the future so far ahead. I don't know exactly what it will take, but I see that it will take a few more leaps, I think, of technology probably as well. On one side, I want to be optimistic. In the last few years we witnessed one of the biggest scientific revolution of the century, like the coming of one new intelligence, like ChatGPT level systems that can do so much. And we thought these were 20 years far away, I would say, or maybe 40. And all of a sudden, here they are. So, I would be very happy if in 10 years from now we have full understanding of the brain. And I think that will take a lot, but I don't think that will happen necessarily in the next 10 years. But I think that will take a scientific breakthrough, unexpected on our causal understanding of the brain.

So, these optogenetic techniques, if they're brought to be working on the entire recording of the brain. So, if we could record all of the neurons and at the same time being able to activate here and there, whatever we reach in a way, while an animal is performing, that would really help us map all of the mechanisms at the circuit level of what is done what, where decisions are taken, information is processed, and map the causal.

Brian Krause:

That makes a lot of sense. By developing some means to actually cause bits of the brain to fire or not at a small scale of granularity, you can probably start to piece together the dynamics and the causality, like you say. That's interesting.

Robert J. Marks:

Well, this has been great. Thank you, Joseph. We're going to continue this conversation. But we've been talking to Dr. Joseph Green and he's a computational neuroscientist. And he has a chapter in the book, Minding the Brain, called On the Limitations of Cutting-Edge Neuroscience. And you never hesitate to advertise your book, so here it is. Minding the Brain, edited by Brian Krause, who is here today, Angus Menuge, and yours truly. And neuroscience is just so fascinating, so we will continue this in the next section. Until then, be of good cheer.

Announcer:

This has been Mind Matters News with your host Robert J. Marks. Explore more at mindmatters.ai. That's mindmatters.ai Mind Matters News is directed and edited by Austin Egbert. The opinions