Dean Buonomano explores the concept of time in neuroscience and physics

He outlines why he thinks integrated information theory is unscientific and discusses how timing is a fundamental computation in brains.

23 APRIL 2025 | by PAUL MIDDLEBROOKS

This transcript has been lightly edited for clarity; it may contain errors due to the transcription process.

[music]

Dean Buonomano
There's this ongoing tension that exists between physics and neuroscience where we like to blame each other. I think neuroscientists like to say, "Listen, we live in a presentist universe. Time is flowing, time is changing, and if that's not what physics is telling you, figure out where the physics are wrong."

IIT is not a neuroscience theory. IIT is a theory of fundamental physics that proposes a new ontology to the very structure of the universe, to the very properties of the universe. IIT proposes that certain configurations of matter are conscious.

Timing should not be seen as a specialized function in the brain because timing is so important to everything we do that it is a universal property of neural circuits.

[music]

Paul Middlebrooks
This is "Brain Inspired," powered by *The Transmitter*. [chuckles] I just hit record there and then totally forgot what I was supposed to be doing. Hello, everyone. Now I remember. It's Paul. Welcome to "Brain Inspired." Dean Buonomano runs the Buonomano Lab at University of California, Los Angeles, UCLA.

Dean was a guest on *Brain Inspired* way back on Episode 18, where we talked about his book, *Your Brain Is a Time Machine: The Neuroscience and Physics of Time*. In that book, he details much of his thought and research and others' thoughts and research, as well, about how centrally important time is for virtually everything that we do. He also describes different conceptions of time and philosophy and how brains might tell time. That was almost seven years ago. His work on time and dynamics in computational neuroscience has continued to this day.

One thing that we discussed today, which actually comes later in the episode, is his recent work using what are called organotypic brain slices to test the idea that cortical circuits implement timing as a computational primitive. It's something they do by their very nature. Organotypic brain slices are somewhere in between what I think of as traditional brain slices and full-on organoids.

Traditional brain slices are extracted from an organism and maintained in some brain-like solution while you perform experiments on them. Organoids, on the other end, start with a small amount of cells that you then culture outside of the organism, let them divide and grow and specialize until you have a mass of cells that have grown into an organ of some sort to then perform experiments on. Organotypic brain slices are extracted from an organism, like the traditional brain slices, but then they are also cultured for some time to let them settle back into some sort of homeostasis set point to get them as close as you can to what they're like in the intact brain, and then you perform experiments on them.

Dean and his colleagues used optogenetics to train their brain slices to predict the timing of the light pulse stimuli that they're using. They find that populations of neurons do indeed learn to predict the timing of the stimuli. They also found that these populations of neurons exhibit replaying of those sequences, similar to replay seen in brain areas like the hippocampus. That's a rough and tumble summary, and we go into a lot more detail in the episode.

We begin our conversation today talking about Dean's recent piece in *The Transmitter* that I will point to in the show notes called *The brain holds no exclusive rights on how to create intelligence*. There he argues that modern AI is likely to continue its recent successes despite the ongoing divergence between AI and neuroscience. This is in contrast to what folks in NeuroAI believe or what many have been arguing for.

We then talk about his recent chapter with physicist Carlo Rovelli titled, *Bridging the neuroscience and physics of time*. In that chapter, Dean and Carlo examine where neuroscience and physics disagree and where they agree about the nature of time. Finally, we discuss Dean's thoughts on the

integrated information theory of consciousness or IIT. IIT has seen a little controversy lately. It's fairly well known, at least amongst my community, that over 100 scientists and a large part of that group calling themselves IIT-concerned, they have expressed concern [chuckles] that IIT is actually unscientific. This has caused backlash and anti-backlash and all sorts of fun expressions from many interested people.

Dean explains his own views about why he thinks IIT is not in the purview of science, namely that it doesn't play well with the existing ontology of what physics says science is all about. It doesn't fit within all the pieces. It can't be examined in the same way that we examine all other science. What I just said, of course, doesn't do justice to his arguments. Again, we hash it out in the episode. He articulates it much better in a few moments.

Lots of topics today. I hope that you enjoy the conversation as much as I did. I want to say thank you to my Patreon supporters. If you support "Brain Inspired" on Patreon, you have access to all full-length episodes, the full archive of all past and present "Brain Inspired" episodes. I have started posting the complexity discussion group meeting recordings there as well.

I've mentioned this a few times on the podcast. We just started a discussion group/journal club to cover the foundations of complexity papers that we discussed in my recent episode with David Krakauer. It's a big group of us. We're over 350 now as of today. Anyway, we just had another one today. I am so happy that I put that together. There are so many people as interested as I am to learn about many of the foundational papers in complexity science. OK, look for all the links and information in the show notes. Here's Dean.

[transition]

You probably hear time jokes all the time. [chuckles] Anyway, speaking of time, do you know how long it's been since you were on "Brain Inspired?"

Dean Buonomano
I'm afraid I do know. I think I actually looked that up.

Paul Middlebrooks
I just looked it up.

Dean Buonomano
Was it 2018?

Paul Middlebrooks
Yes. Six and a half, almost seven years ago.

Dean Buonomano
Good job. Way to go.

Paul Middlebrooks
I think you're in the same office, though. Aren't you?

Dean Buonomano
I can't remember, actually. It's quite likely.

Paul Middlebrooks
You had time-sensitive lighting where if there wasn't motion, the lights would go off.

Dean Buonomano
Good memory. That I don't remember. You were on the forefront of the podcast and you hung in there, so congratulations.

Paul Middlebrooks
I'm still hanging in there. I finally remember reading *Your Brain Is a Time Machine*. This was back when my family, we all quit our jobs and moved into an RV. That's when I started the podcast.

Dean Buonomano
I remember that because I think you were actually in an RV when we first interviewed. Is that possible?

Paul Middlebrooks
Probably so, yes. I remember sitting on the couch and reading your book, actually listening to your other book.

Dean Buonomano
You were also prescient with the RV because you were obviously preparing for COVID.

Paul Middlebrooks
Yes. COVID happened after the RV broke down and we moved into a house. Our timing has been off.

Dean Buonomano
Timing was bad.

Paul Middlebrooks
OK. We have a lot to discuss here. First of all, though, welcome back. Almost seven years. It seems so long. It is so long, it does not seem so long.

Dean Buonomano
Yes. Thank you. It's a pleasure to be back.

Paul Middlebrooks
We have a host of topics to discuss today. I thought maybe we could actually start with the AI side, with the piece that you wrote in *[The] Transmitter* recently. Somewhat I couldn't tell if the tone was pessimistic. Maybe you can tell me, because the point in *The Transmitter* piece is that AI has disregarded brains and more recently, has begun advancing without any regard to dynamics and what are found in our brains.

Then, you suggest I think it's in the title, actually, that that trend will continue and there's no reason to expect that it wouldn't continue. One of the points that you make in your talks recently is that timing has never really been a factor in AI, and especially recently, dynamics, even though you can you can study dynamics and recurrent neural networks, which is a lot of your history, real time is not important in AI and continues to not be important. Do I have that correct?

Dean Buonomano
Yes, I think you got the spirit correct. I would certainly inject a few comments there. First, AI and neuroscience have been sister fields from the dawn of AI. I never implied in any way that AI has disregarded neuroscience. Indeed, AI is based on fundamental tenets of neuroscience. If you look at one of the most fundamental tenets in neuroscience, it's that information, memories are stored in the strength of synaptic connections and the strength of connection rates, and that learning relies on changes in the weights. That tenet is, in many ways, anchored in all of AI based on artificial neural networks.

Then many, many other aspects from convolutional neural networks to things like regularization, which in my mind have aspects of homeostatic plasticity, things like dropout, which have aspects of synaptic failures, things like convolutional neural networks obviously, are based on architecture of V1. AI and neuroscience have had a very synergistic and intimate relationship. I make the analogy that much like if you go back, and I give this example of von Neumann when he was writing down the first code of the first architectures for what we now call the digital computers and von Neumann computers, that he was also inspired by the brain. Obviously, computer science and neuroscience diverged very, very rapidly.

My point is that we are already reaching the point where neuroscience and AI will continue to diverge. I think they're already diverging and they will continue to diverge, and as an indication of that, I gave the example of time. Up into the 2010s, 2011, 2012, when AI started dealing with time varying problems for real in earnest in terms of speech recognition, motion recognition, interaction with external world, the view was that RNNs would be the way for the future because RNNs are how the brain seems to tell time and it seems to be based on internal dynamics.

Then something dramatic happened, which was the 2017 Transformer paper, *Attention is All you Need.* That switched very rapidly from RNN approaches to Transformer approaches. To me, what's absolutely amazing about Transformers is how good they are without any ability to tell time whatsoever.

Paul Middlebrooks
They can tell sequence.

Dean Buonomano
They are timeless, but they can tell order, they can tell sequences, but in a way that's not really time-dependent. This is a good point and so maybe we should clarify. Everybody knows Transformers can absolutely tell sequence. They know the difference between "I am," or, "Am I?" How do they know that? The answer is, it's called positional encoding. They're not processing "I" and then "am", they're processing "I am," or "Am I?" in parallel at the same time simultaneously using positional encoding.

Basically, the token for "I" and the token for "am" has a little code embedded on it, says, "I'm the first token and I'm the second token." Go to ChatGPT and say, "Wait 10 seconds before telling me the capital of France." Either it will do one of two things. It will immediately say, "OK," and give you the answer and not wait, or depending if you ask politely, it might invoke the Python Compiler and actually wait 10 seconds.

Paul Middlebrooks
You mentioned that, I think, in maybe the article and I didn't realize you could do that, you could invoke a Python Compiler in ChatGPT.

Dean Buonomano
ChatGPT, depending on how you're using it will do it automatically. I think it has a little blue arrow where it tells you when it's invoking the Python Compiler. It can, yes.

Paul Middlebrooks
OK. Positional encoding and it's not using time per se, it's using sequence positions, which is sequence. I can't tell if you're pessimistic or optimistic about this, or if you're just pointing out that it's likely to continue. I think what I'm wondering is what to take from that.

Dean Buonomano
I think it was in response, a bit, to this view that AI is only going to evolve, AI is going to only take the next steps with guidance of neuroscience.

Paul Middlebrooks
Who has that view? AI doesn't have that view. Neuroscientists have that view.

Dean Buonomano
Would you say push for NeuroAI has that view?

Paul Middlebrooks
Yes, that's the claim. I think that it's probably largely disregarded in the industry world. I don't know how you feel about that.

Dean Buonomano
I'm impartial, I'm agnostic to that. In some corners, the NeuroAI view is that AI requires neuroscience to continue to advance. Maybe there's some truth to that. Maybe to get the next level, it does. I suspect that because of what I just said, it's amazing how far AI has gone today, largely in the past decade, without taking fundamental principles from neuroscience, as fundamental as time is continuous, as fundamental as the brain processes sequences.

I might prefer not to use the word "sequence", but "ordinality", because sequence implies time. I think I might prefer to use the term "ordinality", which is just labels of first, second and third. The terms here are very vague, but they do generate some confusion.

Paul Middlebrooks
Right. You keep using the word "progress", for AI to progress. Doesn't it depend on how we define what progress is and [laughs] what AGI is, for example?

Dean Buonomano
Absolutely, Paul. Some people's progress might be other people's regress.

Paul Middlebrooks
That's right.

Dean Buonomano
I don't know. It depends if you view eventually achieving AGI or singularity.

Paul Middlebrooks
What does that mean? Does that mean something to you? I know what singularity means better than I know what AGI means.

Dean Buonomano
Again, the irony there is AGI for artificial general intelligence, forget artificial general intelligence, just try to define intelligence. There's no agreement on what caused intelligence. It's not very intelligent to talk about something without defining intelligence. That's true in a lot of fields, right, Paul? You know this, whether it's the consciousness field or the free will field, or even what a gene means. Yes, we have to deal with abstract concepts sometimes that are not very well defined. That's OK. AGI is ill-defined, but so is intelligence. Are humans smart? Is that a high bar or a low bar? Personally, I think-

Paul Middlebrooks
It depends on the human.

Dean Buonomano
-human intelligence is a very low bar for intelligence.

Paul Middlebrooks
OK. We're in agreement there for sure. Maybe we can move on here. I don't know the best route through these topics that we want to discuss because on the one hand, you have teamed up with a physicist to talk about how physics and neuroscience sees time differently. Maybe we should start with that.

Dean Buonomano
I'll add there that we were just talking about Transformers. Transformers might be a good lead into that because they also make this analogy. In Transformer, the past and present are all there at the same time because you're feeding in the whole paragraph-

Paul Middlebrooks
It's parallel.

Dean Buonomano
-to the Transformer simultaneously in parallel. There's a slight analogy there with what Carlo Rovelli and I talk about the distinction between two views of the nature of time. I know we talked about this a mere seven years ago in terms of presentism and internalism. Transformers are very much eternalistic or block universe in that the past and present is there at the same moment, if you will, and the language here gets a bit fuzzy, to feed into the system. A recurrent network doesn't work that way. A recurrent network is inherently presentist.

What Carlo and I, in this piece, *Bridging the neuroscience and physics of time*, were attempting to bridge is the two views in philosophy, physics, and neuroscience of the nature of time. One view is presentism, which is our intuitive view as neuroscientists in terms that we all take for granted. The present is real. The past was real. It's no longer real. We have residues, memories of what happened, but that's all it is. The future is not yet real.

Presentism is the block universe view in which the universe is a four-- In the extreme view, what we call static eternalism is a block universe with four dimensions in which time, in some sense, has already transpired. All of time, past, present, and future, are all equally real, if you will. One way I find convenient to think about this is that now is to time as here is to space, meaning that just as you and I realize that we're in our own location, we're in our here, but other locations are perfectly valid, we accept that. In presentism, we don't accept that I could be in other nows. In the eternalism, yes, you could be in other nows, because all presents are equally valid.

Paul Middlebrooks
I always think of the block universe-- and maybe this is from your book, I don't remember whether you did this, but I always think of it as a block of wax that you could, very thinly, slice through, any slice you want, and then visit the different slices. I'm not sure if you had a figure in your book that made me--

Dean Buonomano
No, I didn't mention wax, but now that you're mentioning it, I think there is a Greek philosopher that somehow makes the analogy with the cutting-edge technology of the time, which was probably wax.

Paul Middlebrooks
There was Descartes' ball of wax, but that was epistemology, not block universe stuff.

Dean Buonomano
Some Greek philosophers, and I think it was Parmenides had this view that it was an unchanging block.

Paul Middlebrooks
Everything is static. I found this chapter difficult to read, and I would also just love to hear how this collaboration came about with Carlo Rovelli, because it was new to me. I didn't know if you guys knew each other. How did it come about?

Dean Buonomano
Carlo and I had been to a couple of meetings together on time, and there's a couple of philosophical symposia on time.

Paul Middlebrooks
It's the philosophy that binds you together then?

Dean Buonomano
Yes, absolutely. There's this ongoing tension that exists between physics and neuroscience where we like to blame each other. I think neuroscientists like to say, "Listen, we live in a presentist universe, time is flowing, time is changing, and if that's not what physics is telling you, figure out where the physics are wrong." Whereas some physicists, and I'm exaggerating here. I don't want to get anybody--

Paul Middlebrooks
This is how you set that up in the chapter as well, it's like a dialogue of the physicist blaming the neuroscientist, the neuroscientist blaming the physicist.

Dean Buonomano
Some physicists, not all, some physicists do view the eternalist view as dogma, and that just take it for granted. In that view, sometimes it's hard to imagine where the flow of time comes from, where we have this subjective feeling of the flow of time. Some physicists view that as an illusion of the mind. They say, "Listen, you neuroscientists, figure this out why we have this illusion of time flowing, because it's not." I'm dichotomizing things here, but that's the idea.

Paul Middlebrooks
That's because math equations work forward and backwards, basically, right?

Dean Buonomano
That's one of the reasons. That's exactly right. This is very counterintuitive to most neuroscientists. Why would the physicists take this view? There's a couple of reasons. One is that the fundamental equations of physics are so-called time reversible. It doesn't matter if you want to run them in the forward direction, you can use that to predict the future, or you can use them in the backward direction to retrodict the past. Also because of relativity. That's independent of relativity, although the equations of relativity are, of course, time reversible.

In relativity, because time and space are a tradeoff, and there's something called Minkowski space, in which the best way to visualize this is that if there's no absolute present, and you're going at one speed and I'm going at another speed, our clocks are moving at different speeds. One way to do this, to go back to your block of wax, is you can imagine all of us in a block. If we're cutting that either orthogonally or diagonally, you can see that we might be on similar planes of simultaneity. That suggests this spatialization of time. That's an interpretation.

As I think, hopefully came out in the chapter, we're both 100% in agreement that there's no strong empirical evidence either way. These are fundamentally open questions, and indeed there's not many ways to even distinguish between those empirically except one.

Paul Middlebrooks
From what I gathered, the goal of the chapter in this case, was to figure out what physicists and neuroscientists agree on, what you still disagree on. One of the things that it seems to me it hinges on, you already mentioned, which is our subjective experience of the flow of time.

Then the other thing that I took from it was, OK, in the mesoscopic temporal scale that we humans live in, it's OK that neuroscientists would have the presentism view, because that's all we can study, because we have only a limited capacity to experience fundamental lower and upper limits of time anyway. Outside of that, where the physicists' philosophical underpinnings of time would come into play, is even outside our epistemological domain, if you will, because we can't measure anything like that, we can't observe it. Does that make sense?

Dean Buonomano
Certainly that we live in the mesoscopic scale, in which relativistic speeds or quantum phenomena aren't on our day-to-day existence, and this makes sense in terms of why our intuitions are absolutely so horrendous and inappropriate for understanding those levels. I don't think the lesson there is that our view of presentism is an artifact of our mesoscopic existence, because the block universe would operate on that exact time scale. The block universe is scale-independent for time.

No, I don't think I would take much from that, and I would say that that's precisely the part of the problem. Unlike gravity near black holes or time at relativistic speeds, the flow of time should occur basically across all scales. My argument is that because the brain evolved to make sense out of the universe in which we live, our mesoscopic universe, and our mesoscopic world is governed by the laws of physics, so the fact that we see it flow, I would argue, reflects reality.

Paul Middlebrooks
You mentioned in the chapter, it's governed by approximately good enough Newtonian physics, and so that's the world we're living in, and therefore that's what our sense of time is, right?

Dean Buonomano
That's what our sense of physics is, and it's what our sense of time is, but the sense of time, again, even in the block universe view, applies to the mesoscopic scale. I would say that no, our perception of time is not a consequence of the fact that we live in a world operating on the mesoscopic scale.

The argument we made, and then I think Carlo and I disagree on this point, is that we perceive it. I would argue that we perceive it because it is actually happening. It is a flow of time. It's important to be clear here, which is absolutely not inconsistent with the laws of physics. The laws of physics allow for presentism, and they allow for eternalism, as long as you're talking about local presentism, not some absolute time à la Newton.

Paul Middlebrooks
I've stopped believing in laws. Is that a real thing? Is a law a real thing?

Dean Buonomano
Have you been getting many traffic tickets since you stopped believing in laws?

Paul Middlebrooks
I've never believed in those laws.

Dean Buonomano
What do you mean you stopped believing in laws?

Paul Middlebrooks
I jumped in there because I was going to ask you what you think about-- You talk in the paper also, or in the chapter, about mathematics. Laws, as far as I can tell, are mathematics. Let's say a constant can change with the evolution of the universe. Would that still count as a law?

Dean Buonomano
I don't know. I think you're picking something that's on the philosophical edge there, and it sounds like you may have read Lee Smolin recently, because that's precisely what he argues, that the constants of the universe are changing.

Paul Middlebrooks
How could they not be? It's all dynamics. [chuckles]

Dean Buonomano
I think they could not be, because the constants are indeed constant. I think they could not be, but they could be, too. I wouldn't say either way, but the bottom line is that, again, going back to our mesoscopic corner of the universe, those laws are pretty goddamn good.

Paul Middlebrooks
Really goddamn good.

Dean Buonomano
We, unusually smart apes, have no right to have really figured out those laws, both at the quantum scale and the cosmological scale, assuming you're not living beside a black hole or a collider. It's absolutely astounding how good those laws are. No, I don't know why you don't believe in laws, but I do, and that was going to come up in our next topics.

Paul Middlebrooks
Right. I think what I actually mean, I just self-elucidated this to myself, is I think I don't believe the same way that I used to believe about what a law is. Maybe that's the right way to put it. I remember Yael Niv saying one of the beautiful things about mathematics is that-- I guess I rail on this too frequently, but she says, "Well, you can take your experiment, you can do the math on it, you can pick up the math, you can move the math and put it down in a different thing and it still works, and that's the beautiful part about it." I like that description of it, and that speaks to laws, I suppose.

Dean Buonomano
You've probably seen me say that, is that the brain is full of cognitive biases. We make wrong decisions. It didn't evolve to do what we're asking it to do, whether it's understanding the laws of physics or understanding consciousness. It didn't evolve. That's not what it's primarily-

Paul Middlebrooks
Some people would actually disagree with that as well.

Dean Buonomano
We can go there, but I would maintain that the brain clearly didn't evolve to understand the laws of physics or to understand the nature of consciousness, not talking about minds. One way around that, to me, I view mathematics as the best debugging device ever invented. Once somebody comes up with a set of equations that describe the incomprehensible, whether it's general activity or Schrodinger's equations in quantum physics, that make no sense. It doesn't matter that we can't understand them, that we can't interpret them. The bottom line is if we use them, we can predict the world around us. Yes, I view the equations as our device that allows us to understand those laws.

Paul Middlebrooks
Those equations, where are you on the platonic side? What? You're shaking your head.

Dean Buonomano
No. I'm not a platonist in any way. I don't think they're reflecting their objects out there in the universe that are--

Paul Middlebrooks
Why is it so perfect?

Dean Buonomano
No, absolutely not.

Paul Middlebrooks
The math is perfect in that way. It's tempting to ascribe-

Dean Buonomano
No, math is not perfect. Math can be perfect in certain configurations. Math can capture truths, but there's a lot of bad math out there. Some things you can write down mathematically. That's just applicable or applied math. Again, we might come back to this topic, as well, in terms of mathematics that is not particularly accurate in some sense in terms of reflecting reality. No, mathematics is agnostic to whether it's good or bad. Mathematics sometimes captures things that reflect how the universe works and other times not.

Paul Middlebrooks
All right. You've alluded to it twice. Are you ready to talk about integrated information theory?

Dean Buonomano
I did allude to it, but you're the boss. It's your show.

Paul Middlebrooks
Oh, yes. I'm the boss. We can be all over the place, so that's fine. Another reason why we're talking today is because you are part of the IIT-Concerned Consortium. Is it a consortium? Is it a group?

Dean Buonomano
I don't know. I think it's just a concerned bunch of people. No, I don't think we're organized enough to be a consortium.

Paul Middlebrooks
All right. We have to define consortium. The problem with this is I don't know if our episode is going to be released before or after the episode I'm about to speak about. I just had on people from a group called Cogitate, whose purpose, from the Templeton Foundation, is to be an unbiased third party to test various theories of consciousness in an adversarial collaboration manner, which means that the proponents of, in this case, two theories of consciousness, have to agree on experimental questions that can be asked where the answers will provide evidence disconfirming one of the respective theories.

The people I just had on were-- I guess they're all postdocs now, I think, maybe. I can't remember. They have been essentially running these experiments that were pre-registered to test the consciousness, I guess it's theory, integrated information theory versus the global neuronal workspace theory. Listeners to our episode will either have heard that already or not. I'm not sure.

Dean Buonomano
Depending if we live in a presentist universe or a block universe.

Paul Middlebrooks
It's all wax anyway. All right. You were part of a group. There's-- I guess it is controversy in integrated information theory world where integrated information theory has gained a reputation as being one of the-- I don't want to say leading theories. You would say that because you're part of IIT-Concerned, right?

Dean Buonomano
Yes. I think it's considered a leading theory.

Paul Middlebrooks
In popular press, let's say, right?

Dean Buonomano
Yes.

Paul Middlebrooks
Therefore it is a leading theory. A lot of people know about it. It seems really cool. Internal causal structure identified with consciousness. Then there was this backlash, and the word that is causing all of the fuss is pseudoscience. There was backlash that a big group of people, not a consortium, suggested that IIT, we can say from now on, is pseudoscience.

You wrote a piece recently, you and many others, hashing out why you do consider it pseudoscience. You justify why IIT is pseudoscience. I think part of the problem before we start talking about this is that in my mind, probably in yours too, the word "pseudoscience" is a slander against something. It is bad to call something pseudoscience, it is, "You're fake news." It's the worst kind of thing to say about anything scientific. Would you agree to that?

Dean Buonomano
Yes, it's a slander. Absolutely. It's not good. I agree with that. By the way, this was obviously talked about a lot in the letter and I think we ended up not using the word "pseudoscience". I can't remember. I think we used the term "unscientific".

Paul Middlebrooks
Oh, OK. That's a little less bad. I was going to say a little better.

Dean Buonomano
They're just words.

Paul Middlebrooks
Yes, but they carry connotations, right?

Dean Buonomano
Yes.

Paul Middlebrooks
Using the word "pseudoscience" in my mind, and I don't know-- now I have to reassess it because semantic drift, but you use the word "pseudoscience", you're trying to just cancel any, validity of anything, but that's maybe just in my head. In the original piece, I didn't revisit this, but I think the phrase was, "Could be considered pseudoscience." It wasn't like, "This is pseudoscience." It was like, "This could be considered pseudoscience from this point of view."

Dean Buonomano
I don't know how useful that direction is.

Paul Middlebrooks
What direction?

Dean Buonomano
Focusing on the words to describe it. I think it's much more fruitful to describe the theory itself, but to build on that, it's important. Words are important. I get that. I think the conclusion was that the word "unscientific" might be less-- Emotions cloud judgment. As soon as you use certain types of words and then people become defensive, then people shift into a different mode where they're using less reasoning. This is the classic think fast, think slow. This is classic systems 1, system 2. Once you bring in an emotional valence to something, it tends to cloud reason.

I don't think focusing on the words is particularly important, but, I understand why people debate on it, but I think it's much, much more useful to focus on why a hundred people signed a letter concerned that IIT is unscientific.

Paul Middlebrooks
I think that that would be fine to start there. There is that original letter and then there's this most recent letter, which, I guess it alludes to the COGITATE work where-- The whole point of COGITATE is IIT has-- you can make predictions based on its theory. Maybe I'll let you say a word about what IIT is and then why it's unscientific, but it makes predictions and then they tested those predictions when that sounds like science.

Dean Buonomano
OK. That's great start to that. First place, let me just put a couple of caveats. I don't think this is a consortium. This is just probably better looked at as a bunch of cats people herded. The fact that they did manage to herd a hundred cats to sign something I think is something that people should reflect on, as opposed to why that came out. I should also say that if you ask a hundred people why astrology is a pseudoscience, maybe they're all going to agree that astrology is pseudoscience or unscientific, but all those people are not going to agree on why.

If I asked you why you may think astrology is pseudoscience, you're probably going to give a different answer than me. I can't speak for the group, Paul. All I can do is share my opinion and that's all I can do.

Paul Middlebrooks
That's what you're here for.

Dean Buonomano
In no which way, shape or form am I representing any group whatsoever.

Paul Middlebrooks
It's important. You said, "Opinion," right there as well.

Dean Buonomano
Yes, absolutely. You brought up COGITATE and maybe we can come back to that. COGITATE is collecting data and collecting data is always a good thing. High quality data that's going to be open source, that's wonderful. I think everybody supports that.

Paul Middlebrooks
Let me just say, it was 250 subjects and they collected magnetoencephalography, MEG, EEG, FMRI, and then they had three predictions that they were testing and it preregistered it all. I already said that. Anyway, great.

Dean Buonomano
All my concerns with IIT have nothing to do with Cogitate. They have to do with IIT and the laws of physics. I think it's important for people to understand, and this is a misconception a lot of people have, that IIT is not a neuroscience theory. IIT is a theory of fundamental physics that proposes a new ontology to the very structure of the universe, to the very properties of the universe. IIT proposes that certain configurations of matter are conscious. This is why some people consider it to be panpsychist, is because under IIT, many forms of matter, whether they're of the neural sort or not, can be conscious. That's the first thing as to why I'm going to focus on the laws of physics here.

The most succinct way I can say why I am concerned that IIT is unscientific, is that you can't propose new laws of physics without integrating those laws into the existing laws, or else you have a free-floating, unmoored, unintegrated rule or law that you can't sanity-check. You don't know if it's consistent with the existing laws or not. Physics is an incomplete puzzle, but as we just mentioned, it's a beautiful puzzle with a lot of pieces already embedded. You can't make up new pieces in willy-nilly without trying to see if that piece you just made up can be integrated to the existing pieces, because then you don't know if you're violating all those other pieces, you don't know if you're violating the laws of physics.

My main objection with IIT is I don't think it fits within the normal scientific method to make up entirely new laws without integrating with the existing, or else it's a free-for-all. What prevents you and me from making up a new law, writing down the equations that describe it, and not integrating it with existing laws, which is required in order to know if there's a violation. In my opinion- and that's just me, in my opinion, yes, I do think IIT may be violating some of the laws of physics.

Paul Middlebrooks
Is the crux here because of the way IIT is developed from the axioms that are supposed to be from phenomenology, from the phenomenological perspective? Is that where it goes unscientific?

Dean Buonomano
No, maybe not. To me, it's really the fact that it's an island. Here there's a bit of irony, because IIT is integrated information, but it's unintegrated with the rest of science. You can't have freely floating theories that don't have any other connection with the other laws of physics. Perhaps it would be helpful if I try to give an example of what IIT is, and where I think it may or may not be in violation of existing laws.

Paul Middlebrooks
That'd be great.

Dean Buonomano
OK. IIT is typically expressed as systems that have internal causal power, which those three words don't necessarily jump out in terms of what they intuitively mean.

Paul Middlebrooks
We're lost already.

Dean Buonomano
I'll do the best I can to actually give a real example. Now, because IIT is not a neuroscience theory, it doesn't require neurons, we can just think of logic gates. Let's think of two logic gates, A and B, and those are such-- They're just called copy gates or threshold gates.

Paul Middlebrooks
They can either be on or off.

Dean Buonomano
They're on or off, and these gates, if they receive a zero, they output a zero. If they receive a one, they output a one. That's all they do. One thing that IIT requires is knowing if the current state of the system has information about past and future states. Now in reality, there's all this partitioning, subnetworks, purviews, mechanisms, and so forth, but for our simple system composed of A and B, it's sufficient to just think about it in the following way. If B is in state 1, does that constrain the past of the system? The answer is yes, it does, because if B is in state 1 now, A had to be in state 1 in the previous time step. That's the only way B could be in state 1.

That's good. That means it has the potential to be conscious. It has the potential to have a positive phi value, but you also have to look in the future. Does B constrain anything about the future? The answer is no, it does not, because B doesn't circle back on itself or onto A. No.

Paul Middlebrooks
It's not predicting anything.

Dean Buonomano
It has no constraints. These constraints are really in relation to above and beyond what a random state would have predicted. The answer is, as you just said, no. It's not predicting anything for the future.

Now here, a neuroscientist, by the way, should ask, "OK, but what do you mean past and future? When in the past and when in the future?" That is a bit of problematic with IIT, because it actually requires you, in some cases, to look at all points in the past and all points in the future within some range. It's discrete. It's well written for a computer because you know what the time steps are, but the brain doesn't work like that. Let me proceed here, but keep that in mind because problems do arise there.

Now, if we connect B back to A, so now we have a recurrent neural network. IIT is about recurrency in many ways, like many theories of consciousness, by the way, which is something I can totally get behind.

Paul Middlebrooks
McCulloch and Pitts liked the recurrence as well. They just didn't know how to handle it.

Dean Buonomano
I'm sorry, who did?

Paul Middlebrooks
McCulloch and Pitts.

Dean Buonomano
Absolutely. There's even a theory called recurrent process theory of consciousness. Now, the state of B will influence the future. It does constrain what can happen in the future. Now it turns out that the system's conscious. We added a connection, and now the system's conscious, has a phi value of one. That's fine.

Maybe it will be useful if we try to do some thought experiments here. One of the things about IIT is that it's space and time-independent. If we get those units and separate them by a huge difference, it doesn't really change. The system is conscious. The second thing that's a bit strange about IIT is it's conscious even if both units are off. They're both zero. Let's get our units and cut the wires between them and use optical beams to connect them.

Paul Middlebrooks
Let's take it to absolute zero Kelvin.

Dean Buonomano
Let's take it to absolute zero Kelvin, although maybe not because, I don't know if it would be responsive then. I don't know if you'd use it as **[inaudible 00:50:30]**

Paul Middlebrooks
All right. I tried.

Dean Buonomano
Are you trying to trap me, Paul? We bring them out to whatever, opposite ends of the solar system. IIT says that, yes, it still has a phi value of one, as long as the connection's good. That's really important. Where is consciousness? It's a bit mystical here, because is it on Earth, is it in Neptune, is it in between? That's just the way it is, but that's fine. That's what IIT predicts.

Here's the next step in our little thought experiment. Let's block the optical path. Now we're going to block the optical path.

Paul Middlebrooks
No communication.

Dean Buonomano
Remember, the optical path is not transmitting anything because they're in zero states.

Paul Middlebrooks
That's right, but we're blocking it without it being active.

Dean Buonomano
Right. You're blocking it without it being active. In our conventional thought process as physicists, nothing should happen. There's no precedent for this in physics--

Paul Middlebrooks
Right. You're going to say the only thing that does happen is that phi goes to zero.

Dean Buonomano
Yes. Very good. It goes to zero. How long does it take to go to zero? By most accounts, it goes to zero instantaneously. Even though these two gates are very, very far apart, it seems to go to zero instantaneously. Now when things happen instantaneously in physics, you have to raise some flags, at least, because you could start running into trouble because there's something called faster-than-light transmission, is a huge no-no in physics. You don't want to get into that. I know you don't like laws, Paul, but it's a law.

Paul Middlebrooks
I like wormholes.

Dean Buonomano
Those are fictitious objects for now, Paul, though. You don't want to violate that in terms of faster-than-light transmission. I don't think that's happening yet. Phi drops instantaneously, but you're not really transmitting any information there. That might be OK. I don't know. Let me now take the final step in the little thought experiment, and let's call it the Middlebrook Split Brain experiment. We're going to get, with your permission, one of your hemispheres and keep it on Earth and the other hemisphere-

Paul Middlebrooks
Take one of my kids. Let's do this with one of them.

Dean Buonomano
[laughs] Do you have extra? Do you have backups?

Paul Middlebrooks
Too many.

Dean Buonomano
No, let's stick with your hemisphere. I don't want to get the kids involved. We'll take one hemisphere, put it on Earth and one hemisphere, put it on Neptune. We have to have these beams connecting them. By most accounts, you'd still be conscious. The thing about conscious entities is you can communicate changes in consciousness.

Now, if we interfere with some of those connections, and if we knew what we were doing-- this is a question, I'm literally asking a question, is that now if we block those connections, that will change your conscious states, and you could communicate, let's say your left hemisphere. The hemisphere that can communicate is over there in Neptune. If we want to go old school, we can put it on Pluto. Up to you.

Paul Middlebrooks
Yes, I'll do it.

Dean Buonomano
Then the question is, now do we have a way to instantaneously transmit information? Because now you can communicate those changes in your phi, in your conscious level. I want to make a couple of caveats here. I think an IIT proponent would make a couple of points here. One is that by adding the delays, that changes this partitioning and maybe your consciousness is just in Neptune or in Earth and it wouldn't work. I think we can get around that because it's a thought experiment and we can match the delays of the inter-atmospheric connections with the intra-atmospheric connections.

Number two, I think a proponent of IIT would point out, "Well, because I'm interfering with the pathway here on Earth, I become part of the cause-and-effect structure." That gets really complicated now because I think the argument might go that the phi goes to zero or is changed by the fact that the system knows that at some point in the future I'm going to mess with those connections. It gets a bit complicated. It's a bit uncomfortably close to clairvoyance for me. My point is the following-- Sorry to give you this whole complex.

Paul Middlebrooks
It's a lot. It's OK.

Dean Buonomano
I know. Sorry to give you this complex, but IIT is a lot. I'm sorry. I think that's why sometimes it's hard to find why it may be violating or not the laws of physics. My point is simply the following, that because IIT is free-floating, because it's unmoored with the rest of physics, I can't know the answer to this question I'm posing. I don't know if it's violating the laws of physics because it's not sufficiently well defined, it's not sufficiently well integrated with the rest of physics to know if those laws are being violated.

That's my main point is that if you have a scientific system in which people can make up new laws, new ontologies without integrating them into the existing, I don't see how that allows for the normal scientific process to unfold because we don't know if they're violating the existing laws. That's my only point.

Paul Middlebrooks
That's your only point. [laughs]

Dean Buonomano
It was a long rant. Thanks for hanging in there, Paul. Thanks for hanging in there.

Paul Middlebrooks
It's been a great show.

Dean Buonomano
I don't know if your listeners will hang in there.

Paul Middlebrooks
All right. I'll turn the tables here real quick. You said all that and yet it could be correct, right?

Dean Buonomano
I guess, it could be.

Paul Middlebrooks
Your point is there's no way to know whether it could be correct because it doesn't-

Dean Buonomano
That's not how science works, it could be correct. The question is how would you determine if it's correct if you don't know? What prevent you and me from, you or me, going home tonight and writing down some equations that are incalculable?

Paul Middlebrooks
If we agreed to a new ontology, it could be correct, right?

Dean Buonomano
I don't think it's a question of we agreeing to the new ontology, is if the new ontology is supported, then we'd agree to that. The first step in that, before we agree to a new ontology, I want to know if it's violating the existing ontologies. To me, that's the first step of the scientific process. I'm willing to say, as I said before, the laws of physics are not immutable. They have room for change, but I need to know if the proposed new ontology, proposed in the absence of any empirical data whatsoever, is consistent with the current ontologies.

Paul Middlebrooks
It's worth saying here that in a Popperian, Karl Popper, sense, the whole business of science is the capacity to actually be wrong, according to the integrated scientific knowledge base that we have.

The other thing that I wanted to say that's very brief, which I find-- I don't know why I find this humorous. After your whole spiel, and we're not talking about Cogitate, but it's interesting that we're talking about Neptune and Pluto and cutting an optic line and instantaneousness or not. Then when it all boils down to it, the prediction is, it's in the posterior parietal cortex. [laughs]

Dean Buonomano
OK. Again, I said the COGITATE project is great. I think collecting data is good, but I didn't say that the way it's being interpreted or used is ideal. I think that is non-ideal because it's being posed as a test of two theories, and as you discussed, that are highly unconstrained. If you read the conclusion statement by Stan Dehaene, he correctly points out that IIT is not really being tested. What's being tested is a loose interpretation of IIT that predicts for reasons, honestly, I don't fully understand that it's in the posterior parietal cortex.

Paul Middlebrooks
For structural reasons, right?

Dean Buonomano
I guess. Remember that, as Dan states, IIT is the mathematical backbone. That's not being tested. IIT lies on a mathematical backbone, which everybody admits can't be calculated, because it's incalculable, because the math, by the time you calculated-

Paul Middlebrooks
Heat death of the universe.

Dean Buonomano
-the phi for a *C. elegans* with 321 neurons, and you started those calculations today, by the time you finish, the sun would have literally expired. It's of limited use, that mathematical structure. Again, this is something else I want to make clear. If IIT proponents are arguing that phi is correlated with consciousness, or a measure of consciousness, I don't think the IIT concern letter would exist.

Paul Middlebrooks
Oh, it's just the identity aspect?

Dean Buonomano
It's, to me, absolutely. They're not really testing IIT, because the, what's applauded in IIT is that it has a mathematical background, but they're not testing the mathematical background, because by definition, phi can't be calculated. At best, it's estimates. I'm going to proceed now.

Paul Middlebrooks
Can we spend a minute, and if you don't want to, that's OK. In the COGITATE case, they ended up testing IIT versus global neuronal workspace theory. Are we OK, as a scientific community, about global neuronal workspace theory? Why would that be OK, whereas IIT isn't? What you're going to say is that, because you can epistemologically test things, and it's--

Dean Buonomano
No, I'd like to back up on that. Can I?

Paul Middlebrooks
Yes. Please.

Dean Buonomano
The answer is quite simple. If you have two theories, Paul, and one requires changing the laws of physics, and the other doesn't--

Paul Middlebrooks
I know, but my--

Dean Buonomano
What?

Paul Middlebrooks
Let's describe global neuronal workspace theory. I don't know if we want to do the Bernard Baars version, or the Stanislas Dehaene version, or whatever. I think the Bernard Baars was more abstract philosophically, and then Dehaene made it more neuroscientific. Would you agree with that?

Dean Buonomano
Yes.

Paul Middlebrooks
The idea, when you write it down on paper, is you have something, some brain activity that ignites, and then is available informationally, is available, that is broadcasted to the rest of the brain, the parts that are necessary for your perceptual sensation of sandpaper or something, when you're running your finger on sandpaper. It's required that there's an ignition, and then it's broadcast. In terms of neurons, I guess it's easy to say, then there should be a lot of spiking activity, and then there should be the proper connections that you can measure with spiking activity in those touch sensorial brain areas, right? Is it simple as that?

Dean Buonomano
I don't think it's as simple as that, but I think it's a fair summary of that. I think the point you're trying to make is that it's also very vague and amorphous, and severely under-constrained. This is what's causing the puzzle, so I'm glad you're bringing it up, because I've talked to people, and they say, you guys are picking on IIT, the IIT concerned letter.

Paul Middlebrooks
You don't want to be known as someone who's just picking on a certain group, right?

Dean Buonomano
People are a bit puzzled by that, because they say, but GNW is a bit fuzzy and vague-

Paul Middlebrooks
Sure as hell is.

Dean Buonomano
-and hand-wavy, as is higher-order theories and--

Paul Middlebrooks
Recurrent processing.

Dean Buonomano
It all is, so people are puzzled. I couldn't see a simpler answer or a simpler dichotomy, is that if your "theories", and we should probably use theories in air quotes, because they're not really developed theories--

Paul Middlebrooks
Shit, now we have to define theory.

Dean Buonomano
They're pre-theories or ideas that we're calling theories, and that's fine. If you suggest a theory of what causes climate change, and that's consistent with the laws of physics, that doesn't require modifying the laws of physics, and I suggest a theory of climate change that causes climate change that requires changing the laws of physics, those two theories are not on the same footing. That's a huge responsibility.

What I think bothers many of us is precisely what you're getting at, is somehow comparing a set of theories, that as imperfect as they are, are not requiring changes in the most successful ontology a scientific humankind has ever achieved, which is the laws of physics, and another theory that does. If COGITATE were to come to the conclusion, which it does, that GNW and IIT are close to 50-50, whatever. Are those two theories on the same level? No.

Paul Middlebrooks
Of course it was going to come out like that. How else could it come out?

Dean Buonomano
My point is if they come out 50-50, but one of them requires changing the laws of physics, who's ahead? It's the one that doesn't require changing the laws of physics.

Paul Middlebrooks
How very Bayesian of you.

Dean Buonomano
Thank you. It is. It should be Bayesian. That's exactly right.

Paul Middlebrooks
Oh, actually in my conversation with them, we were talking about, given the results, how would you even measure which one you believe more? It is a Bayesian method. I think Karl Friston suggested this maybe, which is funny.

Dean Buonomano
The idea there was if you have a lot of things that are graded, you have many measures, how do you do that? You can have a Bayesian approach. As we know in the Bayesian approach, you should probably take in the laws of physics.

Paul Middlebrooks
That's a strong prior. That would be the strong prior.

Dean Buonomano
I cannot, I cannot emphasize how strong of a prior that is.

Paul Middlebrooks
People say it got us to the moon, but we all know we haven't been to the moon, right?

[chuckling]

Dean Buonomano
Fair enough. Good point, Paul.

Paul Middlebrooks
Let's see. I just said Karl Friston. Maybe that's a segue to talk about what you've been up to lately. Have we left anything on the table with-- you did a great job, by the way, of being diplomatic. I don't think you're going to get in trouble. That's good.

Dean Buonomano
I probably am, but Thank you for saying that.

Paul Middlebrooks
Karl Friston, I mentioned him because he suggested this Bayesian approach, which, when I was talking with the COGITATE people, I was like, oh yes, obviously that's the way to do it. Then everyone's response was, yes, everything is obvious in retrospect. When you're in the thick of it, science is hard. Science is hard. It's never obvious what to do. Anyway, Karl Friston wrote, I don't know if it's a commentary on your recent organo-- it's not an organoid, organotypic. What's the difference between an organoid and organotypic?

Dean Buonomano
Organotypic slices have been around for many, many decades since the '80s and let's you get--

Paul Middlebrooks
It's just a brain slice?

Dean Buonomano
Yes. Sorry. In brain slice electrophysiology, people are more familiar with acute brain slices where you extract a slice of the brain tissue of a rodent and then you keep it alive under ACSF, but you can also culture it, and it can stay-- it's a brain tissue slice, and it can stay alive for weeks and months. People now, so there's some papers by the Allen Institute where they've done human organotypics. The tissue is there. You maintain a lot of the structural connectivity, the cell types. We look at it as a little VLSI chip of the brain. It's there, and can we teach it stuff? Organoids are much--

Paul Middlebrooks
Go ahead.

Dean Buonomano
Sorry. Organoids are much more recent development in which you're getting, starting from cells and cell cultures that have pluripotent properties, you can coax them into the early stages of neurodevelopment. That technology is still ongoing. It's a very promising one, but you can't really study learning yet or how adult cortical tissue is processing because those are still very, very early development. The synapses aren't fully indulged, functional, so it's very early development. I think they'll get there, but it's a bit early.

Paul Middlebrooks
It's just not similar enough to in vivo brain structure?

Dean Buonomano
No. Very very different still. They're small 3D spheres that there's--

Paul Middlebrooks
The connectivity, it's hard to get it right basically, in an organoid.

Dean Buonomano
As you can imagine. It's extremely imagined, yes.

Paul Middlebrooks
I just want to back up and say, so I used to extract brain slices from mice when I was first at tech - this is my introduction to neuroscience world - where I would go in and you'd have to sacrifice a mouse. Then you would extract brain, extract part of the brain, put it on a fine-

Dean Buonomano
Mesh?

Paul Middlebrooks
-slicer.

Dean Buonomano
Vibratome?

Paul Middlebrooks
Sorry?

Dean Buonomano
Vibratome.

Paul Middlebrooks
Thank you, vibratome. You'd just get just the tiniest slivers of the brain. You'd put them in a solution that would keep the cells happy and alive. Meanwhile, you've severed many of the connections, the longer-range connections among the neurons. Then you'd let them recover. Then you could do things like patch clamp where you put little pipettes down and listen to the single neurons. You could record local field potentials. You could inject currents and so on. The difference with organotypic is that organotypic, you do that, but then keep growing cells? What is the--

Dean Buonomano
What happens is, as you just said, you sever a lot of the long-range connections. This tissue is just-- each neuron has lost much of its input, so it's in a different dynamic regime. What happens then, our interpretation is then homeostatic plasticity kicks in and it adjusts its thresholds, adjusts synaptic strength to return to some ontogenetically programmed level of activity and dynamics in which you start seeing neurons that act like neurons, where they can-- networks, I should say, that act like networks in which activity can propagate, they can have upstates, they can have dynamics.

Paul Middlebrooks
What's the difference then between a brain slice and an organo--

Dean Buonomano
Typic. Organotypic. In a brain slice, the types that you worked in, generally, they're much more silent. They're not in a dynamic regime in which they can self-propagate activity because they've just lost most of the activity--

Paul Middlebrooks
We wait an hour to record from them or something. Whereas--

Dean Buonomano
No, this is probably days.

Paul Middlebrooks
Days.

Dean Buonomano
Many days.

Paul Middlebrooks
That was a crash course there in brain slice and organoids and organotypic slices. Why did you want to use organotypic slices? You used optogenetics, and maybe we should say briefly what optogenetics is. Actually, let's just say what the question was first. How about that?

Dean Buonomano
The question is, if neurons are computational devices, if neural circuits perform computations, why do we have to study those computations in the intact brain? Neuroscience has benefited from reductionist approaches, whether it's from *Aplysia* or from zebrafish, or from *C. elegans*. Our take on that is just that the best reduced system for the mammalian neocortex is a mammalian neocortex.

We just look at this as our reduced system and that makes it much more tractable. We know the inputs, we know the outputs, we know the experience of the tissue. Then we can control it if we can find ways to give it activity patterns, reproducible activity patterns. That's where the optogenetics comes in. The optogenetics is such that we can stimulate cells or subsets of cells with different optical patterns to see if they can learn and do prediction. There was the connection you were just thinking of in terms of maybe Friston's commentary or something.

Paul Middlebrooks
Yes. I found his commentary interesting because it focused very much in a confirmatory manner for his ideas, but that's not everything. It's just one of many things that you guys found when you did this work. You have the organotypic, brain and you're going to be able to shoot different light frequencies and different light patterns of those frequencies onto different subpopulations of the neuronal cells. Basically you try to emulate, conditioning learning timing of these light patterns, in a traditional way. Allow me to do this and then you can correct me.

You took about 24 hours to just, over and over, bombard this cultured organotypic slice, with these patterns, that paired a conditioned stimulus with an unconditioned stimulus. I should say that more elegantly. Anyway, it's like a classic reinforcement learning paradigm where you're trying to entrain a pattern into the cells. One of my comments was going to be, man, 24 hours at such short latencies, you entrained the hell out of it. I was going to ask you also, maybe perhaps after we clarify what I just said, about how you could fail to see the result that you got.

Dean Buonomano
Fail is always easy, Paul. There's many ways you can fail.

Paul Middlebrooks
I'm doing it right now.

Dean Buonomano
Speaking for myself, I have no trouble failing at any experiment. No, that's not quite our take-home message. If you want me to explain that very briefly, the quickest way you can understand it is we have one stimulus, we'll call it the red light. Then it's just a train of stimuli that lasts whatever, close to 400 milliseconds. Then we have another stimulus. We'll call it blue light because that's what it is. The blue light comes on either at the beginning of the red light or at the end. You have two different groups, an early group and a late group in which only difference between the experience now of the two groups is the temporal relationship between them.

Paul Middlebrooks
They're separate. They're subpopulations. Although I know the red light also slightly stimulates the blue cell, guys. I might have it backwards. Otherwise they're separate populations where when you stimulate with one, you get [mimics sounds] cell firing. That was about 400 milliseconds. Then the blue light was [mimics sounds]. Something like that.

Dean Buonomano
Something like that. You're just training it with different temporal relationships. Remember the auditory cortex doesn't hear. The visual cortex doesn't see. All cortex just receives patterns of action potentials. Is there anything magical about the auditory cortex that makes it better for auditory processing? I don't know if you remember these incredibly popular experiments going back to Mriganka Sur in which auditory stimuli-- I'm sorry, visual stimuli was reoriented into the auditory cortex. Sure enough, shockingly, visual, those cells became orientation selective.

The point of that view is that the cortex is a universal computational device that can process spatial temporal patterns of action potentials. Yes, we're just tapping into that logic. By the way, please, nobody quote me saying that the visual cortex is identical to the auditory cortex. That's not what I'm saying. I'm just saying that ultimately, both sensory or motor, or auditory, or visual cortex are processing spatial temporal patterns of action potentials, because that's all neurons see, is spatial temporal patterns of action potentials. What we're simply testing is whether you can capture some sort of computation or some sort of experience-dependent learning in these isolated autonomous chunks of neural circuits.

After this training in which, yes, you're just doing it over and over again every 20 or 30 seconds for, 24 hours, you want to know is, does the tissue learn? Does the tissue change its structure, its dynamics according to its experience? This is just an early step. I hope, and I'm glad to you, it seems obvious that it will, because it does, because we were happy with that result. What the result is is then you test with just red light now, you just give the train of red light and you want to know, do the cells in different groups behave differently? Is the dynamics of the system different? Meaning that it adapted to its experience, as you would hope it would.

Paul Middlebrooks
Another way to say that is, does it miss the blue light then when it's been--

Dean Buonomano
Let's just back off just one second and first ask, is it different? The answer is that it was different in that the blue and the red light tend to elicit activity earlier in the early group, and later in the late group. One interpretation of that late activity is that it was predicting something would happen, or that it's a prediction error that it predicted was expecting blue light, but it didn't arrive. I think that's what Karl Friston was focusing on, which was not our main point of this, by the way. Our main point was that circuits can learn in an experience-dependent way and to capture the temporal structure.

One of the things that we were most surprised by-- by the way, I should say this is work by an amazing graduate student in the lab, was Ben Liu, who's now at UCSF. One of the surprising things that I never predicted and was surprised when Ben came to tell me, is that there seemed to be replay in the slices in terms of spontaneous activity. He noticed the structure of the spontaneous activity seemed to be different in the early and late group, in which in the early group, spontaneous activity often went on quickly and decayed, and in the late group, that spontaneous activity tended to grow in time and have peaks after the expected time of the bloom.

The overall conclusion is, as you may or may not know, a long line of our research is that timing should not be seen as a specialized function in the brain because timing is so important to everything we do that it is a universal property of neural circuits. One way to test that is to say, if what we're saying is true that neural circuits are intrinsically able to tell time, this is one way to demonstrate that. This is, again, something that might come up in a minute if, depending on the next questions you have to me. This shows that yes, neural circuits, even without a body, can tell time. I'll tie this back up to the IIT because by IIT, these circuits are clearly conscious as well-

Paul Middlebrooks
Cool.

Dean Buonomano
-because that's what IIT would predict, is that you have the recurrence in the circuit, so IIT would also predict that these-- That was one of the ethical implications of the letter, in that it clearly predicts that a lot of organoids fetal tissue and stuff would be conscious.

Paul Middlebrooks
There's time and there's ordinality, like you said earlier. What is the difference in a neuronal circuit? Why would it be telling time instead of ordinality?

Dean Buonomano
In neural circuits, ordinality and time are much more closely connected to each other, because as we talked about earlier, sequentiality requires something flowing in time. When it comes to the brain, I don't normally use the term "ordinality" because ordinality refers to something that's discreet. It's one which is discreet from--

Paul Middlebrooks
I wanted to say sequence, but I wanted to avoid it because it brings, to your mind, it brings in time.

Dean Buonomano
That was applied to Transformers. Not applied to the brain.

Paul Middlebrooks
Then how do you distinguish, I guess, sequence from time?

Dean Buonomano
I don't think sometimes you do. Think of neural sequences, so neural sequences like birdsong studies or [crosstalk]

Paul Middlebrooks
Timing is very important.

Dean Buonomano
There you have sequences. You can have neuron A, B, C, D, E, and that whole pattern can take a second, or you can have A, B, C, D, E, and the whole pattern can take a half a second. The sequentiality is the same, but the timing has changed. That's fine. It's just speed of the dynamical system. There's nothing particularly deep about that comment. It comes up, of course, in that you're using the dynamics. The dynamics is by definition changing in time, and that it provides by changing the speed of dynamics.

Now we have a way to change the speed of our timing, and our I can speak very slowly, or I can speak very quickly. That's the ideas that you have the dynamics, but you can change the speed of the dynamics. There's nothing particularly profound about that. It's just the property of most dynamic systems or many dynamic systems.

Paul Middlebrooks
I can speak slowly too, but that slowdown was uniform. Sometimes I speak slowly, sometimes I speak fast. Then that's birdsong, where the patterning of the temporal internal intervals is very important. Maybe that's a distinction between time and sequence, because then you have timing within the sequence.

Dean Buonomano
Yes. Normally I think of timing as the container. I would say you have sequentiality within the temporal flow, and you can control the speed of that. I don't think it really is fundamental here to understanding of the dynamic. Yes, some things, I would say that everything is flowing in time, and that flow of time can be captured as sequentiality irrespective of the timing, or as truly in the timing in which now you're paying attention.

The way to think about this, is if I ask you a order discrimination task. People study this, I give you A, B, I give you flash of lights. I give you a green light and a red light, and I ask, which came first? That's an order task. If I asked, which one lasted 100 milliseconds, or was the interval between those 100 milliseconds? Now that becomes a timing task. They both require order and sequentiality, but one of them is phrased as a temporal task in that the question requires units of milliseconds. The other sequentially doesn't require units of milliseconds, it just requires ordinality.

Paul Middlebrooks
I had David Robbe on recently. He's actually going to be visiting Pittsburgh in a couple weeks. I've never met him physically, but should be able to say hello to him. He studies basal ganglia. He has come to the temporary, perhaps, conclusion, all conclusions are temporary, that brains do not measure time. That in fact, the way organisms measure time is through behavioral processes. The recent past few years have seen a lot of studies where we're recording lots and lots of neurons, and we can record fine-tuned behaviors from animals while they perform tasks.

He studies timing tasks. Specifically, he studies, the task that we talked about is rats on a little treadmill where the treadmill-- There's a treadmill. They're in the little box and treadmill, there's a little pedestal on the back that they can hang out on. Then there's a reward port up at the front of it. The whole task is just you have to time 20 seconds, and then there's going to be a reward. The rats figure out these behavioral patterns while they're on the treadmill. He thought they would just hang out up at the reward port, basically running at the same speed, and then wait 20 seconds and then get the reward.

Instead, what they did is that they went through a series of attempts, that's anthropomorphizing, but they went through a-- they learned a sequence of behaviors that just so happened to match the timing of when the reward would occur. Then he needed an alternative explanation for this. He's had a history of reading Henri Bergson, was *élan vital* person who has a different conception of time ontologically than Einstein, has famous debates with Einstein. David has come to the conclusion that this more closely matches what he thinks about what time is, that in any way, that time is not measured in brains, but that organisms use their own behaviors as proxies to estimate the time. Sorry for that long-winded thing.

Dean Buonomano
Can I interrupt you a bit? Is that you started this by saying that he argues that the brain doesn't tell time. Is that what you said?

Paul Middlebrooks
Let me play you his question and perhaps--

Dean Buonomano
OK.

Paul Middlebrooks
Let me play you this question. I can play it again if need be.

[audio playback begins]

David Robbe: Hello, Dean. Thank you so much for your inspiring work. My question for you is going to be extremely simple, in fact. If there are population clocks in the brain, who read them? I'm looking forward to hear your answer, and your discussion with Paul, I'm sure it's going to be very interesting. Hopefully we'll have the occasion to meet and chat. All right. Bye-bye.

[audio playback ends]

Paul Middlebrooks
Who's decoding these population clocks? Oh, I said decoding. I shouldn't have said that. We'll just stick to his question.

Dean Buonomano
First place, the question of who's decoding is valid to any form, or most forms of processing, whether it's space or color. I think the question is-

Paul Middlebrooks
That's true.

Dean Buonomano
-coming from a place where people assume there's something special about timing, but it's general of any code, is that who is reading the code? This goes back to little homunculi in the brain. I don't think it's a particularly unique questions timing. I also think it's a question that's fairly well answered, by the way. You don't need a reader. That's the whole idea of a computation. The whole idea that the computation, the code, generates a pattern of activity that generates the motor pattern.

I think sometimes in trying to understand what Robbe's saying, and you phrased it initially as the brain is not telling time, that's a very extreme view. I don't know if he'd really-- sometimes I think he's saying that, sometimes I don't know if he's saying that, but it's also a view that we know is incorrect. Because we know that if I give you a Morse code task, and you're tapping out Morse code, or you're processing Morse code, or we're doing the musical task, we know that the brain is doing that. That's why we're doing this in vitro.

Paul Middlebrooks
If the brain's doing it, but it's constrained by time. I think his point, and I wish he was here, because I don't want to speak for him anyway--

Dean Buonomano
It's hard. I understand.

Paul Middlebrooks
I think his point is let's say I'm tapping Morse code. Of course time exists, we all agree on that. The circuits, the temporal aspects are constrained by my body learning the sequences within the constraints of the neuron pathways, the musculature, et cetera. Then I can get that in order in a rhythm and I don't need to know-- I'm not measuring, this is 200 milliseconds, 200 milliseconds. It's constrained by my mesoscopic, ability to move through the world, my behavioral, brain-to-behavior output.

Dean Buonomano
You and your podcast with Robbe gave that example. I think I know the answer to the confusion, but let me just get there. You gave the example of one Mississippi, two Mississippi. We count. We use the body.

Paul Middlebrooks
In the podcast?

Dean Buonomano
In the podcast.

Paul Middlebrooks
Oh, you listened to the podcast.

Dean Buonomano
Of course I did. You have this situation in which we use our motor commands, our motor lip movement to time something. There's no doubt whether who's controlling the lips. I think nobody's arguing that the lips control the brain. I think we all acknowledge that the brain controls the lips, even if there's feedback, but we can control that feedback. We can cut that feedback, and you still--

Paul Middlebrooks
Good take.

Dean Buonomano
I could read your mind where you were going. We know, because we can-- there's many neurological disorders in which that feedback is corrupted and performance drops, but there's no doubt just who's controlling who. There's no argument there that the brain is generating the dynamics. What I think the problem is, what I think the misconception may be is what time means to the most of the field. When most of the field says timing, what we mean is essentially dynamics.

Sometimes I think his comment is that the brain doesn't have access to time. The brain can't tell time. Clocks can't tell time either. What is a clock? A clock is just a mechanical system. It's just a dynamical system that changes in time. Time is an abstraction that we use to standardize change. There's a great quote by Ernst Mach, the physicist that says it's utterly impossible to measure change using time because time is an abstraction we arrive at by measuring change. When he says that time is not in the brain, that's a bit of semantics, but fine, but time is not in a clock either.

Time is just an abstraction for change. When you have a temporal task, what it means is when we say 200 milliseconds or 700 milliseconds, then what you really mean is that as measured by a clock, this amount of time has elapsed, this amount of change has elapsed, and we need to match that. Sometimes we match that using the body. Sometimes we don't. I can give you an examples of some embodied timing or not.

Here at UCLA, we have an elevator. We have the high-technology pole. The elevator is very annoying because you have to flash your card and within something like 905 milliseconds, something I haven't mastered yet, you have to press the button of the floor you want to go to. How do we do that? What we do is sometimes change the speed of our arm in order to match that. Sometimes the brain uses the body to help us do timing. There's no doubt about that. I don't think anybody would disagree with that. It's incorrect to imply that the brain is not ultimately controlling the body and doing the dynamics, but that's--

Paul Middlebrooks
The brain is part of the body also.

Dean Buonomano
Yes. Fair enough. Thank you for that correction. Either way, yes, the brain and the rest of the body is controlling that. Let's say I'm counting and if I count in binary, so 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16.

Paul Middlebrooks
Oh, jeez. Show off. What a show off.

Dean Buonomano
Show off. Sorry, sorry, sorry. Sorry to do that. Is that embodied counting? No. My brain is controlling my body to help me count. Do you want to call that embodied counting? I don't know. I guess you could if you wanted, but let's never get cause and effect mixed up there. It's clearly the brain that's tapping into the body, even though the brain is part of the body that's controlling that. I still struggle to understand the point.

Paul Middlebrooks
Could this be a case where you're both right? That's often the case. You have single neurons in your organotypic slices, and they can tell time, right? With an approximation.

Dean Buonomano
Absolutely. Again, I think what Robbe would object to that saying is what they're doing is changing. Time is a word we use to quantify change. What my view is, is Robbe is using the word "time" in the way most people don't use it, and that's generating a bit of tension because he thinks that maybe clocks have access to time. No, nobody has access to time because time is an abstraction that we arrive at by measuring change in this case. In physics, we're going to avoid getting the physics for a moment.

Paul Middlebrooks
I've done enough speaking for David here. I can hear him in the back of my head melting about what I said.

Dean Buonomano
Fair enough. Someday I hope I'll get to chat with him one on one, and he can try to address this.

Paul Middlebrooks
Anyway, my thought was, a single neuron, it has all these changes, processes. Those changes take time. The synaptic vesicles, when you're entraining it, you're--

Dean Buonomano
Training.

Paul Middlebrooks
Entraining it.

Dean Buonomano
I don't know. Entraining--

Paul Middlebrooks
I get entrained by my children's music all the time. Then later I hear the same stupid song that I heard, that I don't enjoy, but it gets stuck in my head and it has become entrained in my head. Absent stimulus. A single neuron, it is a living body. An identity and not an identity, but an autonomous thing. You could make the analogy too, the human body is a whole. If I'm timing things as I drum or something, I have all these internal processes that are contributing to my eventual behavior.

I have a question here. A single neuron then, flashing the lights on it, it's getting used to-- recycling the synaptic vesicles in my example at a certain rate because it needs, it needs to. That entrains the ATP production, et cetera. That rhythm happens and there is no time there. It's literal changes in physical processes happening within the living cell. There is no time there. You said, let's avoid physics, but then we have to talk about physics.

Dean Buonomano
What do you mean there is no time there? I think that--

Paul Middlebrooks
Timing is the change. Let's say I'm pushing the limits. If I'm doing it for 24 hours, I'm probably pushing the limits of the cell. You guys, you stimulated in the sweet spot. You could have stimulated too much, and it wouldn't be able to keep up, and it wouldn't be able to learn, or you could have stimulated too little and it wouldn't-- There's a sweet spot of the stimulation within the regime that's possible of a single cell and then populations of cells to keep up and then to be entrained? That's not a question, but that's my question.

Dean Buonomano
Just to clarify, in the first place, these phenomena are circuit properties. They're not single cell. They're circuit properties that require neural circuit dynamics. It's a bit confusing. Sometimes we use the difference between neuronal dynamics to mean single cell properties, and neural dynamics, which now means circuit properties. This is definitely-

Paul Middlebrooks
Apologies for that.

Dean Buonomano
-neural dynamics. Now, what is neural dynamics? Neural dynamics is simply any dynamical system, whether it's a computer or a ball falling down a hill that's governed by the laws of physics, that plays out in time. I think the word "time" generates a lot of confusion. This is my concern in this debate, is that some people take time as being something that's out there. This generates confusion.

When we say it lasts 500 milliseconds, yes, we say that our stimulation protocol was set up, so the time difference between this stimulus and that stimulus will be 500 milliseconds. What is 500 milliseconds? 500 milliseconds is the time it takes a quartz crystal to oscillate 17,000 times. I think this is why these debates are so confusing, because time is such an important concept that we tend to forget that it's not out there in the external world. We've talked about this in the past, Paul. Time is the most common noun in the English language.

Paul Middlebrooks
I was going to ask you that again, because I referenced that. I remember you saying that, or maybe it was in your book. People don't believe me that that's the case.

Dean Buonomano
There is this thing called the internet. People can look up to double-check that.

Paul Middlebrooks
That's where you find all the facts.

Dean Buonomano
It's the common noun. I don't want to say it's the common word. Hopefully, I didn't say the word.

Paul Middlebrooks
I say noun. I'm careful to say noun when I do it.

Dean Buonomano
It wouldn't surprise me if I messed that up at one point. Anyways, it's generally stated as the most common noun. It's a bit ironic to say that it doesn't exist. Hopefully it does. I would push back to say that it does have a purpose. The purpose is that it really anchors much of our lives and the changes that are occurring. When I say that there's no doubt that the brain tells time, what I mean is there's no doubt that the brain has dynamical processes that allow it to synchronize, to anticipate, to match, to decode changes happening in the external world.

In both those cases, in the internal case and the external case, we happen to use the word "time" to quantify those changes. I get it that sometimes it causes a certain sense of confusion. This goes back to Bergson, too. I think he was just using the word "time" in a way that didn't match how most people use the word "time", particularly nowadays. That's why it's not a particularly fruitful line of inquiry.

Paul Middlebrooks
When's your next book going to come out? Are you writing a book right now? You're a book writer. You're an author.

Dean Buonomano
I'm a scientist. I hope to write another book. I'll keep you posted. You'll be the first to know.

Paul Middlebrooks
Keep me posted. Last thing. We started off talking about AI and how time is irrelevant, basically, in AI. Of course in robotics, timing is important. Does AI need time?

Dean Buonomano
I just want to, again, avoid confusion. I never said time is irrelevant in AI. What I said is that it's amazing to me how successful Transformers are in that even though their architecture doesn't allow them to tell time. I think this is your point, is that obviously, in robotics, self-driving car, time is important.

Not only that, by the way, is when you use speech processing, so if you ask ChatGPT, but rather than typing in the prompt, you're now using the speech recognition system, and you say something like, great eyes or gray ties, or they gave her cat food versus they gave her cat food, where the temporal structure is there. The ChatGPT can pick up some of those temporal differences.

It's not really the Transformer. There's a front end that's doing the speech recognition. Obviously, to do speech recognition well, you need to be able to look at the intervals, the temporal structure there. It's more than just the ordinality. When you type it in, it's primarily just the ordinality. I want to make it clear. I'm never saying time is not important there. Your question is is time important for a--

Dean Buonomano
The interesting thing is we use the term "dynamics". We talk about dynamics of systems that are not biological. You can talk about the dynamics of an artificial neural network. Then it's not really dynamics. It's your imposed dynamics. Simulated dynamics, yes.

Paul Middlebrooks
Simulated dynamics.

Dean Buonomano
Absolutely.

Paul Middlebrooks
Time is so important to everything that we do. As you say in the Transmitter article, you turn off a computer, you turn it back on, and it's fine.

Dean Buonomano
It's a very different type of computation. The answer to your question is AI is doing incredibly well without-- self-driving cars, obviously, you need to have temporal structure and have to do speed and stuff. I think they have dynamics in one way, but it's mostly discrete processing. The answer is there's many ways to cope with time. There's many strategies to deal with that, including just doing a feed-forward network with delays. I don't know if we're wrapping up soon. I think we are.

Paul Middlebrooks
We are.

Dean Buonomano
Maybe to wrap that around and go back to consciousness. To me, consciousness is a biological process.

Paul Middlebrooks
I like it. I agree.

Dean Buonomano
Now, now, Paul. It's a biological process that's defined by how it evolves in time. To me, consciousness is like music. It only exists in the flow of time or life itself. You look at life, doesn't make sense to ask if a frozen organism is alive, because life is defined by change, by metabolism, by reproduction, by-

Paul Middlebrooks
Entropy.

Dean Buonomano
-entropy, thank you, and so forth. My guess, and that's all it is, this is not a theory.

Paul Middlebrooks
Oh, come on. Put theory on it. It'll make it more fancy.

Dean Buonomano
No, I'm not going to do that. Is that the one thing we do know is that consciousness is a biological process. At least the only example of consciousness we know in biological processes are things that flow in time. I think it's helpful to, from the get-go, as the consciousness field attempts to go through all the different channels, theories here, is to try to ground itself on the rule of time. IIT, as I understand it, is not a process. It's a state.

Paul Middlebrooks
That's true.

Dean Buonomano
As I understand it, IIT doesn't really require change or time in that. If everything is static, it would still be conscious, as I understand it. Again, it's hard to really know, because it's not very well-defined. There's no other units there. Now, then, in *The Transmitter* article, yes, I speculate that, to me, AIs implemented on standard digital computers, which don't have continuous time, which are discrete in terms of their computations, I speculate, and that's all it is, is speculation, that those would not support consciousness, because they don't support continuous time.

Paul Middlebrooks
Bunomano theory, everyone.

Dean Buonomano
Damn it, Paul.

Paul Middlebrooks
Dean, I guess I'm going to say thank you very much for coming back on. I guess I'll see you in 2032?

Dean Buonomano
Sounds good, Paul. It's a date.

Paul Middlebrooks
Oh, my God. Thanks, Dean.

Dean Buonomano
Thank you, Paul. It's been a pleasure.

[music]

Paul Middlebrooks
"Brain Inspired" is powered by *The Transmitter*, an online publication that aims to deliver useful information, insights, and tools to build bridges across neuroscience and advanced research. Visit thetransmitter.org to explore the latest neuroscience news and perspectives written by journalists and scientists. If you value "Brain Inspired," support it through Patreon to access full-length episodes, join our Discord community, and even influence who I invite to the podcast. Go to braininspired.co to learn more. The music you're hearing is "Little Wing," performed by Kyle Donovan. Thank you for your support. See you next time.

[music]
