THE TRANSMITTER | BRAIN INSPIRED

John Beggs unpacks the critical brain hypothesis

Beggs outlines why and how brains operate at criticality, a sweet spot between order and chaos.

21 MAY 2025 | by PAUL MIDDLEBROOKS

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

[music]

John Beggs
The heart of criticality lies in the laws of physics. There's something about setting up a brain or a flock of animals, or the ear, setting them up at this point where information that comes into the system is not extinguished. It's preserved for as long as possible. It lingers, and you don't overamplify the response, and you don't overdamp the response. You just let the information echo within the system for as long as it can before it dies out with a power law tail.

The reason why I think the cortex is probably most likely to be critical is because it has to simultaneously optimize multiple tasks. It's got to be good at transmitting information. It's got to be good at storing information. It's got to be good at dynamic range. It's got to be good at computing. All these things at the same time.

The field has benefited from the blowback. I have benefited from the blowback. I have grown in my appreciation for the subtleties and the nuances of criticality.

[music]

Paul Middlebrooks
This is “Brain Inspired,” powered by *The Transmitter*. You may have heard of the critical brain hypothesis. It goes something like this. Brain activity operates near a dynamical regime called criticality, poised at this sweet spot between too much order and too much chaos. This is a good thing because systems at criticality are optimized for computing. They maximize information transfer, they maximize the time range over which they operate, and a handful of other good properties.

John Beggs has been studying criticality in brains for over 20 years now. His 2003 paper with Dietmar Plenz is one of the first, if not the first, to show networks of neurons operating near criticality. It gets cited in just about every criticality paper that I read. I think every single paper that I have read cites this paper. John runs The Beggs Lab at Indiana University Bloomington.

A few years ago, he literally wrote the book on criticality called *The Cortex and the Critical Point: Understanding the Power of Emergence.* I highly recommend this book as an excellent introduction to the topic. It is filled not only with great explanations, but also points to a lot of the historical and very recent literature related to criticality. John continues to work on criticality these days.

On this episode, we discuss what criticality is, why, and how brains might strive for it as some homeostatic setpoint. We talk about the past and the present, of how to measure it, and why there isn't a consensus on how to measure it. What it means that criticality appears in so many natural systems outside of brains, yet we want to say that it's a special property of brains. These days, John spends plenty of effort defending the criticality hypothesis from critics. We discuss that and much, much more. You'll hear that John is super scholarly about the subject, even when I ask him about topics outside of his main wheelhouse.

He drops a lot of references, and I link to many of the papers that he mentions. I link to them in the show notes. I really enjoyed speaking with John, and I continue to feel lucky to get to speak with so many interesting and really kind and well-meaning people in this broad, broad field. You would think, based just on the guests that I have on this podcast, that when you walk out into the world, everyone will be kind and well-meaning. Oh, I'm not sure. Is that so? Is that so?

We can only hope it's trending toward that. Anyway, thank you to *The Transmitter* for helping support this podcast. Thanks to all of you Patreon supporters, who do the same. Carry on and be well, and enjoy John.

[transition]

John, I have been awash, much thanks to you and others, in criticality stuff for the past couple of months. I'm working with data sets. First of all, nice job writing the book. It's really readable, it's very enjoyable.

John Beggs
Thank you.

Paul Middlebrooks
The Cortex and Criticality, is that the-- [crosstalk]

John Beggs
Cortex and the Critical Point. Yes.

Paul Middlebrooks
The Cortex and the Critical Point. Since then, once you start thinking about low dimensionality and manifolds in neuronal research, you see them everywhere. Now I see criticality everywhere.

[laughter]

I want to get to the bottom of this. What does it all mean? Anyway, I'll point people to that book. Was it 2003 when the Beggs and Plenz original--?

John Beggs
Yes.

Paul Middlebrooks
Is that the original neuronal avalanche?

John Beggs
That's when we first talked about neuronal avalanches, but as I explained in the book, that wasn't the first time that anybody came up with connecting neuroscience to criticality. There were many people before that who were working on it.

Paul Middlebrooks
I see. That was in a petri dish and across electrodes-

John Beggs
Yes.

Paul Middlebrooks
-in a culture preparation. That's, oh, 22 years ago now. You had your 20th anniversary. Is that the main thing that you think about is criticality these days? What I want to know is how your views have changed over time. First, we should talk about what criticality is.

John Beggs
Sure. Yes.

Paul Middlebrooks
All right. Let's start there, then. Why do I care about criticality, and what is it?

John Beggs
First, what is criticality? Criticality is a special setting on complex systems, and let me start out with a simple example. Let's imagine I have a whole bunch of neurons that are connected to each other, and now I can excite one neuron. I want to measure something really simple. I want to find out if I excite that one neuron, how many other neurons does it excite in the next time stamp? If exciting one neuron leads to more than one neuron being excited, and so on and so on, then what you'll get is I activate one, then I get two, then I get four, then I get eight, and activity will just spread.

This is the real, simplistic version of it. That's not good for a brain because it would lead to seizures. On the other hand, if I have one neuron that's excited and it leads to less than one neuron being activated in the next time stamp, then activity is going to die out. You don't want to over-amplify, and you don't want to dampen the activity. What you want is something where if I stimulate one neuron, it leads to one other neuron, which leads to one other neuron on average.

Sometimes it'll activate two, sometimes it'll activate one, but the average is about one. That's really simple to explain, but it actually has really complicated and interesting implications. If you're going to transmit information through, let's say, a network of neurons, you don't want that ratio. We'll call it the branching ratio. You don't want the branching ratio to be greater than one because then what happens is I have some input for an input layer, and then it goes through a layered network.

At the output, it's all going to be saturated. Then you won't be able to guess from the output what the input was. It will be lost because almost all cases, the output's going to be totally, all neurons are on. That's bad for information, but it's also bad for information if it's damped for more obvious reasons. If I activate one, it's going to die out, and I'll have all zeros at the output, and so you don't get anything. If I want to transmit information through a multilayer network, it's best to do it with a branching ratio near one, as close as possible. That's also where they run nuclear reactors.

You talk about your branching ratio. If it's greater than one, you've got a bomb. If it's less than one, it's subcritical, and you don't generate energy. You want to keep it steady, percolating along like that. That's the real simple picture of what criticality is, and a little hint at why it's important. If you want to transmit information through a neural network, it's best right at that spot. One more little bit that I should throw in, maybe if people like complex systems, and I know you do, and you've been working on those things, you've heard of the Mandelbrot set?

Paul Middlebrooks
Yes.

John Beggs
That basically is very similar to this idea of a branching ratio equal to one. What you do is you take a complex number and you square it, and you see if it's equal to this complex number by itself plus some constant. What gives you the Mandelbrot set is if you iterate this again and again and again, and it doesn't grow and it doesn't shrink. If you zoom in and you see these pictures, a little snowman and fractal snowman at smaller scales or bigger scales, that little boundary line is determined by something that essentially has a branching ratio of one. It has complexity to all scales. A very simple rule can lead to really complex outcomes.

Paul Middlebrooks
I guess I'll just jump right into it. It seems like criticality is everywhere. What I want to figure out is what's different about criticality in brains relative to other systems. There's fractality everywhere. Is it special to living systems? I know Per Bak, the original sand pile stuff. There's criticality there, which is not a complex life system. Then the idea is often criticality is stated as being at a phase transition, right?

John Beggs
Yes.

Paul Middlebrooks
At a transition between states. You and others have written about perhaps there's a homeostatic set point in populations of neurons or neurons basically trying to maintain near criticality, whether it's exactly at criticality, which we'll get into--

John Beggs
We'll get to that later. Yes, sure.

Paul Middlebrooks
What is special, then, about brains with respect to criticality? Why do I care about brains-- if it's everywhere, then if it's also in brains, is it that special?

I see criticality everywhere now.

John Beggs
That's an excellent question. It is often raised by reviewers of papers and grants-

Paul Middlebrooks
Oh, okay.

John Beggs
-that's likely to deal with it in certain ways. Let me see if I can start with the brain business. Brains are inherently interested in information processing. What I talked about earlier was just this idea. If you want to transmit information through a network, you want to be near the critical point, and then you'll minimize your losses. Other things that brains do would be, let's say, computations. This is a little bit harder to measure, but some people have been playing with this wonderful idea of reservoir computing. I don't know if you're familiar with that or not.

Paul Middlebrooks
Liquid state, liquid-- [crosstalk]

John Beggs
Liquid state machines, that's another echo state network. This whole idea is you take a system and now you ping it with some stimulus. Now it's going to propagate through the system. Now, imagine if the system is too excited. I ping it, and now everybody just turns on, so it's an explosion. That would be a branching ratio greater than one, or if it's a very damp system. It's just molasses. I dropped something in here, and then no waves go anywhere.

That's an overdamped system. That'd be a branching ratio less than one. It turns out if you want to get optimal computing-- and there's a famous paper by Legenstein and Maass, there's also a Bertschinger and Natschläger. I can give you the references, and you can link them to this later.

Paul Middlebrooks
Cool.

John Beggs
Basically, what these people have done, it's really beautiful work, is you take a network, a reservoir, and you want to use it as an input-output mapping device. You can get the best input-output mappings, the most versatile input-output mappings, if you have a reservoir that has complex dynamics. What they do is they take their artificial networks and they're tuning them, and they tune them basically to the critical point. When they get at the critical point, it can compute the widest variety of input-output mappings.

Paul Middlebrooks
It's a capacity issue?

John Beggs
Exactly. It's computational capacity. Information transmission, computational capacity. If you talk to Woodrow Shew, he's going to talk about dynamic range. Brains are faced with this issue of dynamic range. We go out to, I don't know, a beach in Belize, a white sand beach in Belize, and we're getting trillions of photons. You lock yourself in a room and you're trying to go to bed, and you black out your curtains. Your eye, once it's adapted, can still detect one photon. This is many orders of magnitude.

If you want a system that can basically respond to huge dynamic range, again, you want to be at the critical point. For sensory reasons, you want to be there. Mauro Capelli has also looked at this kind of thing. He's published a paper a long time ago that was, how do you optimize dynamic range? Operate at the critical point for sensory systems. Computational capacity, information transmission, dynamic range for sensory stuff. The other thing about being at the critical point is if you're subcritical, you go into an attractor. You're sucking into some really stable state, which might be good for memory recall.

Paul Middlebrooks
Oh, see, I was thinking of the critical point as itself an attractor. Is that the wrong way to think of it?

John Beggs
It could be. You could rig it up that way. If I go back to the simple branching model, I could stimulate some neurons and then see how activity travels through the network. Now, if I've got a network that's got varying strengths of connections, so some are stronger and some are weaker, and so forth, what will happen is when I stimulate this thing, it'll go through a bunch of paths. It will tend to pick the path that has the strongest connections, most often. It'll sometimes go through the paths that have weaker connections. That's fine.

If I make the network really subcritical, in other words, I make all the connections really weak, I really don't get any repeating patterns at all because the signal just dies out. If I make all the connections strong, then what happens is I do have an attractor, but it's really just one giant state. I stimulate, and the whole brain goes on and everything's lit up. I have the widest variety of potential repeatable paths above chance when I'm at the critical point.

You can show that with computational models. I had a paper with Clay Haldeman back in 2005 where we did that with different models. These repeating pathways, you have the widest variety of them when you're at the critical point. That's related to memory storage. If you want to store information with networks of this type, you're actually better off being at the critical point. Transmitting information, computing, dynamic range, storing information.

Paul Middlebrooks
Universality is another one that you talk about in the book.

John Beggs
That is another interesting thing. Maybe we want to bracket that for a minute, because that's more of something that is good for us in terms of understanding the brain, but I don't know that the brain itself directly benefits from it. It is a juicy topic.

Paul Middlebrooks
One of the confounding things about this notion of the critical brain hypothesis is there's no one measure. You've done a lot of work measuring what are called neuronal avalanches, which is what you were talking about when activity goes from, let's say, zero to then one unit is active, and then two, and then one, and then three, and then four, and then zero. That's defined as a neuronal avalanches. The way that one measure of one indicator of criticality is if the distribution of the sizes and the durations of those neuronal avalanches follow a log scale-free plot dynamic.

John Beggs
Exactly.

Paul Middlebrooks
You mentioned branching ratio, which is another indicator. Long-range temporal correlations are another indicator. Why do we have to triangulate about criticality? The thing is, different analyses have occurred over time. I told you offline, I'm talking with Woodrow Shew as well. This is a cottage industry of creating new analyses. Where are we with all that stuff? Why isn't it easy? [crosstalk]

John Beggs
Why so many toothbrushes? Why can't we all just use the same toothbrush? First of all, I would say this, that in part, there's a good reason why people want to come up with different measures. This is scientific rigor. Then there's another reason, which we'll get to in a minute. In the beginning, when I started out as an assistant professor, I was naive. I looked at the data that Dietmar and I got, and I said, "Oh, it follows a power law. Therefore, it's critical." I was reading Per Bak's, *How Nature Works*. Very modest title, explaining the whole gamut of criticality?

Paul Middlebrooks
I have the book, but I haven't read it. Do you recommend it? I can take this out of--

John Beggs
I recommend it as a historical perspective of the ferment at that time. They were explaining everything from stock market crashes to evolution, extinction-

Paul Middlebrooks
Everywhere. It’s everywhere!

John Beggs
-to the brain, to everything. I think it's a useful way of framing things. Per Bak was later criticized, we can talk about that more, for perhaps being a little too sweeping in his claims. Anyway, this idea of measuring whether something's critical. I started out, and I was swept up with Per Bak and other advocates of this. I said, "Oh, as long as I got a power law, then it's critical." I said that in talks and almost said that in papers. Then I started getting a bunch of blowback. This is why I'm an assistant professor, really severe criticism, and so forth. Then I was about thinking, "Oh, my gosh, I'm not going to get tenure. I'm getting ripped to pieces. It's not going to survive."

Then I started looking at other measure, and fortunately, I had really good colleagues who were talking to me about things. One of them was Karin Dahmen. She's a condensed matter physicist over at University of Illinois. She and colleagues, Myers and Sethna, came up with this idea of the crackling noise relation. What that basically says is, let's say you get avalanches and you get the exponent for the avalanche sizes, and you get the exponent for the avalanche durations.

Those things can be related by a simple algebraic equation to another power law, which is the average avalanche size for a given duration. All these three things form power laws, and there's a simple algebraic relation called the crackling noise relation that relates them all to one another. You should really only get this if you're close to the critical point. That was one of the things that rescued me. I collaborate with her, and we said, "Oh, yes, you can get that," and our data ended up showing that.

That's a more stringent test for whether you're really at the critical point. Because you can get scale-free power laws, even if you're not at the critical point, from a different type of process. Let me give you a quick example of that. There's this observation that meteorite sizes or crater sizes on the moon follow a power law. You get lots of little ones, you get a decent range of medium ones, and a very small number of huge ones. How do you explain that? One way of explaining that is through something called successive fractionation.

Let's say you have a stick, and now you crack that stick at some random location. Then you take the fragments and you crack them at random locations, and you do this N times. Now, what you're going to get is you're going to get some small number of fragments that were not short. That rarely happens that this one didn't get cut in half, but you're going to get a lot of little ones that are like dust. If you do this, you get a power law of lengths. You just take a stick, you could do a MATLAB program and just divide it, pick a random spot, then divide the fragments, and keep doing it, you'll get a power law of lengths.

You could call it critical, but it's really just some process. A meteorite's coming in, it gets hit by something, it fragments into two parts, and it fragments again, and then whatever ends up hitting the surface creates a crater. There's no need to posit a phase transition or a critical tuning thing that allows you to have optimality right there. It's just something that's produced by randomness. What you've got to do is you've got to distinguish a phase transition power law from a power law that could come through some other way. This crackling noise relation that we just talked about, that's something that you're not going to get with these other mechanisms.

Paul Middlebrooks
What does the crackling noise look like if I'm not near criticality?

John Beggs
What'll happen, and this is something that Keith Hengen and colleagues have looked at. He's at Washington University, Saint Louis, they're looking at homeostasis of criticality. They basically say, "Let's plug in the exponents for the size and the duration and see what it predicts the other exponent should be. Then let's empirically measure the other exponent, and we'll see the difference between them." He calls that the difference in criticality coefficient. The distance of criticality coefficient, the DCC, tells you how far away you are. You're right. You aren't always going to be at criticality, and you can measure how far away you are from criticality using that.

Paul Middlebrooks
Isn't criticality itself an infinitesimally-- It's an abstract notion that you can never be exactly at criticality or else we will all explode or something, right?

John Beggs
Exactly. You're correct. Yes, you're correct. It's really a concept that only works for extremely large numbers. It works in the thermodynamic limit where something is infinite. Now, you can approach it, and so you can get a power law over some range. Maybe you get it over two orders of magnitude, or if you're really good, there's some work that Michael Breakspear and others have done on seizure-type events in infants. If infants are exposed to hypoxia, so they have episodes where they don't get enough oxygen when they're coming out of the womb, they can have these bursts, and you can capture these things with EEG caps.

These bursts, if I remember right, they have five orders of magnitude in terms of fractal scaling. They'll see a burst, and they'll see another one bigger, another one much bigger. Five orders of magnitude. You get a power law of five orders of magnitude, but it's not infinite. In the ideal, you're only going to get something that's a perfect power law spanning over all orders of magnitude if you're at the perfect critical point. We're never at the critical point. We're never going to be there for lots of reasons. We're finite size, but the other thing is I'm constantly being driven by sensory inputs. That will perturb the system away from criticality, but homeostasis might bring it back.

Paul Middlebrooks
Now I'm being selfish. What I want to do is share my screen and show you all my data work.

John Beggs
That's okay. We can do that.

Paul Middlebrooks
No, this is a podcast. I'm recording in mouse motor cortex and basal ganglia while they're doing various things. I get exponents for the size and duration distributions, and I get the crackling noise exponent. They are logarithmic, but they're slightly away from criticality. 1.2, whatever that range is. It's frustrating because I don't know, is 1.6 not critical? What is that range, and how close to critical am I? Does that matter? The slope of the line, does that actually matter as long as it's a straight line, as long as it fits?

Anyway, this is what I wanted to ask you about. I think I didn't articulate this well when I was speaking with Woody this last week is, so we can't ever be at perfect criticality because that would be across all scales. You'd have infinite dynamic range. Does it make sense to ask if a system is in criticality within the dynamic range for which it is functionally operative? Our brains operate at a certain speed, our behavior is at a certain mesoscopic temporal scale. It would make sense then to have a fairly narrow dynamic range in which, let's say, a population of neurons is acting in a critical manner relative to the information that's sending to other parts of the brain or to the spine or otherwise. Does that make sense?

John Beggs
Oh, it totally makes sense to me. If criticality is a place where healthy operation occurs, then you'd expect the brain to, on average, get near it, but not always be exactly at it. In the same sense that my heart rate has a certain optimal range, but sometimes, it's going to go up, and sometimes it's going to go down. I have mechanisms that bring me back into this range. My blood pressure, or my pH, or anything like that. These things are within some band that we consider healthy.

That's what I think you can see in beautiful experiments by Keith Hengen and his colleagues. They basically take animals and they perturb them. One of the examples they gave is they close one of the eyes and they record from the contralateral visual cortex. They show that right after this thing being closed, it goes subcritical. Then they look at that crackling noise relation, and he uses the distance to criticality coefficient, and then over a little bit of time, I think it's slightly more than 24 hours, basically, the brain restores itself to being critical, even though it's not receiving any inputs.

Paul Middlebrooks
I think if I-- You can correct me, the criticality markers occur before the-- When they suture one eye, then the contralateral side of the visual cortex, the firing rates go down also.

John Beggs
Yes.

Paul Middlebrooks
The metric of criticality resurfaces before the firing rates catch up, which is really cool.

John Beggs
Yes, that is pretty cool. We're still trying to understand that. One of the things it suggests is that the distance of criticality is at least as important as firing rate. The brain gets that in line really quickly before firing rate comes online, back into its zone. Yes, we don't know fully what that is, but I saw that paper and I reviewed it and I wrote the intro piece for it in *Neuron* 2019, Ma *et al.* I love that paper. I just thought, "Wow, he's really nailed it." He's really shown that if you perturb, the system will come back. They've got much better stuff since then. I don't know if you've seen this, but they were looking at a tauopathy model of Alzheimer's.

Paul Middlebrooks
Oh, no.

John Beggs
This also appeared in Neuron, I think, maybe last year. What happened is they have these mice and they express tauopathy. That causes them to lose their memories. Little mice, they only live for two years. What they're doing is they're looking at cell-to-cell correlations, and they're looking at firing rates, and they're also looking at distance to criticality. What they find is the distance to criticality, that metric predicts better than any other low-level metric looking at single cells.

When the rat will go bad, how bad the symptoms will be? It's an emergent biomarker that is a population signal that is a better indicator than anything you look at in terms of single-cell firing rate or pairwise correlations. I think it's very useful for, potentially, biomarkers in humans. You can record things with EEGs on humans or MEG.

Paul Middlebrooks
You originally got excited about this idea, and you're an assistant professor, and then you start getting the blowback. It seems like these days, you're still writing pieces defending, right? You're defending the idea.

John Beggs
I'm still in the fetal position.

Paul Middlebrooks
Okay. Where is your--

John Beggs
No, not really. Not really. I think this is part of the scientific process is to have skeptics and to throw stuff at you. I can tell you a story about Alain Destexhe. He is the gadfly. He's constantly biting me. Socrates called himself a gadfly of Athens. He was asking them questions. I think that was good for them, and I think it's good for me, too. He and I are friends. How bad can it be if he invites me to Paris twice to give talks, and he takes me out to dinner, right? He writes things that are--

Paul Middlebrooks
Yes, but he'll knife you when he's doing it, right?

John Beggs
I don't know. He and Jonathan Touboul have been extremely nice. They would send me advanced copies of what they were going to put out, and they'd say, "Hey, look, I want your feedback on this." These guys are scientific gentlemen. They'd send it to me, and I'd say, "Well, I disagree with this. Blah, blah, blah. Blah, blah, blah." They'd incorporate some of my comments, but not all of them, and then I'd write a rebuttal piece or something like that. I think that the field has benefited from the blowback.

I have benefited from the blowback. I have grown in my appreciation for the subtleties and the nuances of criticality by the things that Alain and Jonathan have challenged me with, and others. There have been others out there as well that have said, "Hey, how do I know this isn't that?" That, in part, led to all this proliferation of measures. How do I really know if I'm critical or not? That's still going on even now.

Paul Middlebrooks
One of the issues that I've thought about, it seems like, because there are a lot of different measures now, if I want to find criticality in my data, I can find it by hook or by crook, and that worries me.

John Beggs
Oh. I would disagree with that. One of the things that we did very early on is, for example, we take the power law distribution and then we'd shuffle the data in time, and show that it would not produce a power law distribution. I think you should always have some kind of control measure. If you shuffle your data in time, your criticality measures will, at the very least, drop in quality. Your distance to criticality will go up. That's something that Keith Hengen does, that's something we do, Woody (Woodrow Shew) does not, a lot of people do that.

They're constantly looking at it. I think you have to be rigorous about it. Now, if you have too small a dataset and you shuffle, and then it looks the same, then I would say, "You have too small a dataset." Then you're just not really in the game to play criticality. If it's looking like there's no difference between the actual data and the shuffle, then let's not talk about this, but you have to have numbers large enough to play this. That's why, as people record more and more neurons, it's more of an option, I think, to look at.

Paul Middlebrooks
Where are we right now in the critical brain hypothesis? Are we in the heyday now, or are we in maximum criticism?

John Beggs
[laughs] Peak. It's passed.

Paul Middlebrooks
You think it's passed?

John Beggs
I don't know. No, no. Here's the thing. A lot of the work that's happened roughly up to this time has been, I would say, more fundamental science, where you're looking at the scientific question of, is it viable to say the brain might operate near criticality? What does that mean, scientifically? How do we measure it? How do we verify that? How do we make sure that we're not talking rot? Then, once it gets accepted to at least some degree, then medical people start testing it to see if it's related to neuropathologies.

In 2022, there was a really nice review by Vincent Zimmern in *Frontiers,* and what he did is he just basically looked at the medical implications of the critical brain hypothesis. He said, "Okay, it could be relevant for schizophrenia, for depression, for Parkinson's, for epilepsy." He just went down the whole list of things. That paper, even though it was publish in '22, at least on Google Scholar, has something like almost 180 citations now.

I think what's happening is clinical people are picking this up and they're beginning to say, "How can I find out if, once somebody has a bout of depression, do they go subcritical? How can I find out if someone is showing signs of schizophrenia?" Maybe they're, in some ways, departing from criticality. I don't know. I'm not a clinician, but I'm delighted to see that people are picking this up and starting to show interest in it. That might be the second wave of this stuff. It would be much less focused on the fundamental science of whether-- In some sense, when they start using it, at least in their minds, they think it's settled. That it's good enough to at least try as a hypothesis for diagnosing things.

Paul Middlebrooks
I'm trying to think also about causality in criticality. Is criticality causal or is it epiphenomenal? How do we think about that in terms of function?

John Beggs
It could be epiphenomenal, that's possible. That it's just something that happens, it's like noise that comes out of a-- You use a radio, and then if you were to look at just one resistor in it, measure the voltages, you'd say it's just noise. It could be a byproduct of what the brain is really doing. Now, if that were the case, though, then you wouldn't expect it to show such nice homeostasis, so when you perturb it, it comes back.

It's an epiphenomenon. Why does the brain care about it more than firing rate? I would offer those as things. Now, the truth is we don't have enough data to know how to answer exactly on that question. It's a valid question. I think we need more data, I think we need causal interventions. If I perturb criticality, does it always come back? What Keith and his colleagues did is they basically did, what I would say is a negative perturbation. You take it where it's critical, then at subcritical, and it comes back. What if I do a positive perturbation? Somehow, I throw it into supercritical state, does it fight its way back down? Is it symmetric both ways? I think that's a little bit harder to test.

Paul Middlebrooks
Why is that harder to test?

John Beggs
Because it's a little more unnatural. At least, the first way that it occurs to my mind is you put a bunch of bicuculline on the brain, and so now you knock out the inhibition, and you're going to have seizures. If you have a bunch of seizures, does the brain fight its way back from seizures to coming down to critical--

Paul Middlebrooks
It has to in order to survive, right? You can't stay in that state, so it has to go somewhere.

John Beggs
Right. Now, maybe it does, but maybe the mechanisms of coming back down are very different from the mechanisms of coming back up. They could be different, but I agree with you. If you're going to survive, you've got to somehow overcome seizures. I don't know how it does that, and there might be different time constants on the feedback, coming from positive and coming from negative. There's interesting questions, and you could even come up with a generic model and phenomenologically, is it better to come back fast from below, or do you want to come back faster like-- Here's seizure territory, if this is critical here and you're up here, that's seizure territory.

You might expect this is a more urgent thing that you've got to deal with, so come hell or high water, I'm going to get that thing back down. Whereas here, you don't want to come back up too fast, maybe you'll screw up your wiring and you'll erase memories or something, so you do it over 48 hours. I don't know. These are still open questions. They're prompted by the idea that maybe the brain is homeostatically tuned to operate near criticality, and then maybe that's a golden age of research where people are taking it-- It's no longer under scrutiny like, "Is this a real idea or not?" Then, people are using it as an operating, working hypothesis to investigate clinical things.

Paul Middlebrooks
I'm not supposed to be biased as a scientist, but I love the idea of the homeostatic setpoint being criticality, essentially, but then how? How does that setpoint get set, and how does it achieve it? That's the big question, right?

John Beggs
Yes. Keith has some ideas about it. You should have him on this podcast as well. What I can say, probably without giving too much away, is that they've done a lot of really high electric count recordings in mice. They have also done some really sophisticated models, working with people from the Allen Institute. I think they'll be coming out with some statements about how that works.

Even their 2019 paper, their Ma *et al.* paper in *Neuron*, they did some decent models, and I should say Ralf Wessel is on some of this stuff as well. Ralf Wessel, he's a physics guy out of Wash U in St. Louis. One of things that they found was they created a number of different classes of models. What they found was that the classes of models that matched the data best were ones that had inhibitory neurons basically driving the recovery. That's a big clue. You could think that the excitatory neurons drive the recovery, but it's--

Paul Middlebrooks
You mean by settling down or speeding up, or whatever, as an inherent process? The excitation inhibition balance seems to be critical-- forgive the term here, but it seems it's always a critical parameter in models. What you're saying is it's actually the inhibitory neurons that's allowing the system to get back to that criticality?

John Beggs
Right. That's at least the latest that I know about it. I'm not an expert in this, but I think it's fascinating, and, yes, they will probably come out with some stuff that indicates exactly what type of inhibitory neurons, and what are the dynamics of this. I'm sure they've got a model of how it works. That's really great question is, how does it come back, and is it working the same way above criticality and below criticality? Do those inhibitory neurons only pull up, and you have different neurons push down? I don't know. There could be distinct circuits for that.

Paul Middlebrooks
That's one of the fun things is there are so many unanswered questions, so much work to do that it's ripe. A lot of fruit to pick.

John Beggs
Yes.

Paul Middlebrooks
As an aside, we're talking about brains, but I want to expand the conversation. For example, if I record my mouse wandering around in a box, I can look at the kinematics, the behavioral data. Some of the tracked positions and some of the metrics we used have these long-range temporal correlations. The behavior itself is scale-free, parts of it anyway, in addition to the neuronal activity. Stepping way back, I'm part of a discussion group, and we've been talking about evolution lately. I think people talk about evolution as if it's a force, but really it's like a description of what has to happen for things to survive. Is it right to think of criticality in that same way, that to be alive is to be-- You have to be at criticality, just to be a living organism?

John Beggs
That's a really interesting question. I don't know enough to speak authoritatively about that. I certainly have some opinions, I'll offer them in a minute, but a really nice review is by Mora and Bialek, and I think the title of it is something like *Are Biological Systems Poised at Criticality?* They look at things like the immune system. There's work by Magnasco *et al.* on the hearing system. It's poised at a critical point at a Hopf bifurcation. It's amplifying things, but not too much, so it's just hovering at that spot. They look also at the neuronal avalanche stuff. They also look at flocks of starlings, or I should say murmurations of starlings, or a murder of crows.

Paul Middlebrooks
Oh, yes.

John Beggs
Crows don't act like that. There's an English word for each type of flock of birds [laughs], which I've already exhausted my knowledge. Anyway, they show that there are lots of things, also like bacteria and different things that swim. Schools of fish. A lot of these things seem to organize at a point where they're operating near a critical point. Why is that? It might be for information processing purposes. Let's say if you're in a murmuration of starlings and some hawk is coming down, if you had maximum susceptibility, then that would allow-- Let's say a bunch of these starlings notice the hawk is diving.

They turn, and then the turn propagates through the entire flock very quickly, best at the critical point, and then they break up. For survival, it might be useful. If information processing is important for evolution, then being near the critical point would be something that evolution could push toward in multiple ways. The way I'd like to think about it is it might be something like wings. Wings appear in dragonflies, bats, birds, of course, and even some snakes and lizards.

Some lizards have certain ways of flying, and some snakes, I don't know, in Southeast Asia, they make their bodies really flat, and they fall from trees in a certain way. Flying squirrels, too. What the heck is going on? I think there's some evolutionary pressure to say, okay, let's use Navier-Stokes equations to gain some distance here, and we can fly. We can fly if we have the following things. You get some surface that catches the air and maybe creates lower pressure on top, and now you can fly. Evolution independently arrives at this in multiple organisms at multiple epochs in evolution because it's just there in the laws of physics.

Paul Middlebrooks
It's best. It's the best thing to survive.

John Beggs
It works. If we were to take this as an analogy, you might think that, okay, the heart of criticality lies in the laws of physics. There's something about

setting up a brain or a flock of animals or the ear or a swarm of bacteria, setting them up at this point where information that comes into the system is not extinguished. It's preserved for as long as possible. It lingers, and you don't over-amplify the response, and you don't over-damp the response. You just let the information echo within the system for as long as it can before it dies out with a power law tail. Now, new information that comes in, these things can be jointly processed. I don't know, but it could have something to do with the laws of physics. That's my generic hunch.

Paul Middlebrooks
It has to have something to do with the laws of physics.

John Beggs
Yes. In that generic sense. Let me put it, maybe a little bit more forcefully. You're challenging me to say a little bit more than that. The laws of physics have very interesting properties like symmetry. If I do an experiment in misorientation now, I rotate, it's the same there. If I do it here, I translate it over here. It's the same. If I do it yesterday, or I do it today, or tomorrow, it's the same. These symmetries are related to very deep things in physics through something called Noether's theorem, which says that we're always going to minimize the action. The action, in typical sense, it's the difference between the kinetic energy and the potential energy. The difference between that integrated over time, you always want to minimize that.

Paul Middlebrooks
That's a weird law. I've always found that strange.

John Beggs
It is weird. What's funny is I'm talking to all these people in my department and others, and I say, "Why is the action always minimized? Why is that?" They go, "Oh, we don't know why. We just know it always is." That's something in the laws of physics. There's a deep connection between minimizing the action, which we find works for strangely classical mechanics, quantum mechanics, and general relativity.

Paul Middlebrooks
Another way of saying that is you always take the shortest path, right?

John Beggs
Correct. For people who may not know what the minimization of action is, if they listen to the podcast, it's always taking the minimum path. This minimization of action is a fundamental observation of physical law. It leads to these symmetries. These symmetries lead to conservation. Conservation is like the branching ratio. If the branching ratio is one, things are conserved. If the branching ratio is greater than one or less than one, things are not conserved.

Conservation, if you have a system that's conservative and it's wired up in the right way, so it's some population, you've got links between all these guys, it'll be critical. It obeys this idea of conservation, which is linked to minimization of the action and linked to the symmetries that we find. There's an interesting paper by Lin and Tegmark, and the title of it is something like *Why Does Deep and Cheap Learning Work So Well?* This is related to deep learning.

Why the heck can we get away with just firing up a network with all these layers and just training it on a bunch of examples? Now it seems to have the concept. It's not just memorized everything that it's seen; it now has the concept of what a cat's face is, and it can see new versions of the cat's face that it's never been exposed to before, and correctly identify it. It really seems to be capturing the concept, almost like Plato's forms. What is the essence of cat-ness? It gets it.

Why does it do that? One of the things that they argue there is that because there are these symmetries in nature, because things are rotationally invariant, a lot of things are like that, translationally invariant, there are all these symmetries. Basically, when you want to write an expression to explain the dynamics of the system, the equation that you write for the dynamics is very compact because of these symmetries. They say, "Oh, the Hamiltonians are very compact." They're small, and we know this is true.

If you want to explain the whole world from quantum mechanics to general relativity and classical mechanics, you could write three equations on a t-shirt. Schrödinger's equation, general relativity, and then minimize the action. You cover almost everything there is to cover. Why is the world so compact in that way? It's related to conservation and symmetries, and that allows us to learn things. I think also, my hunch is, is that it's related to criticality somehow.

Paul Middlebrooks
I had David Krakauer on the podcast a few episodes ago, and he has written this book. I asked him because he talks about broken symmetries a lot in the book.

John Beggs
Sure.

Paul Middlebrooks
I read the book, and I'm like, "Yes, broken symmetries," and I realized I didn't know what he was talking about. It's one of those things where you

read the phrase over and over and then you realize, "Oh, wait, I actually don't know what he means." What you're talking about with the symmetries are idealized physical laws, but anything interesting that happens is a broken symmetry, so there's this conundrum.

John Beggs
Sure. I'm glad you brought this up. All right. If everything is perfectly symmetrical, often it's totally boring. For example, at the Big Bang, you get a certain amount of anti-matter and matter created. If that's totally symmetric, if it actually is perfectly balanced, they cancel.

Paul Middlebrooks
Nothing--

John Beggs
We have nothing. The brain, we talked about E/I balance. You've got excitation and you've got inhibition. If they are perfectly balanced at all moments in time, we have no activity. The thing that happens typically is there's a pulse of excitation, and it's followed by a ring of inhibition afterwards. There's a slight asymmetry, and that asymmetry allows activity to go through. The same is true with anything at a phase transition point. At a phase transition point, typically what happens is, let's take this.

Let's say you have a bunch of molecules and they're in a volume here, and they're just bouncing around, so they're in a gas. Now what happens is, basically, you have for every little voxel in there equal probability that there's a molecule in there. In some sense, it's isometric, it's symmetric. Now, if I slowly condense this thing into a fluid, what you get is fluid appearing on the bottom. Now there's a much higher probability that things on the bottom are going to be occupied than things on top. You've broken the symmetry.

Paul Middlebrooks
You've reduced the entropy as well, correct?

John Beggs
Yes. This is related. Now, what you do is you break the symmetry, but right at that point where you're breaking the symmetry at the gas-liquid phase transition, that's where you get power laws. That's where you get all the interesting stuff, right at the point where you're breaking symmetry. I'm not going to contradict Mr. Krakauer, head of the Santa Fe Institute. I know him and we've met before, and he invited me out once, and it was great. I think that he is absolutely right. You got to break the symmetry. He's totally right. Now, the symmetries, though, are fundamental and interesting, and you need them in order to get close to this phase transition point.

Paul Middlebrooks
We are far from thermodynamic systems, let's say just a human being, right? We're far from thermodynamic. We're constantly, is the right way to think about it, that we are shooting for symmetry, but we're far from thermodynamic equilibrium, so we're never going to get there. It's just a constant battle to be at that homeostatic set point. Man, that was a mouthful.

John Beggs
[chuckles] Yes. There are a lot of ideas in there, and people are grappling with them. These are the things that we're all talking about. First of all, I'd say, let me just make one distinction. There's a difference between what I would call an equilibrium model and a non-equilibrium model. An equilibrium model is a model where, let's say you've got some system and now you just let it cool down. You're not driving it, you're not adding anything in there, you're not heating it up, you're just letting it cool down. That type of model is very useful, and it can approximate many things in the brain.

For example, Bill Bialek and Elad Schneidman have used maximum entropy models to map the Ising model onto the brain. You can get a lot of traction with things like that. It's really wonderful work and it even points toward criticality, but the brain, as I'm sure they would agree, is fundamentally a non-equilibrium system in the following sense. It's constantly receiving inputs. This would be like your system and you're not letting it cool down, but now you're heating it up. You've put it on a hot plate and now you're turning this little hot plate up and you're driving the system and when you drive it, you can get boiling or you can get roll cells of the fluid formation.

There are all kinds of structures that can appear when you drive it, that don't appear if you don't drive it. You've got equilibrium models and you've got non-equilibrium models. You can approximate the brain with an equilibrium model, and the Hopfield network is one of those things, and maximum entropy models are, but a non-equilibrium model is one that's being driven. I think if you take a look at a little patch of cortex, it's being driven.

Paul Middlebrooks
It's being driven also by itself, right?

John Beggs
Correct, both things. It's being driven by itself, and it's being driven by other inputs from other areas. You need a non-equilibrium model. If you have a non-equilibrium model, then it gets more complicated because people have made great progress in equilibrium statistical mechanics, but non-equilibrium statistical mechanics is not a settled discipline by any stretch.

Paul Middlebrooks
Why is that? Just the complexity of the puzzle.

John Beggs
It's very, very hard, right? It's super hard. Now, we have attempted to say these ideas of criticality at a phase transition point. We use a branching model that's a non-equilibrium model, but even more than that, it's being driven. Any patch of brain is getting driven. We've gotten into this idea called quasi-criticality, and that's essentially trying to accommodate this idea that a given patch of cortex is driving itself, as you point out, but it's also being driven by external inputs, and that's going to push it away from the critical point.

Paul Middlebrooks
What's the difference between quasi and near criticality?

John Beggs
You could have a classical system that's not being driven outside. It's just near criticality because maybe it's not at the critical temperature yet. I could have some system that has-- Let's take water. You take water and you put it at the right pressure, and now I bring it to some temperature. I basically can bring it close to the critical point. That's an equilibrium system. I'm not putting anything in. I've got this little cell, I've got water in there, and it's a mixture of gas and liquid.

Now, I can change the temperature, and I can either bring it right to the critical point or I can be slightly away from the critical point. Now, that's an equilibrium system that is slightly away from critical. Now, imagine I had that system and now I put a little hole in it, and now I can put water droplets in. If I take that system and I put it at exactly the critical point, and now I start squirting little water droplets in, that would be what we're trying to describe by quasi-criticality.

We're saying that through homeostasis and everything, what the network does is the network brings itself as close to critical as possible, but it's got a little leak [chuckles] and water's dripping in there. Now it's not really at the critical point anymore. It's close to the critical point, but the reason it's pushed off is because I'm squirting water in there. It's a non-equilibrium system.

Now things are happening that aren't quite right. It's definitely not perfectly symmetric. We're breaking symmetry, and that's where interesting stuff happens, but the picture of it being apart from criticality because it's driven is slightly different from the picture of it being apart from criticality because it's not at the critical point yet. The temperature's not there.

Paul Middlebrooks
I see. This is going to be a bit of a left turn here, but we're going.

John Beggs
Go for it.

Paul Middlebrooks
Why the cortex? Why not the whole brain?

John Beggs
That's an excellent question. First of all, I think that other parts of the brain could be critical. Sub-cortical parts could be critical. There's a really nice paper out by Miguel Muñoz and colleagues, and it was in Proceedings of the National Academy of Sciences maybe last year, something like that. What they did is they looked at neuropixel recordings that were put in striatum, amygdala.

Paul Middlebrooks
I'm afraid you're going to review my paper because I have striatum recordings, and this could be a huge conflict of interest on this podcast now.

John Beggs
[chuckles] Anyway, they record from all these different regions and then they have a way, not using avalanches, but using principal components and things, to look at these different regions. What they say is that a lot of these regions are really close to the critical point. There might be an argument for that. I don't know enough about that to rule one way or another. I'm just commenting just to say that some people would argue that many parts of the brain are.

Now, here's why I think it might not be that. There's nice work by Viola Priesemann and colleagues on neuromorphic computing, and what they do is they take this neural-like chip and they train it on many different tasks. They train it on complicated tasks, and they train it on simple tasks. Then they ask, after they train it, is it close to the critical point? What they find is that for really simple tasks, it's not close to the critical point, but for really complicated tasks or more complicated tasks, it is closer to the critical point.

Paul Middlebrooks
That's because you need the capacity to accomplish a more complicated task.

John Beggs
That's the thought. Let's say I've got a specialized circuit, and let's say all I need to do is keep a rhythm. There's this pre-Bötzinger complex in the brainstem, it just keeps a rhythm, and it's just got to keep you breathing at a certain level, and maybe gets inputs about your oxygen, and maybe breathe faster or breathe slower. That thing might be, maybe I'm wrong, but I have no reason to believe it needs to be critical. It just needs to be reliable.

Paul Middlebrooks
You don't want your heartbeat to be critical, right?

John Beggs
Maybe not right. Now, on the other hand, heartbeats do have long-range temporal correlations, so there might be something going on. [chuckles] You can't escape it. It's everywhere, but I wouldn't expect that the pre-Bötzinger complex has to be critical in the same sense that the cortex does. The reason why I think the cortex is probably most likely to be critical is because it has to simultaneously optimize multiple tasks. It's got to be good at transmitting information. It's got to be good at storing information. It's got to be good at dynamic range. It's got to be good at computing all these things at the same time.

A given patch of cortex, before you're born and exposed to the environment, you don't know what associations you're going to learn. You might learn that red means stop and green means go, but if you live in another country, maybe it's different, or maybe you learn in America, we're going to drive on the right side of the road, but in England, you drive on the left side.

There are a bunch of arbitrary associations that we all have to learn to get along in the world, and that cannot be pre-wired into the cortex. The cortex has to be generic. It has to be a generic computational unit that's largely specified by its inputs. If you have something that's not highly specialized, but needs to be ready to do anything decently well, jack of all trades, then being critical makes sense.

Paul Middlebrooks
I'm born, we've been talking mostly about like neuronal activity, but there's criticality and structure as well. You talking about being born there, you're massively connected when you're born, and then there's a pruning that takes place.

John Beggs
Yes.

Paul Middlebrooks
I bet you know the answer to this. Is that pruning? Does it get to like so eventually we end up at like a small world network, optimal state, and I know small world network is related to criticality, but is that pruning itself? Do we end up with a fractal critical structure?

John Beggs
This is something that has been investigated in primary cultures. There's a paper from I think 2010 by Tetzlaff *et al.*, and Ulrich Egert is on that, and one of the things that they do there is they grow these cultures. They basically take neurons from, let's say, rat hippocampus or cortex, I can't remember which. They enzymatically dissociate them. They're floating in a solution, and then they pour it down over a 60-electrode array. They record these guys from right after making the culture to four weeks later, something like that, and what they find is that generally the picture is, it goes through some pruning. It goes through exuberant connections, and then some of these things are pruned.

Paul Middlebrooks
They're not tasking the culture with anything. It's just letting it grow.

John Beggs
They're just letting it spontaneously be active and listen to itself. Over that period, it does eventually approach criticality. If I remember the cycle first, it goes super critical, which would make sense if you're overconnected, and then it goes subcritical, which would make sense if you've pruned a lot, and then it gradually approaches criticality from below if you start now learning and you strengthen those connections that have not been pruned, you get closer and closer to that. That's an interesting thing.

There have been models to look at that kind of thing. A really good question would be, does that happen in deep learning? You train a really deep network, and you find out, definitely they prune them, definitely the weights change, but does that recapitulate the same thing that Tetzlaff *et al.* found back in 2010? Does it basically act like it's super critical at the beginning, and then go subcritical, and then gradually approach from below? I don't know that. That's an open question.

Paul Middlebrooks
Since you mentioned AI, this is ostensibly a podcast about neuroscience and AI. I've only read about criticality in biological organisms in my studies thus far, but AI, you can turn the computer off, come back the next day, turn it on. There's no necessity of ongoing dynamics, and so in that sense, criticality is not part of that story. I suppose if you run the model. Maybe do you know, is there work in artificial intelligence looking at criticality and whether tuning to criticality improves model performance, for example, or generalizability?

John Beggs
There is a little bit of work on that. It seems that they are unaware of the work in neuroscience. In a sense, they're independently discovering it.

Paul Middlebrooks
That's good.

John Beggs
That's exciting.

Paul Middlebrooks
AI always unaware of neuroscience,

John Beggs
Always unaware. Whatever. There are a lot of smart people working in that area, and then they eventually hit on this. It may be that that's where you want to poise your networks for optimal training. In fact, I think there were two papers in *NeurIPS* that a colleague of mine showed me. I'm blanking on the guy's name, but I could probably send you the paper titles after this. They were looking at that.

Basically, they were looking at something related to that, which is, if you set up the connectivity matrices in such a way so that you have an input vector. You've got all these Xs that are ones and zeros, and you take the length of that input vector, and now, as it goes through the layers, you want to look at the length of that input vector. These networks will learn best if the input vector length does not grow and does not shrink over time. It's preserving, it's conserving.

Paul Middlebrooks
A branching ratio of one.

John Beggs
Branching ratio of one or a conservation principle. Something like that. That's very interesting, and I look forward to hearing what people have to say about those types of potential linkages between these areas.

Paul Middlebrooks
Just zooming out, I'm not going to ask you what your take on current AI is. What do I want to ask you about this? I didn't prepare any questions about artificial intelligence with you because I'm so wrapped in the biological world, but do you look at AI and think, "Oh, they need criticality?" What are your thoughts there?

John Beggs
They're doing really well without criticality. [chuckles] I'm very open to that kind of thing, and I'm increasingly interested in that. I have colleagues who are interested, colleagues who are getting me interested in it. We'll see how it all turns out.

Paul Middlebrooks
Like we mentioned before, you spend a good deal of your time in defense mode because of objections and criticisms to the criticality hypothesis. Are there any that are more worrisome to you, or what is the big obstacle right now that keeps you up at night?

John Beggs
At the beginning, all of them are worrisome to me. [chuckles]

Paul Middlebrooks
You have tenure now.

John Beggs
Yes, I do, so I can do anything. I think that what happens is, they should all be taken seriously. I think the latest one that I addressed was interesting work by Alain Destexhe and Jonathan Touboul. I don't know if you got to see this or not, but essentially what they said was the following. Maybe it would be instructive for me to go over it.

They said, "Hey, look, you think that satisfying an exponent relation is important for criticality? You think that fractal avalanches at different sizes so that they can all be collapsed and look like the same shape at different scales. We call that avalanche shape collapse. You think that avalanche shape collapse is really important, but guess what, John, we can get this with an Ornstein-Uhlenbeck process."

Paul Middlebrooks
What's that?

John Beggs
I'll simplify it. Let's say we're flipping a coin, here's our timeline. It goes like this. Every time you flip a coin and you get a heads, you go up one, and

time is like this. I go up every time I flip a head, and every time I flip a tail I go down. Now, this thing's going to go up and down, up and down over the line. You can look at the times where it crosses over the line from when it goes up to when it comes down.

You can look at those excursions, and you can average them over many scales, and they're fractal. I've done it. I've created a program in MATLAB and ran it a billion times, and sure enough, I get superb avalanche shape collapse. I get good avalanche power laws. They satisfy the crackling noise relation. That's all wonderful. It's a fair coin.

Paul Middlebrooks
What you're talking about right now is in *Frontiers in Computational Neuroscience* 2022, right?

John Beggs
Right.

Paul Middlebrooks
Addressing skepticism of the critical brain hypothesis.

John Beggs
Yes, exactly.

Paul Middlebrooks
I'll point to it.

John Beggs
You've seen it. That is a fair coin. It's a critical process. It is exactly a critical process. What they were basically arguing is, how do we know that this isn't what's going on in the brain. You just get a random walk, it's a coin flip. We get random activity and neurons and they created a model and it matched all this stuff.

Paul Middlebrooks
However, if you randomize your neural data, then you get different exponents, you get different numbers, which is--

John Beggs
That's a good point. That's an interesting thing. Let me just address, first of all, their stuff, and then maybe we can go back and address this other business of shuffling the actual data. In terms of taking their model seriously, the coin flip model or the Ornstein-Uhlenbeck process, which is like that, except it's got a potential well. It's a little bit more likely to come back to the center than a coin flip. I had to admit, it satisfied everything. Then I went through a period of shock and mourning and whatever.

Paul Middlebrooks
Seriously, when you get news of that, you're like, "Oh, another one I have to deal with."

John Beggs
Right. I ran all the simulations. I'm like, "Holy smokes, what the heck is going on?" Then I thought about it more and more, and I'm glad they challenged me. I thought, "What if I do admit that it's critical? What if I do admit that the coin flip is critical? Because in a sense, it's at a point of symmetry because the chances of getting heads is exactly equal to the chances of getting tails. If you slightly break the symmetry, if you slightly make the heads more likely than the tails, then you lose the power laws, or you begin to lose the power laws. It's a phase transition between equal probability to something that's not equal. Right there is where you get the face transition.

Now, here's what I convinced myself of. The coin flip is critical. The Ornstein-Uhlenbeck process might be critical, but it's a different type of criticality. This is what I said, the criticality that we get in the brain is a result of the collective interactions of all the neurons, so it's emergent. Like the criticality you see in a flock of starlings, it's emergent, but the criticality that they're claiming might work for the brain is just the coin flip.

If it is just the coin flip, and so we're being driven by random stuff outside. If I now go into the brain and I cut the connections, and it's all from being driven by a random source from outside, that's like an Ornstein-Uhlenbeck process or a coin flip. If my neurons are just responding to that, some of them go on and some of them go off. If I cut the connections and I still get power laws, then I agree with them.

The brain is doing nothing but reflecting the random statistics of the input stream it receives, but we know from experiments that if you go in and you block with AP5 and CNQX and other things like that, if you block synaptic connections, you don't get power laws anymore, you don't get avalanche shape collapse, these things fall apart. What seems to be essential for the brain is the neurons and neuron couplings.

The same thing is true in a model. If I go into a model and get it to run a criticality, I've got integrating fire neurons, all the standard stuff, and now I cut the connections, it also falls apart. That's not what I think is going on. Although, how they always say, but see Ilya Nemenman. Ilya Nemenman at Emory would tell you a contrary story. He's got a good argument. It continues to go on.

His argument is basically, the criticality is coming from some other region over here, and it's very low dimensional. I said, "What if I create a culture and I have two wells, and now I cut the connection between them? According to your hypothesis, that should severely dampen the criticality over here. I have done that experiment, but this is something we could look into. Is it being driven by an outside source or is it intrinsically generated by the network itself?"

Paul Middlebrooks
Do you think of the different Brodmann areas of the brain? Let's just say prefrontal cortex as a whole needs to be a criticality, visual cortex itself needs to be a criticality. They can be at different phase, different levels near criticality, and then they have to coordinate, or is it just across the whole cortex? How do we [crosstalk] if I add one neuron to a population of neurons and it destroys criticality, right?

John Beggs
I think actually there is a gradient of criticality in the brain, and there's some data to suggest this. There are people looking at autocorrelation timescales. Let me see if I can remember this paper. I have a good book here I can look it up.

Paul Middlebrooks
Oh, there it is.

John Beggs
I forget what I said though, when I wrote the book, but then I forget what the references are, so I got to look at this.

Paul Middlebrooks
Oh, there's a lot in there. I'll just say it's such a readable-- It seems to go fast. It doesn't seem like a long book, but there's a lot in there.

John Beggs
Oh, good. It went really slow for me. [chuckles]

Paul Middlebrooks
I bet.

John Beggs
Anyway, Murray et al, they look at the autocorrelation timescale. Anyway, this book, by the way, you don't have to pay for it. It's also open access, so you can get all the PDFs for free if you go to the MIT press website, but Murray *et al.*--

Paul Middlebrooks
How was the book's reception then first of all, before you go into this? How has the criticism been of the book?

John Beggs
I've not received anything that said it was awful or whatever. I think mostly everyone who's gotten in contact with me has said they really liked it. So far, it's been pretty positive.

Paul Middlebrooks
Were you asked to write it, or is this an endeavor that you chose?

John Beggs
They asked me to write it.

Paul Middlebrooks
They did?

John Beggs
Yes, they invited me and I was thrilled. I'm like, "Yes, great. I'd love to do it."

Paul Middlebrooks
It's a cool book.

John Beggs
Thank you very much for your help with it. Murray *et al.*, what they do is they are a bunch of primate electrophysiologists and they all get together and they record from different parts of the primate brain, and they look at autocorrelation, time constants. What they find is that in the frontal cortex, the autocorrelation time constant is longer, and in the visual cortex it's shorter. There's work by Viola Priesemann, and I've been itching to see her get it published. She and Jens Wilting have come up with a very good way to measure the branching ratio at short time intervals. You should also look at their work. I can--

Paul Middlebrooks
This is MR. Estimator.

John Beggs
Yes, exactly.

Paul Middlebrooks
MR. Estimator.

John Beggs
Yes, exactly. That they have applied to different brain regions, and they haven't published that, but at least as far as I know, I talked to her a little while ago, and she says, "No, we haven't published it yet, but I want to get to it." She got distracted by COVID and then published amazing work in science on that. Now, she's coming back to neuroscience. Anyway, they had results that more or less mirrored the Murray *et al.* in terms of time constant. In other words, the branching ratio is closer to one when you get to prefrontal cortex, and it goes a little bit further away from one as you go back.

Paul Middlebrooks
Subcritical?

John Beggs
It's subcritical. Yes, it tends to be subcritical. One of the things that she consistently says is that you rarely see the branching ratio go over one. It's almost always close to cortex, but not over the line. Because if you're over the line, then maybe you hyperexcite and you risk seizure.

Paul Middlebrooks
What would that mean? That would mean like evolutionarily, a different species who maybe doesn't have a prefrontal cortex like we do, their most anterior or newest brain region should be close to one, and then everything else can be a little bit under. The longest time range should be the most abstract concept related brain regions, right?

John Beggs
Exactly. That's what I would predict. That's what I think should happen.

Paul Middlebrooks
A snake, it will still have some part of its brain that is near one.

John Beggs
Maybe. I haven't seen snake record. I've seen turtle recordings. You can talk to Woody about that and zebrafish recordings. Those both look like they're near the critical point.

Paul Middlebrooks
That makes our prefrontal cortex look a little less impressive is if turtle cortex has a branching ratio, then that's the best brain area.

John Beggs
Yes. It's not all about criticality, though. We've got a pretty big prefrontal cortex, and they don't. It's receiving a wider variety of inputs that have been processed through many more layers. Of course, you know this, you're a [chuckles] multi-neuron electrophysiologist, but anyway, yes, I agree with you. If they're all operating near the critical point, what do we gain from that? I don't know yet. There are still all these other questions that have to be explored.

Paul Middlebrooks
What's your level of confidence these days about the critical brain hypothesis? Is that an all-time high? You seem like a humble individual, so I think that you probably have a measure of self-skepticism as well, and how has it changed over time?

John Beggs
I think we should always be skeptical of things. We should always be questioning, how do I know this is true? More or less-- What is it? The Dunning-Kruger curve, have you ever heard of this?

Paul Middlebrooks
Yes. Where are you in that? [chuckles]

John Beggs
I've followed the Dunning-Kruger curve. In other words, the first time I saw our data, I was convinced that the brain and all brains were absolutely critical.

Paul Middlebrooks
You have to have that to pursue something, right?

John Beggs
[chuckles] Right. I was totally convinced. Then I went through a really low period where I'm like, "Oh my gosh, I don't know, squat." [chuckles] That was before I got tenure. Unfortunately, we published a little bit on time, then I got tenure. Now, maybe-- I don't know. Hopefully, I'm trying to have the right mix of skepticism and confidence. I'm confident to the extent that it's being picked up by a lot more people. It's producing seemingly useful results. It's a great joy to see someone like Woody Shew go so far with it, someone like Marco Capelli, someone like Keith Hengen. There are a lot of people out there who are playing with it. Dante Chialvo, he knew Per Bak long ago, so he's really old guard.

He had models about the critical brain when I was still in diapers, almost. [chuckles] He's really one of these people who's been around for a long time thinking about these things and thinking deeply about it. I think it's good to see it being picked up by more and more people, the diversity of things, but I think everything can be oversold.

Paul Middlebrooks
That's the thing.

John Beggs
You don't want to do that.

Paul Middlebrooks
Yes. It sounds like a theory of everything for the brain, like the free energy principle. Like, here's the answer, and it sounds sexy and it sounds cool. Do you resent that at all? There seems to be a lot of hype around it.

John Beggs
Being sexy is not bad. Being hyped is not bad as long as you're grounded.

Paul Middlebrooks
I wouldn't know, man.

John Beggs
[chuckles] Right. I should retract everything. I don't know either. I've been married to the same woman for a long, long time, and I'm fine with my sexiness as long as she thinks it's okay. Here's what I would say, as long as you're grounded in testable hypotheses, as long as it can be refuted, and this can be refuted, right?

Paul Middlebrooks
How can it be refuted? Going back to what I said earlier, I feel like I can finagle my way to find criticality signatures in my data if I just-- It's like p-hacking almost. That's what I worry about.

John Beggs
I'll give you a couple of examples where it's not critical. Go to the cortex even, and take a look at recordings from layer II, III versus recordings from layer V.

Paul Middlebrooks
I've got to send you my data afterward, because I'm going to disagree with what you're about to say, but go ahead.

John Beggs
That's fine, but I'll at least say, at least in some people's data, layer V doesn't produce avalanche shape collapse, and it doesn't produce satisfactory crackling noise relation. That at least gives me comfort that it's not always everywhere.

Paul Middlebrooks
Why do you think that is?

John Beggs
Maybe you'd be able to correct me on all this. Here's my naive idea of what's going on based on what I've read in other people's work. Layer V might be doing something different from layer II, III. Layer II, III, as we know, is where all the cortical connections travel.

Paul Middlebrooks
Getting a lot of input from other areas.

John Beggs
Basically, if you want to go from one mode of cortex or one region of cortex to another, to another, to another, you could travel along layer II, III

connections and go through the visual stream, for example. Layer V is outputting to targets that are often distant or subcortical. One of the things that at least some people have found, there's this paper by Peters *et al.*, I think it was 2014 in *Nature*. What they saw was that there's this increasing orthogonalization of the outputs of layer V so over time, they basically get separated. You can imagine you're driving a car, and you want to be sure that when you put your foot down you definitely know the difference between the gas pedal and the brake pedal.

You don't want those things to be neurally overlapped, where 50% of the time, you hit the wrong pedal. You want them to be completely orthogonalized. The output of motor cortex, if it's coming through deeper cortical layers, that's going to perform a different task. That's orthogonalizing the outputs. I wouldn't expect that process to be critical, but if you're in some deep network that's basically processing all the way from on off cells to edge detectors to face receptors in the fusiform gyrus, then you probably have a different type of goal. Your goal is to allow information to travel all the way through with minimal loss. Then you want to branching ratio close to one, then you want to be close to criticality, et cetera.

Paul Middlebrooks
Maybe a layman's simple way of saying that, and somewhat the way that I think about it is if your job is to do something, enact something, you might go away from criticality. Whereas if your job is to receive, process, and transmit for further processing, you might be better served near criticality.

John Beggs
Yes, that's my intuition for it. Now, there are papers where people talk about focusing on a particular task, and as they focus on-- There was a paper in *NeuroImage*. I don't know, maybe within the last year where they gave people audio and visual tasks. What they did, according to the paper, I haven't read it carefully, but the abstract at least said that when they focused on a visual task, their visual cortex came closer to criticality. When they focused on an auditory task, that got closer to criticality. These are the types of things that you'd want to look at to see if the theory can be refuted. Is that what you'd expect, or is that not what you'd expect?

I think there are lots of ways you could refute it. Why would there be homeostasis? You wouldn't expect that if criticality were a valuable thing. What's the relationship between IQ and proximity to criticality? There have been papers on that. I don't know how good they are, but certainly, we could talk about in the future. If criticality is important for information processing, then we would expect people who do better on behavioral tasks would be closer to critical. In those areas.

Paul Middlebrooks
In what brain area?

John Beggs
You got to pick properly. Maybe you give them a purely visual task, or maybe you give them a purely auditory task, or something like that. I can at least conceive of ways to refute this. A negative response would be that, "Guess what? They go in and out of being close to critical, and it has random relationship to the task that you give them." If there's just chance relationship, then criticality seems to lose on that. On the other hand, if it's statistically significantly related to when they focus, then maybe we got a different story.

Paul Middlebrooks
John, what have I not asked you that I should have asked you? What should I have asked you about criticality that I--

John Beggs
Oh gosh. I don't know. Oh, maybe here's something. What's the difference between homeostasis and criticality?

Paul Middlebrooks
Good question.

John Beggs
Can I just offer a little opinion?

Paul Middlebrooks
Hold on. John, what's the difference between homeostasis and criticality? [chuckles]

John Beggs
I would be happy to answer. I'm glad you asked. I know many people have come to me and said, "Criticality is just the same thing as homeostasis." My answer initially, when I heard that, I thought, "Wait a minute. I know they're not the same, but I had to think about why they weren't." Here's my answer to that.

Homeostasis is the process of returning to a set point. If you have a thermostat in your house, and you want to set your temperature at 68 degrees, if you open up a lot of windows and you lose a lot of heat, your temperatures are going to go down, and now the furnace turns on. Likewise, if you come from the other direction. Being at 68 degrees, it's just a point. There's nothing magical that happens in 68 degrees. Homeostasis is the process of getting to 68 degrees.

Now, what is criticality? Criticality is a point you can get to through homeostasis, but the critical point has special properties. That's different from homeostasis. Homeostasis and criticality are linked because it seems that the brain uses homeostasis to get near the critical point. When you get to the critical point, magical things happen. You have power laws, you have scale-free activity.

Paul Middlebrooks
Isn't the hypothesis, then, that homeostatic setpoint is at criticality for that particular reason, right?

John Beggs
Yes, I would agree. They're definitely related. All the beautiful work that Gina Turrigiano and colleagues did on homeostasis and synaptic scaling and things like that, that basically shows that the brain has a process of getting to some point, and it wants to return to that point. Whether you perturb it above or below, it's going to come back and get there. That didn't say much about the point itself. It said that there is something that the brain is doing to get to some location. What is that point?

Now, Keith Hengen has worked on-- He thinks he and Woody have a nice review that I think will be coming out soon that basically posits that criticality is the end goal of what homeostasis is driving at. That's related, but they're two different questions. What happens at the critical point is not told by the process of getting there. Once you get there, you get things like fractals, the Mandelbrot set. You get infinitely repeating patterns on multiple scales, information transmission that is with least loss. You get a correlation length that is infinite in the Ising model. That means all scales can communicate. That only happens at the critical point. The critical point has special properties, but homeostasis is the way to get there.

Paul Middlebrooks
Let's talk about internal reference signals. Cybernetics, or the thermometer example. You said it, I externally set it to 68, and then it's a machine, and that I've built it so that it goes to 68 with feedback.

John Beggs
Right.

Paul Middlebrooks
One criticism, people like Henry Yin, who's been on this podcast before, where cybernetics got it wrong, is that what they're missing is that the set point is, we generate it ourselves. It's self organized set point. One of the questions is then if criticality is this space that we want to be near, and therefore we have a homeostatic signal to get there, how do we even begin to think about how that occurs?

John Beggs
I guess I could imagine some of the mechanisms that Gina Turrigiano and colleagues have looked at. If the activity is too high, some way, somehow the neurons set realize that.

Paul Middlebrooks
You're going to die off if you're not in the right spot, essentially.

John Beggs
Right. It's got to do that. If the activity's too high, what happens is the neurons begin to pull their receptors now that they have less input from other neurons. The receptors are pulled in, or vice versa. If the activity is too low, they start inserting receptors into the membrane so that any glutamate that's out there, they just grab it up and then--

Paul Middlebrooks
You have to know about your neighbors, then, and their activities, right?

John Beggs
Yes. The neuron itself may have a set firing point. Somebody has probably worked on this. I just don't know the literature well enough, but I think that it may have a firing rate set point, and then it realizes if it's above or below that, it either removes receptors or adds receptors, and maybe those set points aren't established very early on. Maybe you have to get the network to evolve into a state where it's roughly mapping the external world correctly. Now it says, "It looks like we're running properly," and it locks these guys in. I don't actually know. Gina would know. She's someone you should have.

Paul Middlebrooks
[chuckles] It's going to become the criticality podcast instead of--

John Beggs
Definitely Keith Hengen, definitely Woodrow Shew, and Marco Capelli. I can put you in touch with any of these people if you don't know them. Ralf Wessel is related. There are a lot of good people out there doing cool stuff.

Paul Middlebrooks
It's just such a beautiful-- You're a physicist by background, right?

John Beggs
Yes, mostly, but I have to be careful. Not entirely. I got my undergrad degree in engineering physics, then I got a master's in engineering physics. Then I went off into the Peace Corps, and I said, "I really want to be a Professor." [chuckles] I came back, and then I actually studied neuroscience in a lab where I was patch clamping. I was patch clamping neurons and looking at time constants and-- Exactly.

Paul Middlebrooks
That looked like a joint thing there.

John Beggs
People might not know. People might think that everybody who does patch clamping smokes pot, but no, this is-

Paul Middlebrooks
That's right.

John Beggs
-basically breaking a seal through a little tube. It's hard to believe that, but by sucking the pressure, you break through and then you record from the neuron its voltage. You become electrically one with the neuron. [chuckles]

Paul Middlebrooks
Yes.

John Beggs
That's what I was doing. I was doing stuff that you might call it biophysics. I had this real strong inclination to see the world through physics. When I started studying neuroscience, I was always trying to ram neuroscience into a statistical physics package.

Paul Middlebrooks
As a physicist would do.

John Beggs
Yes. Even before criticality, I wrote a paper on a statistical theory of long-term potentiation and depression.

Paul Middlebrooks
I saw that.

John Beggs
That was my idea. Let's get this into some stat mech framework. My preferred framework is physics, but the most curious and interesting object in the whole universe is the brain.

Paul Middlebrooks
Oh, isn't it so beautiful? It's so beautiful.

John Beggs
It is.

Paul Middlebrooks
It's amazing. It really is.

John Beggs
It is.

Paul Middlebrooks
Oh, last question here. Maybe you don't have an answer to this, and that's okay. We got brains. That's cool. They're complex, interesting things. We have behavior, itself is complex, and we have minds. Is criticality related to our minds at all? Is there a link there?

John Beggs
Ah, yes. Again, you're taking me out of my field of knowledge.

Paul Middlebrooks
That's out of everyone's field of knowledge.

John Beggs
Yes. I would say that if criticality is related to information processing, then yes, criticality has got to be related to mind. Now, is it related to consciousness? That's another aspect of mind. There have been people like Giulio Tononi and Olaf Sporns.

Paul Middlebrooks
What do you think of integrated information theory? That's been in the news, and it's on this podcast, like this past episode and upcoming episodes and stuff. I'm sorry to derail--

John Beggs
No, it's good. I interfaced with their ideas a little bit earlier, where they had something called neural complexity. We did a paper on looking at their measure of neural complexity and comparing it to criticality.

Paul Middlebrooks
I use your code. I used that, not my code.

John Beggs
Oh, okay. I shouldn't say it's my code. That's Nick Timme. He was a graduate student. He wrote the code. He was very good. Nick's code measures neural complexity. It turns out that they peak at roughly the same points, but they're not exactly the same thing. You can create a model that has criticality but not neural complexity. Has different curves for these things, so they're not identical.

In that sense, some of the early ideas that Giulio Tononi, Olaf Sporns, and Gerald Edelman were coming up with that might be related to consciousness. Maybe it's related to criticality, at least in that rethink. I think we don't know enough about what consciousness is, but I certainly would expect that consciousness is related to emergence. That it's a collective phenomenon produced by large numbers of neurons interacting with the world. In that sense, maybe criticality will be relevant.

Paul Middlebrooks
IIT says that logic gates that aren't even on have consciousness, right?

John Beggs
Yes. I don't know enough about integrated information theory. I've got to study that. It's on my list of things I want to do, but I haven't kept up with it.

Paul Middlebrooks
[chuckles] John. Thank you so much for going down this road with me.

John Beggs
Thank you for asking me.

Paul Middlebrooks
Lots of twists and turns.

John Beggs
Sure. Thanks a lot.

Paul Middlebrooks
Keep up the good work.

John Beggs
Thank you.

[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 the transmitter.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 hear is a little slow, jazzy blues performed by my friend, Kyle Donovan. Thank you for your support. See you next time.

[music]

Subscribe to “Brain Inspired” to receive alerts every time a new podcast episode is released.