Michael Breakspear and Mac Shine explain how brain processing changes across neural population scales

Breakspear and Shine find a scale-free property of brain activity that is conserved across diverse species, suggesting that a universal principle of brain activity underlies cognition.

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

[music]

Michael Breakspear
I hope that we are starting to put down some of the principles of how the brain works that are just a trend, but are actually going to stand the force of time. Like some of the things that were achieved in physics in the early 1900s.

Mac Shine
The dynamical systems to me was such a bedrock, and it is now of all my work, because it says something a little deeper, which is, what could a neuron even do or a population of neurons do given the information that they have or what they want to do next or something. You can frame problems from a neuron's perspective. I found that so much fun to think about.

Michael Breakspear
Plasticity plus criticality in very general circumstances can lead to self-organized criticality in which you no longer need to tune anything, because the system will self tune.

Mac Shine
I think we need to keep that humility as we go through, and realize it's a really difficult thing to go across scales, but as Michael said, we need to take the neurological detail and the computational so what, and try to make them connect as many ways as we can. It's exciting to me.

[music]

Paul Middlebrooks
This is "Brain Inspired" powered by *The Transmitter*. What changes and what stays the same as you scale from single neurons up to local populations of neurons, up to whole brains? How do tuning parameters like the gain in some neural populations affect the dynamical and computational properties of the rest of the system? Hello, everyone. I am Paul. This is "Brain Inspired." Those are the main questions my guests today discuss.

Michael Breakspear is a professor of systems neuroscience and runs the systems neuroscience group at the University of Newcastle in Australia. Mac Shine is back. He was here a few years ago. Mac runs the ShineLab at the University of Sydney in Australia. Mac and Michael have been collaborating on the questions I mentioned using a systems approach to studying brains and cognition.

The short summary of what they discovered in an early collaboration, is that turning up or down the gain across broad networks of neurons in the brain affects integration working together, and segregation working apart. They map this gain modulation onto the ascending arousal pathway in which the locus coeruleus projects widely throughout the brain distributing noradrenaline.

At a certain sweet spot of gain, integration and segregation are balanced near a bifurcation point near criticality, which maximizes properties that are good for cognition, which you know about if you've been listening to this podcast over the past few episodes. In their more recent collaboration, they used a course graining procedure inspired by physics, to study the collective dynamics of various sizes of neural populations going from single neurons to large populations of neurons.

Here they found that despite different coding properties at different scales, there are also scale free properties that suggest that neural populations of all sizes from single neurons to whole brains, can do cognitive stuff that's useful for the organism. They found that this is a conserved property across many different species, which suggests it's a universal principle of sorts of brain dynamics in general. We discuss all that, but to get there we talk about what a systems approach to neuroscience is, how systems neuroscience has changed over the years, and how it has inspired the questions Michael and Mac ask.

Thank you for being here. You'll find links to Michael and Mac's information, and many of the papers that we discuss in this episode in the show notes where I also link to a few good reviews for an overview of this approach to neuroscience. At braininspired.co you can also find out how to support this show through Patreon. Thank you Patreon supporters. All right. Here are Michael and Mac.

[transition]

You both are in the business it seems, and you can correct me if I'm wrong, of studying brains and cognition across scales, across levels. That is a systems approach however, and when people ask me this, it's hard. I don't like to be pinned down to one certain phrase. Like when someone asks me what I do, I generally say I'm a neuroscientist, but something feels unsatisfactory about that.

However, Mac last time, actually the episode title, which was three and a half years ago, was Systems Neurobiology, because that's what you say or have said that you do. Michael, at least on your website it says you do systems neuroscience. I'm wondering what the difference is between systems neuroscience and systems neurobiology. Michael, you want to take a crack at it first?

Michael Breakspear
Thanks Paul. I think systems neuroscience is just a very broad umbrella term. You can start with detailed neurobiological descriptions of the brain, and start stripping away some of the details. In our group, we're searching for unifying principles and underlying theoretical frameworks for describing how the brain works. Sometimes we strip away a lot of detail, and we don't leave that much of the neurobiology left.

It would seem to be doing a disservice to use the term neurobiology if we're studying very, very simple mathematical systems that somehow relate back to what we think is happening in the brain. That degree of abstraction, and that search for underlying principles, that might relate something that we're looking at in the brain to other systems that aren't the brain living systems, or even complex behavior and non-living systems. Maybe that lack of bound is why I still use this term systems neuroscience after 25 years of starting to use it.

Paul Middlebrooks
Mac, it sounds like what Michael does just encompasses what you do, and then some more. Do you agree with how he describes it?

Mac Shine
I think that's how I think about it. I think of systems neuroscience as this really big broad space that defines a perspective that you would take on the system, the ways that you'd like to frame your questions, because of the ways that you'd like your answers to look. I think the reason that I pushed on the neurobiology part back in the day, I probably do so less now-

Paul Middlebrooks
Oh, okay.

Mac Shine
-that I've come to learn just how broad of a church system neuroscience is, was really to differentiate the fact that I just at the end of the day really cared so much about the biological implementation level details than the math side, which I wasn't coming from.

I think a little bit imposter syndrome. I didn't want to come across as a card carrying computational neuroscientist with a background in applied mathematics or physics, but rather to say, I want to take really seriously the complexity of all the biology, and then think how on earth could that give rise to something as interesting as cognition or attention or consciousness or something.

Really starting from that point, to differentiate it from other groups that I saw that were really advancing the computational side, I wanted to push on that. I think in retrospect it is, I think all encompassed by this term system zero science in general.

Paul Middlebrooks
What is the systems aspect? Mac, you and I were just conversing a couple weeks ago, and you said you really like to, I think you said geek out maybe. I'm not sure if that's a phrase you would use, but on the neurophysiological details and the biophysical details, and the cell types and all that jazz. In that sense it's almost like a bottom-up approach.

That belies that actual systems approach. Then there's neuroscience neurobiology, but then both of them are preceded by systems. What's special about a systems approach, and this is for both of you, relative to the historical approaches in neuroscience and neurobiology.

Michael Breakspear
I think the biology matters, and maybe there's a nice tension in the field with the neurobiology versus just the pure systems neuroscience, is that the implementation does matter. You might get underlying universal principles like self-organization, complex dynamics, multi-scale coordination, but when you look at it in the brain, there are specific differences and additional levels of complexity that matter.

Like neuromodulation on different timescales. The way the actual substrate of the brain can change in response to its activity, because of long feedback circuits. These do matter, and they differentiate the brain as a system from other systems. I think there's a nice tension there that

sometimes the implementation is really important, and pulls the framework of what you're looking at back towards the specifics of the brain, and even that circuit that you're looking at.

In other instances, you can strip away some of that detail and say, "Hey, in this way the brain actually looks a lot like other self organizing systems." That's, I think being part of the creative use of the collaboration with Mac when we're having these discussions, and my tendency to pull away to a theoretical framework, and Mac's groundedness in bringing back all the beautiful neurobiology.

Paul Middlebrooks
You never think, "Oh, I wish he would leave off talking about the 1700 neuron cell types. What does it matter? I want to abstract them away."

Mac Shine
All the time.

Michael Breakspear
Sometimes if we're writing a paper together and we're doing this iterative process of editing, there'll be more detail in it after Mac's edits, and then some less detail I go through, and then we converge on some sweet spot. Mac's already thinking of brain from a systems perspective, whereas sometimes you find yourself in a forum and people are like, "What about the gap junctions?" Or, "What about this cell type?"

That's when the dialogue begins to fracture. When there's a insistence that a particular piece of detail is really important, that's where the systems' perspective disappears or is no longer that useful. What we've seen over the last few decades is that the systems' approaches finding more and more applications and showing more and more deeper insights, I think.

Paul Middlebrooks
It has been a few decades now. Mac, was I interrupting you? Were you about to add something?

Mac Shine
Oh, no. I was just going to echo what Michael said. I think it's really fun to go back historically and read some of the really impactful papers that set us on the particular course that we're on now. I just so happened to be reading the old Hopfield and Tank paper from 1986 when I was four years old the other day.

It's a beautiful paper. Really one of the starting points of much modern deep learning neural network parallel distributor processing ideas. Hopfield and Tank say in their paper, it's amazing how far you can get by just ignoring all the details of the neurons. Then they go off and have this really interesting hypothesis that led to this idea of parallel distributed processing.

It's also not clear that all those details don't matter, too. It could be that the gap junctions end up doing something computationally that we don't understand yet, that will end up being really beneficial, but we don't know. This is why I really like to push on the neurobiology to say there's all this hard, one stuff, all this broken symmetry down at the level of biology, and some of it will end up mattering for the computational outcomes that the brain's capable of, and some of it won't.

Right now, we've gotten really far with one subset of it, which is really, really cool and interesting, but let's go and see what else we can get from the rest of it. Maybe there'll be some stuff to learn there. Odds on there will be, because evolution and natural selection really wants to find good answers. It's like water finding the cracks on a pavement. It's going to find the solutions that work.

I think at the end of the day, taking that historical perspective, but as Michael said, not just saying, "Oh, look, I found a little squiggly wiggle on the brain, maybe it works." Think to yourself, what on earth could that do for the system, and how could it benefit an entire organism wandering around its world? I think the thing that's really exciting to me is that we're now starting to get these paths between the scales that weren't there for a long time.

You can now say, "Look, here's one hypothesis about what a chemical like noradrenaline could do. Let's go see if we can make that go up the scales and make some predictions and test them." It's not to say that it's the perfectly correct way. The atomic model of the solar system was really popular for a long time before we realized that we were thinking about the solar system the wrong way.

We had this earth-centric rather than a solar system-centric perspective. I think we need to keep that humility as we go through and realize it's a really difficult thing to go across scales, but as Michael said, we need to take the neurobiological detail and the computational so what, and try to make them connect as many ways as we can. It's exciting to me.

Paul Middlebrooks
You were actually alluding, with the noradrenaline, to one of the first projects, or maybe the first project that you guys had together. I want to get to that and reflect on that. What Michael was saying earlier sounded like also it's really dependent on what perspective you're taking and what questions you're asking, whether you're going to add the details from gap junctions or not, or whether you're going to consider them within the system, because it depends on the timescales you're asking about.

If you're going across scales, you have to somehow look at all of it, and there need to be systems properties. It does seem the noradrenaline thing that you were just mentioning, Mac, in that particular system, yes, it eventually goes across scales, but it begins by studying noradrenaline. Does it depend on the perspective that you're taking for the particular project you're working on, I guess, is the question.

Mac Shine

I think undoubtedly that's a lesson I learned from Michael. I still vividly remember a talk he gave at *Human Brain Mapping* about the observational lens that we take as a scientist, and that if we don't realize that we're looking through a lens, we can think that the way we're looking at the system is just the way that the system is. If we step back from that and take ourselves into the picture, then all of a sudden we're like, "Oh, actually the whole world isn't just 40 Hz oscillation, synchronizing together."

That's just one way that you have to look at the system, and that might be how it looks through that lens, but if we go to look through a different lens, through the BOLD signal, or if we go all the way down scales and we're looking at electron microscopy, it's all the same system. We're just looking at it in these different ways.

I think that's a really important step for us to take as a field, to step out of our epistemological clothes into a more ontological space where we can say, "What is the actual system underneath all of this? How am I bringing something to the table? How is it different, or how does it make a different prediction than what you would've made? How can we reconcile those?" To me, that's a real important step for us maturing as a field in neuroscience.

Paul Middlebrooks

Michael, you've been doing this for a couple decades. You mentioned it's fun to-- You both alluded to the early systems papers that are a few decades old. My historical context is that, probably because of the reductionist molecular biology revolution that the mainstream in neuroscience, yes, the systems approach has always been there, but it's maybe waxed and waned a little bit, but I grew up in a very reductionist perspective approach to neuroscience. That seemed like the mainstream. Historically, thinking back, Michael has it waxed and waned this systems approach in popularity?

Michael Breakspear

I'd say more waxing than waning. You're right, I did medicine in the year-- It sounds like I'm very old, late 80s, early 90s, and everything that we were taught in neurophysiology and physiology, and neuroscience was this reductionist approach. At that stage, the single gene leading to the single knockout neuro leading to the phenotype was still very much the dominant way of looking at neural disorders, including schizophrenia. It was also the time of the so-called grand-- What was it, the grandfather cell?

Paul Middlebrooks

Grandmother.

Michael Breakspear

Grandmother cell, sorry. A very particular person in your life. When I started doing my PhD in '99, I basically was in physics and maths at that stage, I wasn't really interacting a lot with the neuroscience community because I wanted to learn all these beautiful nonlinear dynamics. At that stage, it was apparent already that there was some pioneers applying this for decades in neuroscience, people like Walter Freeman, Paul Nunez, many others. I worked with Peter Robinson and met Vik Jirsa who was starting to do neural field theory.

This was an active field, but it was published in Biological Cybernetics, physics journals. They'd been the rise and fall of chaos theory in a very cartoonish manner. When we started then applying our mathematical tools to neuroscience, if it was time series analysis, you could get it published, but if it was a model-- I remember sending a paper into *Journal of Neuroscience*, probably 2005 or something, and one of the reviewers saying, "Look, this looks like an interesting paper. I don't find any fault in it, but this type of paper just does not belong in the *Journal of Neuroscience*."

Paul Middlebrooks

What year was this?

Michael Breakspear

I think it was about 2005. There was still a lot of gatekeeping from the type of modeling material, all the type of material that could get into a mainstream neuroscience journal.

Paul Middlebrooks

Now you can't publish without having a model.

Michael Breakspear

Yes. With that in plus computational biology. It was hard to get the mainstream neuroscience field to engage in what we were doing, and we started doing epilepsy neural field modeling, and bit by bit through presenting the work internationally and making connections. Suddenly or slowly, these types of papers started to get published. *Journal of Neuroscience*, *Human Brain Mapping*, *NeuroImage*, *Cerebral Cortex*, and then there was a watershed moment in the mid-2010s and 2015. From then on, people like Mac and others have really pioneered what the field can do and come into mainstream neuroscience. That's my recollection of it anyway.

Paul Middlebrooks
Mac, are you part of mainstream neuroscience now?

Michael Breakspear
Publishing in mainstream neuroscience, or at least pioneering work with that lens, we talk about lenses and apertures to look at the brain, but there's also a lens and an aptitude to look at neurosciences as a process. That's what I'm talking about.

Mac Shine
It was really interesting, Paul growing up in the human brain imaging community, where Michael was already a really well-established and respected person in that community. I think Michael, by your own admission, you were a little bit different than a lot of the other groups. There'd be groups out there doing whole brain imaging, and look at this map that I found, I found this network of the brain or something like that. It was anti-correlated with that other one.

Then you'd go to Michael's tutorials, and you'd be like, "Here's a bifurcation in a little simple toy model of a neuron. Now, what happens if we scale it up?" He was really coming from this completely different perspective. I remember when I first saw you talk, actually, Michael, it was my first ever *Human Brain Mapping* at Quebec City. This is really way back in the day.

It was an education course again, and it was just overwhelming. As someone who came from a background in medicine as well, we'd been told all these really beautiful stories about physiology in different organ systems, like how cells in the heart can combine together to form a pump, sort of systems-level. You'd go to the brain, and you wouldn't find those same stories. In part it's because it's so complicated.

Neurons are this incredibly difficult mathematical object on their own. Then when you put them together into groups, they do even more complicated stuff. Then even that's an approximation of the actual complexity, you see at the whole systems level as it's interacting with the world. It's just this bewilderment that you have when you go, "Wow, how is this thing that I'm scanning in an MRI scanner that looks like this, how is it related to this thing that's completely different?" It feels like it's different languages.

I think something I've always admired about you, Michael, and why I was so excited to start working together back in the day, was that you were fighting that fight for that community. This community, you could see how it could benefit from that constraint that was imposed by the reality of neural or non-neural systems as they were, but they had this beautiful access to the whole brain as well, in a way that a lot of other communities didn't have.

I think it's that connection between those that's so exciting. The new scientists that are coming to the field now, you can now come in with a mathematics background or a physics background, and immediately get shown four or five just fascinating questions that you can come in and try to solve. Now you don't have to go and stare in a telescope at some binary star system that's wiggling around each other in a weird way, often in the distance.

You can just come right in front of you and look at the neural data, and try to see if you can apply principles from your training to that problem and help solve it. I think the fact that that's been set up by this tension that's been held onto, rather than just their fields separating, they've tried to come together, and I think you really were a pioneer in that front, Michael. To me, that's what's so exciting about this next generation that we have in front of us, these new students.

Michael Breakspear
I think there's a really good analogy there, Mac. If you look at the field now with neuropixels and other wide field, whole field calcium imaging, and optical imaging, and this is a paper that Mac and I recently published as well with Brandon Mann, the complexity is overwhelming. You need whatever tools you can avail yourself. PCA is going to get you a little bit of the way, but if you really want to understand the dynamics, you need to bring in the tools that have been used by physicists to understand other complex systems. Where normalization tools, which we used in this paper, bifurcations, face transitions, lots of beautiful methods that are available, and [crosstalk]

Mac Shine
About science and graph theory.

Michael Breakspear
Yes. Without those, the data itself is so complex. I think this is a good analogy for the brain, that if the brain was really as complicated as all of its components looked at through a microscope would suggest, we wouldn't be sitting here having a conversation, and as three different brains making sense of what each other is about to say, and thinking and the narrative, and moving with fairly low-dimensional arcs through space. Our motor system has to collapse, all these bewildering degrees of freedom.

One of the really important principles from physics is the notion of the enslaving principle, and how many systems with many degrees of freedom can contract down or collapse onto low-dimensional manifolds, and be enslaved by the order parameters on those manifolds. It's like a reverse causation where now you've got the mass action of many components on a low-dimensional manifold.

The tiny little systems, the tiny little guys, the neurons can't fire independently anymore, they're trapped in this low-dimensional manifold. Their firing rate can go up and down slowly, but it's really enslaved by what's happening around them. Through this massive collapse of the degrees of

freedom, we're still maybe left with seven or eight or nine degrees of freedom. It's very hard to imagine that in your mind. Even four-dimensional space is pretty bewildering.

That's required for humans to converse and to make sense of each other, and to make sense of other brains. If that's what the brain is doing, collapsing down its degrees of freedom to make sense of other brains and other agents in the world, then that's a sensible method to use in the laboratory. There's a nice hermeneutical principle there, I think.

Paul Middlebrooks
You just said a lot, and I was furiously taking notes because I wanted to interject and ask questions about almost everything that you just said. One of the things that came to mind as you were talking, the physics inspiration has been very successful. You've used it, from dynamical systems theory essentially, and studying what you call studying complex systems.

It made me wonder, and this is a naive question, because I love the dynamical systems theory stuff too. I love making things low-dimensional. I love manifolds, all of that stuff. There's this huge influx of physicists into neuroscience, and you know the cartoon where it's "Why are you looking for the keys under the light post? It's because that's where the light is."

I wonder if there is, in a sense, and I know right now it's a burgeoning time, and there's a lot of influx in using inspiration from physics to study the complexity of brains and brain functioning. I wonder if, in a sense, there is this, "This is what we've done in physics. Let's apply it to this other system." It's a sort of analogous, like looking, that's where the light is, is in physics. I'm not trying to insult anyone, but I'm wondering if that even makes sense.

Mac Shine
I would say the analogy is a bit closer to-- The physics approaches are all flashlights, or torches, or different types of light-emitting objects that'll have their own challenging lens with them. One of them is a laser pointer, the other one's a big waffly light, and your job is to then go around looking in the space to see whether or not any of those particular lenses are helpful for you to find the keys. I think the drunk looking for the keys under the spotlight to me is a little bit more apt to think about the way that in neuroscience, we've been looking at parts of the system and assuming that the whole thing is in there.

Paul Middlebrooks
That's the obvious example.

Mac Shine
Yes.

Paul Middlebrooks
Sorry, I didn't mean to interrupt you, but [crosstalk]

Mac Shine
No, no, no. Not at all. I think it's hard not to do it right, because if we are as pedantic as we need to be, then everything is just overly defined and hyper precise. Like the classic example is, we found examples of this, this, and this, and then you have to go dot dot dot in mice. Another one that I love is this idea of calling a neuron that's involved in pain, a pain neuron.

Pain is a concept that we use to describe how a person, an agent feels. It's a very, very complex coarse-grained description of that state. That neuron might very well be involved causally in that process, but I don't think if you took that neuron out and put it in a dish that there'd be any pain anywhere near it. It's just part of a big complex system.

We're not going to hold each other accountable to being that pedantic, but we have to be careful, because otherwise we start looking for a function in a microcosm of the system, of which it's not really responsible in the way that we want it to be. Now, all of a sudden, as scientists, we've already got hard enough jobs going across all these scales and learning all these technologies. Now, we've got to be philosophers too. I think it just suggests that we need to be really humble and bring a lot of different perspectives onto problems and be open to the idea that neuroscience is still really young. We need to learn how to think about the system and ask the best questions of it.

Michael Breakspear
I think Mac answered that very nicely, but physics itself is not static. It's not a set of tools or one illuminating light. It's interesting to see, as physicists find their place in the neuroscience world, how the questions and challenges of the brain are also motivating new types of physics and new types of mathematics, particularly this interplay of stochastic high-dimensional systems and low-dimensional systems.

We know the brain certainly not evolving on a very low-dimensional chaotic-like manifold. There is a lot of noise, and that sometimes the noise is really important, and this high-dimensional noise gets in and explodes what the brain's doing, and that's really important as well. There's quite a lot of interesting stuff happening about stochastic dynamics, the interplay of high-dimensional systems and low-dimensional systems.

That's partly being motivated by the questions that the brain and the field of neuroscience is asking. I do think we're in a pretty golden period at the moment. We've got the beautiful tools of physics. We've got wonderful new ways of measuring the brain, neuroimaging, behavioral neuroscience, and of course these whole-field recordings, and then we've got AI.

What's really interesting in AI is these very high-dimensional deep neural networks are doing brain-like things. At the heart of them, we don't really know why and what they're doing. There seems to be interesting dynamics going on in deep neural networks, like phase transitions and latent manifolds, and dimension reduction and all this stuff, which is brain-like, and maybe not what the brain's doing, but it's also very interesting for physicists.

We just saw a Simon's Foundation grant last week and the physics of artificial intelligence. That's a very fertile area to be working in. I hope that we are starting to put down some of th

e principles of how the brain works that are just a trend, that are actually going to stand the force of time, like some of the things that were achieved in physics in the early 1900s.

We could mention a few of those and explore a few of those. None of them are watertight, but I think we're in a pretty good period at the moment. I'm still quite young. Thank you very much. I'm excited about the future as well. I'm excited about my own future in neuroscience.

Paul Middlebrooks
Michael alluded to this most recent collaboration that you guys were a part of 2024, the MUN et al. paper, which looks at a multi-scale organization across multiple different species. I want to come to that. Let's see, that's 2024. Your first collaboration was in 2018 is when the publication happened, which was six years before. This was the noradrenaline work basically that Mac was mentioning earlier. Mac, you had told me that was your first foray into that computational world that now you're an expert in, of course.

Mac Shine
Whatever that means, yes.

Paul Middlebrooks
That started your collaborations. How did that come about, and how have you guys weaved in and out since then?

Mac Shine
It's a great question. This takes me back to my postdoctoral years. I was at Stanford University working with Russ Poldrack who was just a fantastic mentor for a young scientific mind. He really showed me how to do modern science at a really high level. He is a big proponent of open science practices and working in a really community effort to make the science as a whole better.

I'll have to give Russ a lot of credit here because we've essentially been working on a purely neuroimaging project together. I was interested in dynamics in brains and dynamic network perspectives. People had been applying things like graph theory to neuroimaging data. Michael had some really beautiful pioneering work in our space.

Again, the Honey et al. paper was a real groundbreaker there and really opened up our eyes to the fact that the patterns that we've been looking at in the neuroimaging, we've been taking these static snapshots of something like 10 minutes of data and describing those static snapshots. If you start to go and look down at subsections of the data, it actually looks quite different, and people had just started really pulling back the curtain there a little.

We did some work there, applied graph theoretical tools to this data, found some really fun patterns in the data that fluctuate in these different extremes, and spent, God, the better part of a year, really just scratching our heads, trying to think what on earth could be causing these fluctuations in these graft signature of the brain.

Roundabout way, we came across this hypothesis that it was the ascending arousal system in the brainstem. Sending these projections up to the whole brain, dowsing them in neuro chemicals that could then change how well the two different areas could communicate in neural terms. We got to this hypothesis, we're like, "Oh, cool, this is really interesting."

As soon as you get to that point, you think, "What I've just done is described the system." I'd really like to causally interact with it. Back then opto-fMRI was really just starting. No one at the time that we knew had done stimulation of the locus coeruleus, the main noradrenaline hub in the brain. There was a group at Stanford that had done some cool stuff on the zona incerta but nothing was quite there.

The other two options we had, one was to go and do pharmacological FMRI. We ended up working with Sandra Newhouse on a really fun project where people were given atomoxetine, which is a noradrenaline reuptake inhibitor. You can look at what happens in the brain with that. The other idea was to go into computational modeling, where now you're much more abstract, so you have to make a lot of assumptions, but you have direct control over exactly what you do in the system.

I talked with Russ and said, "Look, I'd really like to get into this." There was this really cool technology, this tool out there called the Virtual Brain, which was created by Vic Usin's colleague, Petra Ritter, which was out there and available. It's a bit of an easy access for someone without the mathematical background.

You can immediately download it, start playing around with all the different types of models that you care about, but it's also just so overwhelming in terms of all the choices you have to make. I played around with a little model, Russ and I had made one, we'd swept a parameter that we thought related to noradrenaline, the slope of the sigmoid transfer function.

The first thing we learned together was like, "Man, if we touch this one parameter, everything changes, and then we touch this other one, nothing changes. Rather, all of a sudden it changes." We're like, "Whoa, whoa, whoa, pump the brakes. We need to get someone in who can help us do this at a high level." Russ, I think in his infinite wisdom, suggested that I reach out to Michael, who at that time I was terrified to speak to, because I used to really put Michael up on a pedestal.

Since we've had a few beers together, the pedestals now diminished slightly, as I'm sure mine has as well. It was such a great idea. Russ basically said, "You're about to head back to Australia. Go link up with Michael. Let's do this properly." Michael had a really talented grad student, Matt Aburn, who jumped on the task, took all the code, redid it, we built it up again.

It was really my first foray into this space. We used what's called a FitzHugh–Nagumo model, which is a simplified reduction of a reduction of a Hodgkin Huxley. It's got the elements of neurons in it, but it's simple enough that you can still play with a couple parameters. It was beautiful. It was just this thing that you could control, these super complicated equations, tiny changes that had interpretable outputs.

It was just this completely different world to me, and it really opened my eyes to this access, to this intuition about what's happening deep down at the micro scale, in this case, probably about the mesoscale, where you can now say something interesting about what's happening. It was a completely different language than the one that I'd learned doing neuroimaging, even in the context of cognitive tasks.

I really owe Michael so much because that eye-opening moment for me was like, "Oh, wow, I need to learn more about this. I need to bring people in around me as I build my group that can help me do this at a high level." It really was a huge catalyst for my early career. I don't know what it was like from your perspective, Michael.

Michael Breakspear

I think that story brings out a few really important insights. One is bringing a question that you want to ask into the modeling space. As we all know, the human brain projects set out to model the brain, but there was not really a question about what to model. When Mac and I started having a chat about this particular project, it was clear that Mac wanted to understand how changing neural gain in a population of neurons could change the integration and segregation as the system as a whole.

There was already a really well-defined question. The second aspect is to take a model and strip it down to something that's simple enough, but not too simple. Not a Hodgkin-Huxley model, but a FitzHugh–Nagumo model, but nothing simpler than that. Otherwise, you end up with something that's basically not a neuron at all anymore, so just finding the sweet spot in terms of the modeling.

That's one thing, is bringing a question and taking the model at its right level. I think the other thing is, if you just, not working with someone who's had some experience in doing this, the model can be as complicated as the original system. In fact, there's this acronym, the best model of a cat is a cat, and preferably the same cat. I mean, it's a nonsense statement because the model of the cat, which is the same cat is as complicated as the original cat.

There's this Luis Borges short story about cartographers making more and more complicated model maps of the world until they set about making a one-to-one map of the universe. The adage of all of that is that models can be very complicated and can be too complicated, but the advantage of being in the modeling space is you can use your insights and make reductions and make simplifications, and come up with something that'll work.

When you're working with someone like, in this case, Mac and Russ, you've got a question to ask, you can make pretty good progress and you can get some really nice, innovative insights as I think that paper brought to the field, that changing neural gain in individual neurons change their individual bifurcations. You didn't need to get them all to bifurcate because they're quite heterogeneous.

When you get enough of them to go from one type of behavior to another, at some point then the whole collective action of the whole ensemble goes through that phase transition. There's a really nice insight that Mac brought to that paper, but then you could say, "Well, okay--

Mac Shine

Can I say something about that, Michael, because I think it's worth clicking on this? That was such a fascinating thing to me that I hadn't heard about since Year 11 chemistry or something. You learn in Year 11 chemistry about boiling water, and it goes from a liquid phase and you put temperature on it consistently, and then all of a sudden it starts to boil, this phase transition.

Never in my wildest dreams that I think that that was what was underneath the hood. I wanted to see whether or not if I changed this neural game parameter, and then I took this neural model and I put it through a balloon model, so it looked like the bold signal, could I see the integrated segregated network change? You could with some caveats.

It's a very complicated model. There were pockets where it looked like that and other pockets where it didn't, but the really crucial thing that Michael just brought up was that the change happened abruptly. A small smooth change in a control parameter and a rapid change in an order parameter. That is indicative of a concept that you've brought up on the podcast already, Paul, of criticality, which was not something that I was familiar with as a scientist.

I'd never really come across the term before, I hadn't been exposed to it. What a gift to be given as a curious scientist, something you've never even thought of before. Could be an explanation for what you're trying to understand. To me, it was just this beautiful moment of like, "Now I've got to go learn all this other stuff." It's like the curse of knowledge a little bit. Now I know that the world is more complicated than I thought it was beforehand.

I say this to my students all the time, "Try to surprise yourself, try to put yourself in contact with a world where you don't already know what the answer even looks like." I think that's been a clarion for my group. We've been trying to really hold onto that notion of being afraid of not knowing what's going on and leaning into it rather than running the opposite direction. I don't know if I've ever thanked you enough for that, Michael, but it was such an incredible exposure for me to that whole new world in a lovely way that made sense.

Paul Middlebrooks
Let me jump in here real quick. I want to make sure that we actually-- I don't think we described the ascending arousal, like what the paper was about. What I mean is how does it relate to the brain?

Mac Shine
Oh, oh. There's actually a really beautiful paper from back in 1990, Servan-Schreiber and John Cohen was the senior author, where they were trying to understand this classic Yerkes–Dodson relationship that's really pervasive in the psychological community, which is that your performance is optimal at the intermediate stages of arousal.

If you're too sleepy, you don't do very well. If you're too stressed out, you don't do well, but right in the middle you do quite well. We all know this from taking exams. They've done this really elegant model where they used the slope of a transfer function, connecting nodes in a little toy model, to demonstrate that you could get this sensitivity to the conditions in the model at an optimal point.

There was a working regime, and if you went too far to the left or the right of that regime, then you're in big trouble. The way they did that was by mimicking this slope with this transfer function. We took that notion, we knew that there were transfer functions in the neural models. Michael and I chatted a little bit about it, and then there's ways that you can sweep that in optimals starting point that actually ends up being one of the hardest things to do in a model, is to know where to start, like, "Where should I start to look?"

If you're in the wrong regime, pushing the parameter might do nothing at all. I think Eve Marder's work is a really beautiful example of this. If the temperature is in this zone, it's doing one thing and you flip it, it's a complete different thing. That was the idea was to take this mathematical abstraction, put it onto a biological system and see whether or not it could recreate the imaging patterns we saw. The answer was, yes, it could, but it also then did all this other cool stuff that then became the start of lots of other questions.

Paul Middlebrooks
This is the criticality road that you were just mentioning for one thing?

Mac Shine
Yes. Really, it's where neuroscience meets non-linear dynamics. I think that's neurobiology and non-linear dynamics, the intersection between them. Here's a thing that a brain has been shaped over evolutionary time to be able to do with cells. It can do these set of things, and here's the language that you need to describe them.

Before that paper, the most of the work that I'd done was really descriptive of a system.

Oh, here's the things that are sub-ball super threshold, or here's the set of all the things and here's the configuration. The dynamical systems to me was such a bedrock, and it is now of all my work, because it says something a little deeper, which is, what could a neuron even do, or a population of neurons do, given the information that they have or what they want to do next, or something? You can frame problems from a neurons perspective.

I found that so, so much fun to think about the implications of that and why would you have a basal ganglia, which looks so different than a cerebellum or a thalamus, what kinds of constraints would they impose on dynamics? Off you go into this weird and wonderful world of trying to guess what the emergent property of a system will be. It really opened my eyes to that possibility.

Paul Middlebrooks
I don't think you mentioned the locus coeruleus and the ascending arousal system, which is the neuromodulatory analog of these computational systems, even when you're describing how it relates to the brain. One of the things that you guys claim in the paper, I believe, is that this, and you mentioned criticality just a moment ago, Mac and I did as well, that this ascending arousal system controls the nearness to criticality in different regimes. Do I have that correct? Then-- Oh, go ahead, Michael.

Michael Breakspear

I probably used the word tune, so don't think--

Paul Middlebrooks

Tunes. Okay.

Michael Breakspear

Yes, because the behavior is emergent, so there's not really a controller. There's tuning into different regimes. Sorry, just to jump in there about that.

Paul Middlebrooks

That's all right. I was thinking about control theory. Control theory might tune it. Anyway, my question though, was, and I'll use your language, Michael, then what tunes the tuner? What tunes the locus coeruleus there? If this parameter, is it hunger? Is there a circular causality there?

Mac Shine

Yes. I think it's a good circular causality, and the kind that Michael was talking about alluding to before this idea of constraint closure, that you can imagine that their locus coeruleus activity and the release of noradrenaline will have an impact on downstream targets. Anything that expresses a particular type of receptor, which will then change the excitability or the susceptibility of that individual neuron or population of neurons, which then have to compete in this massive space, almost like a little new species for their survival.

Then they do that by setting projections elsewhere, which then recruit more help, or they send projection that then ends up bringing online the competitor that then kills them off because it's a better fit to the challenge at hand. In some ways, these things end up being the constraints that imposed on the network change the way that the network can evolve.

Then you can play the other game and say, "What are the areas that are downstream to the locus coeruleus?" There are local populations. Michael Bruchas's group just had a beautiful paper in Nature where they study the peri-LC coeruleus, a bunch of GABAergic neurons that sit just underneath and nestle in on the locus coeruleus.

They found that the different subpopulations of the peri-LC coeruleus are actually inhibitory of the locus coeruleus targets in particular behavioral contexts, and can actually be recruited by different neuropeptides. Now you've got this really beautiful complex circuit, which is neurochemical and electrical, all working together to change the context in which different subsystems of the brain can be brought online to energize the rest of the system.

I think this is just going to get even more complex, as then you look at the cortical targets or the habenula inputs to the operculum, all this different complexity will become just overwhelming. What we need is more hypothesis-driven science there. We need to be able to say, "What if I knock out this receptor? What if I take an animal that's now the dominant one rather than the passive one, and I turn this neuropeptide on? What change? Is it via this circuit?"

As Michael was alluding to, we have the tools now where we can do this. Not saying it's easy; it's just doable. The kinds of technology we need is there, and now we need time to really refine this.

Michael Breakspear

I would pull back in the other direction as well and say these critical transitions are information-rich. That's in the paper that we're talking about, the Shine 2018 eLife paper. In some of the subsequent figures, we pulled out that at the boundary between these different regimes is when the capacity of the system to do computational work is highest. It's a very complex, information-rich, emergent dynamic that comes out of the system.

Then you just need to imagine a slower feedback circuit that says if you're information-poor, if the noise is too flat, just slowly increase the gain and you'll slowly push into this information-rich regime. Whereas, if the noise is too high, if the entropy is too high, just pull back the neuromodulation in the other area and--

Paul Middlebrooks

You need something to read the noise in that case, right?

Michael Breakspear

Well, you need something to read out the behavior that the system's exhibiting and feed that back. I'm going to come to that in a moment. At that level, I think I put my systems neuroscience hat back on and say, "Look, it doesn't really matter what the circuit is. This could be enacted in a whole way of different mechanisms, but what's important is the overall information landscape and the notion that if you read it out, you have a slow feedback circuit, it'll be self-tuning." It could be some of the mechanisms that Mac just mentioned, or it could be a mechanism that we don't yet have on the board, but the principle will be the same.

Then there's another principle is if you put a bit of plasticity into the synapses, the system will start self-organizing even without these slow feedback loops. Plasticity plus criticality in very general circumstances can lead to self-organized criticality, in which you no longer need to tune

anything because the system will self-tune. This is another principle that came out of physics, self-organized and criticality, not just criticality, in which you can tune the system to be critical, but self-organized criticality.

In pretty generic circumstances, you just need a little bit of plasticity, local plasticity, blind the global behavior of the system, and you'll start to get these self-organized critical dynamics. In fact, that was exactly the topic of the paper, 2009, I think, Rubinoff, blah, blah, blah, and Breakspear, that was gatekept out of mainstream neuroscience because it was a modeling paper that wasn't relevant to what neuroscientists were currently looking at. That's a funny reflection to have.

Mac Shine

It's also an example of what we call Breakspear's law in my group, Paul, which is that anytime we think of an idea, he's got a paper from 15 to 20 years before that already had that in.

Paul Middlebrooks

That's awful. That's awful, Michael. I--

Michael Breakspear

I know. I have become a bit of old mate who goes, "Yes, thanks, that's an interesting idea, but actually--" I try and be humble. In fact, that way, if you go back and you look at the old Breakspear paper, they actually didn't do what I believe. I'm filling in a lot of details here, and also would've been very rudimentary, if at all, compared to the full richness of what's being proposed in the modern era.

Paul Middlebrooks

[laughs]

Michael Breakspear

If I was doing important stuff back there, it was because I'd learned stuff from other people and other luminaries or luminaries of the field who many times now have been forgotten. People are rediscovering prescient insights that people were having in the '80s and '90s that I was drawing from.

Paul Middlebrooks

Michael, you don't know me, but I've had a few guests on to talk specifically about criticality. I've become very interested in it in my own research, but I'm also interested in metastability and the interactions between criticality and metastability. I'll point people to this, I think it's Nature Reviews Neuroscience, or Nature Neuroscience, the Hancock et al. 2025 paper, which I really enjoyed and talks about the differences between criticality and metastability and multistability, and [chuckles] all these terms and dynamical systems that you can really go down the rabbit hole in.

Maybe we'll come back to that, but I'm aware of the time, and I want to make sure that we fast forward six years and talk about this most recent effort that we've all, I think, discussed a little bit already in the podcast in our conversation. This is led by Mon, is 2024, and this is the multiscale organization in which touches on a lot of the principles that we've been discussing. Maybe one way to lead into that is, I wanted to ask about--

All right, Michael, you've been studying neural mass models for many years now. There are neural mean-field models, neural mass models, and other-- There's this iterative coarse-graining procedure that you guys use in the 2024 paper. Then there's that kind of set of analyses. You mentioned PCA earlier, Michael. Then there's the popular low-dimensionalality reduction techniques these days. PCA, other non-linear dimensionality reduction techniques, but the neural mass models and so on that I mentioned before are our dimensionality reduction techniques in a sense.

I'm hoping that you would just contrast neural mass model approach, the mean-field approach as a dimensionality reduction technique approach, as opposed to PCA and some of the more common and popularly commonly-used techniques. I think people, the neuroscience community, are less aware or less well-versed in the neural mass models and iterative core screening approaches, et cetera.

Michael Breakspear

Well, that's a huge area to explore, but try to simplify things. At some level, you say, "Oh, this is the level I'm going to engage this particular complex system. I'm going to take a mean, I'm going to assume the mean, and maybe the variance of this system in this local area will tell me a lot about how the system is behaving." Then, using different approaches, you can then come up with this neural mass or mean-field approximation at that particular scale.

It might be the scale of a cortical mini-column or column. It might be the scale of a couple of centimeters. Then you get some principled approach for what that small system, how it behaves. It's inputs, it's local operations. There's normally a sigmoid curve in there. There's some filtering, and then there's an output.

Then Mac mentioned the Honey paper. This is the 2007 paper. We took hundreds of these little local patches, and using best possible connectomics at that stage, or CoCo, Mac, we wired them together. Even though we had three or four dimensions in each area, we still had 80, 90 these days, 500, 1,000 of these little populations. You've still got a system with, say, 3,000 degrees of freedom, but that's when the nonlinear dynamics takes over. Then you get principles of synchronization, you get underlying low-dimensional manifolds, but that's what the dynamics are giving you.

There's no mathematical tools anymore. The system itself is finding these low-dimensional manifolds, these symmetry manifolds, and contracting onto them. We had a paper, Roberts 2016, on metastable brainwaves, which had these 500 nodes. It contracted onto something maybe eight or nine dimensions you get these metastable traveling waves, breathers, rotating waves.

Paul Middlebrooks
The breathers, are those the one that diffuse?

Michael Breakspear
Yes, diffused out. These are things that applied mathematicians have been fighting in other systems. That's where you no longer are forcing a dimension reduction. The underlying symmetries in the system are then doing the work of doing their own dimension reduction. That's the really beautiful stuff.

For a long time, these assumptions had been somewhat untested. There was a leap of faith saying, at this stage, at this level, we can throw away a lot of the statistical moments, the kurtosis, the var, the skewness, et cetera, et cetera, and we could just look at the mean and maybe the variance. When I was visiting Max Group and Brendan Munn was saying, "Oh, we can use iterative coarse-graining on these whole field recordings," this is a way of no longer having the leap of faith.

By iteratively coarse-graining, starting to see what are these beautiful whole field recordings, what are the statistics, what are the moments-- We know from- and this is what Mac wrote in the paper, the single cells are very sparse. They spike, and then they don't spike, and they don't spike together at the same time very often. They're very kurtotic and relatively uncorrelated.

Paul Middlebrooks
Define kurtotic, please. Just real quick for the audience.

Michael Breakspear
A kurtotic system is a very high, central moment. It's very peaked close to the mean, and it's got very long fat tails. Most of the activity in that system will be close to the mean, but every now and then, there'll be what you call an extremal event, which is a long way from the mean, much further than the standard deviation would normally suggest.

Paul Middlebrooks
All right, thanks.

Michael Breakspear
Through iterative course grading, and I'll let Mac tell some of the narrative as well, you begin to build up the statistics, or you let the behavior of the system to tell you the statistics. What we found in that system is that the mesoscopic scale, finally, the system became more or less calcium. The kurtosis, more or less what's called the excess kurtosis, disappeared. It wasn't at a massive scale. It was a few scales down the mesoscopic scale, which was beautiful for someone like me, because we were doing mesoscopic neural mass modeling for years.

Paul Middlebrooks
What's mesoscopic in the brain? What counts as mesoscopic?

Michael Breakspear
Tens of millimeters. Millimeters to tens of millimeters.

Paul Middlebrooks
This is spatial. I just want to say the iterative coarse-graining approach, it's very basic and simple. Just to say what it is, you take one neuron and then you find the next neuron that is most highly correlated with that first neuron, and then you pair them and you sum their activity or something. Then you look for something that's the most-- You're just iteratively taking the most correlated signals with each other and building up from there. Iterative coarse-graining sounds fancy, but it's a really simple procedure, from what I understand.

Mac Shine
Yes, that's right. It's actually fun to go back one step here to explain how we came to this. I think many really fun things in science, it actually comes from disagreement rather than us all joining together in a kumbaya circle and talking about how great we are. When Brandon joined my lab, he had a background in physics, but had worked on neuroscience applications in some really fun work looking at the Hearst exponent in thalamic recording or cells that they had in slice from marmosets. He was already neuroscientifically quite literate.

All of his intuitions about how the brain worked were built on the things that he'd seen in spiking data. He'd seen this cell's spike, it was very different than this other cell's spike or something, and it would change with the brain state. For me, all of my intuitions were based on whole-brain imaging, like taking the bold signal and all of its smooth lilting, it looks like tides sloshing around or like a lava lamp or something. I had all these stories about robustness, and he had all these stories about sparsity and efficiency.

When he first joined the lab, we just had a series of five or six, just arguments on the whiteboard, "No, you can't be right. There's no way that's right." "Well, it is." He'd say, "No, it's not that."

Paul Middlebrooks

Oh, I thought you meant someone would post something one day and then the next day there'd be an argument or statement.

Mac Shine

Oh, no, I wish we were that clever. No, we're not mathematicians, Paul. We're biologists. Anyway, we had these very respectful, but great debates. Brandon, in his own unique way, thought, "Screw that, I'm going to go find a thing that can actually test and show Mac that he's wrong, that it's not the system looks the same. It looks different depending on how you look at it." He had read really beautiful work by Bill Bialek, which introduced this idea of the renormalization group, a particular way of doing it, which is this course-granting procedure.

In their case, they actually renormalize at every level so that they can then directly relate the different features to one another across scales. What we wanted to do was not renormalize. We wanted to be able to say something about how the variance or the timescale scaled as we went up the different levels. As you say, it's very, very simple. You take all of your data, you create a correlation matrix of all the different unique elements in the system, whatever scale it is. You then rank it according to the highest correlations. You pair those two neuros together, get the sum, and then you do it again and again and again.

Every single level, you're potentially going up and iteratively doing this procedure. It's very similar to something like hierarchical agglomerative clustering, but what you're doing is you're forcing yourself to only pair once at every scale. In agglomerative clustering, what you do is you say, I'm going to pair neuron one and two, but they're the most correlated with neuron three, so that I'll pair one and two with three, and then maybe one, two, three, four will become a blob, and then you keep going until you have done the whole system.

That's great, and that's going to show you local populations, but it doesn't tell you how something like the variance scales as you go up the different levels. The other thing that we did was a little bit different. A lot of the old core screening work that was done in physics that was really helpful, say for taking the movement of a fluid and trying to come up with a mean field approximation of it does local spatial clustering. What they'll do is they'll say, I'm going to take the pixel of my video here and right next to it, and I'm going to put them together, and then I'll do it for 4 pixels, and then 8, and then 16 until I get to the whole level of course screen.

What we did instead is we use the temporal similarity as our guide. We're not doing special course screening; we're doing temporal course screening, so we're looking for Pearson correlations. One of the things you find really early on is that it's not just like blobs that grow in size in the brain. You start to see very early these distributed networks that cross anatomical boundaries. Part of the tectum is connected to the cerebellum is connected to the brain stem, it's connected to the telencephalon of these little marble zebrafish.

Paul Middlebrooks

These are with no lags between the brain areas? I'm sorry to interrupt, but it's just--

Mac Shine

Oh, no, no, it's all zero lag at this stage. We played around with a bunch of different things using different measures than the Pearson's, playing around with the lag, playing around with using triplets rather than pairs. There's a whole bunch of things, but in all the cases, we see basically the same type of idea, which is that there's really no privilege scale. I think Michael's comment about the mesoscale being where you start to see the Gaussian statistical representation is accurate, but one of the other really interesting take-home is that the variance scales as a function of the size of the system.

It really has this beautiful relationship where it's not scaling such that all of the different elements of the system are just noise, like we would maybe assume in a mean field approximation. If we did that, then what we'd expect is that the variance would scale with a particular exponent. If it could be the opposite thing, which is that there's no unique information at all. Whatever you see at the bottom is exactly the same thing that we see at the top, but it was somewhere right between those two extremes.

There's a really beautiful set of intuitions that come from that, a little bit like back in the day with the modeling paper in eLife, where you can start to say to yourself, "How would you build a system that could do that? How could you build a system that would have that feature?" A lot of our toy models, if you went to build them, let's say I put a network of a particular size in my system, well, guess what? It's not going to scale like that because it's going to scale together until the network, then the resolution of the core screen is above the level of the module you put in, and then it'll scale randomly as if you're just adding random noise in.

Now, all of a sudden, you have this tool that you can use to interrogate what a built system looks like. That's why it was so exciting to me. I'd come from this world of whole-brain imaging and networks and constructing no models of networks, like small world networks and things, and we could now start to say how many of those could actually give rise to what we saw? It was this really, really great opportunity to essentially take an argument and then let it play out with data and just see how-- Again, I didn't see this coming was surprised and really excited by where that led us in terms of a conversation, what's coming next.

Michael Breakspear

I think Mac summarized it really beautifully. I think what is interesting is that it is a really simple tool. It's just diadic, in other words, pairwise, pairing at each scale, and off you go. To this day, I still wonder, because we have Noel models to see if it's trivial. If you just had any fast scaling with the same microscopic statistics, would you get this? The answer is no. You do need this delicate balance of order and disorder and synchrony and non-synchrony to get it.

That's what the brain's doing. It's just tuned in the way that Mac described. Because the tool's so simple and all of this pops out, I think that's the power of this insight. It's just a very simple tool. It's not high-dimensional nonlinear PCA stuff going on here.

Mac Shine
Can I actually just add to that quickly, Paul?

Paul Middlebrooks
Yes.

Mac Shine
I just want to give a shout-out to just how hard Brandon, in particular, just worked on this because this was a classic example of him having to explain a complex physics concept to me, a non-physicist, and then deal with all of my annoying questions. After he got through me, he then had to deal with Breakie's annoying questions, and then had other collaborators' annoying questions. It was a herculean effort.

Probably the thing that I love the most about this is, I think, again, one of the more surprising results. We challenged Brandon to take this really simple approach and go look across species. We started out in a larval zebrafish, which is just beautiful. Some data from Ethan Scott, a collaborator, where they essentially have access to pretty much the whole larval zebrafish's brain--

Paul Middlebrooks
That's why you started off with that data set, because you could see calcium signaling across the whole brain essentially?

Mac Shine
That's exactly right. We wanted to get the recordings that we got the most of the brain in the highest resolution. Brandon then went and said, "Oh, actually, there's some really beautiful mouse data out there." I think he took a string of data set, analyzed that data, found the same scaling relationship, and we're like, "Oh, wow, okay." Again, this was now just V1 rather than the whole brain, but still amazing to get that many neurons.

Then we thought, well, how far else- where else could we go and look? He looked at ferret data. He ended up looking at Drosophila, an insect. He even went down to a C. elegans worm, and he found the same scaling relationship. This is fascinating to me.

If you look under the microscope at these different nervous systems, they couldn't look any more different. Drosophila's nervous system. All of the cell bodies sit on the outside like in a shell, and all the inside is neuropil, and their synapses look fundamentally different than a mammalian synapse. Yet, somehow, some way, the temporal organization of their system followed the same type of principle that we saw in both the whole and a part of a mammalian brain.

As someone who grew up with an evolutionary biologist father and has a deep reverence for thinking about nervous systems across phylogeny, what a beautiful result to think about why would you want to organize a brain like that? I think the answer's really simple. The world is really complicated, and it has many, many different things in it of lots of different shapes and sizes, different lengths. Some things happen quickly, some things happen slowly. Some things are small, some things are big, and you don't know what you're going to need to deal with next.

What you need to do is have a system that can handle whatever comes in next. You should be optimally susceptible to all different shapes and scales such that when the change happens, you're ready to take advantage of it and adapt to it with plasticity or whatever inbuilt features you have to solve that problem.

I know that's an oversimplification, but what a beautiful idea to explore further and to think about and to start to ask questions about in this space. It took me into a world that was way beyond where I thought we were going to get when we started from an argument between each other on a whiteboard. A really lovely paper to be involved in.

Paul Middlebrooks
Let me see if I can summarize the take-home, and then you guys please correct me. At the very microscopic scale, a few neurons, a single neuron, a few neurons, the activity between those entities is fairly uncorrelated, sparse, et cetera. As you grow to the larger scales, meso scales, then you start to get a robust and more highly correlated signals at these larger meso scales.

What this is good for is, we talked about criticality earlier. It puts you at a critical point where you can do integration well and you can do segregation well, so at the small scale promote segregation, the larger scale promote integration. You're at this point, bifurcation perhaps, critical point where you could go either way, depending on the needs of your current environment. Where did I go wrong there?

Mac Shine
I think the thing I'd add to it, I'd love to hear your thoughts on this, Michael, is that it's that, but all the way up and down. It's like that old "totals all the way down" idea. At every single level, you need to be perched at this balanced state. You can't really afford to have hyper-ignorant neurons. There are a fat tail of that kurtotic distribution that Michael mentioned before. They do have zero-lag temporal correlations with one another. It's just that the vast majority don't.

If we just took the vast majority to be the coding principle of the brain, we might be persuaded by something like Barlow's sufficient coding hypothesis and say, you know what? Neurons don't want to do the same thing as any other neuron. That's a big problem. Why waste the energy?

If we go the other direction and only measure local field potentials, all that we'll see is that everything's correlated together. We'll say, "Oh, what you really want is redundancy and robustness." If I lose a neuron, so what? I'll just use one of the other ones that I've got in my big pool. The answer is, it's actually not either one, but it's both. The system as a whole is poised such that it can have the benefit of both of those different features of efficiency and robustness all the way across the scales from down at the micro scale to the macro scale. It's a really beautiful idea.

Michael Breakspear

I agree. As a nice example of that, if I hold up a phone and it slips, the firing of a single cutaneous slip receptor can help, very quickly, me detect what the slip is. There are these examples where a spike can matter. A single spike or few spikes are really important for the way that we adapt in the world because of these complex multi-scale systems that we're embedded in.

If we did that, if it was always, if every spike mattered, there's so many billions of spikes, we wouldn't be sitting here having this sensible discussion. The way the brain appears and other nervous systems, as Mac has pointed out, the shared unifying principle is to allow information to percolate very quickly up across all these scales by having this multi-scale balancing act.

Mac Shine

It's actually led to some fun discussions with my wife recently. She's been listening to a lot of podcasts on cults. You think about a cult socially as a group that pushes itself off from the rest of the world and just listens to its own little story. There's no outside input. That would be like having a particular scale that's imposed on it, and it doesn't allow inputs from all the different areas to affect it. It doesn't allow the little word from the side to nudge it in the right direction.

I think it actually makes you start to think about how we can construct this world that we live in with the internet leading us into this gossip-heavy "he said, she said," or whatever" gender said" ideas. How do we build that better so that we are now open to actually communicating more effectively and not getting swayed by these little pockets? It's really fun to think about all this. In other words, the physics is a tool that you can use to look at whatever system you want. It doesn't just have to be a brain. It could be anything.

Paul Middlebrooks

Mac and Michael, I know we started late, so this is going to be on the shorter end for episodes. Michael, I know we've only touched on this one facet of your work, and we didn't talk about Egan's trapping and things of that nature, aligning maps, all of that good stuff.

I know time is of the essence. It's Monday morning for you guys. It's Sunday night for me here in the Northern Hemisphere. Maybe we can just end on this if you will. We talked about these modern approaches and the system's way of thinking and approaching, but we've also talked about how complex everything is. For both of you, what are you stuck on right now? What obstacles are you facing at the moment? What's in your way? Mac, do you want to go first, and we'll end on Michael, perhaps?

Mac Shine

Yes, what obstacles are in the way?

Paul Middlebrooks

I put you on the spot [laughs].

Mac Shine

My strategy as a group leader is to meet my students/research group where they are and find out what's interesting to them and let their natural motivation for a problem lead the question that we're interested in. One of the benefits of that is that I don't have to spend very much time motivating my group. They're all really passionate about what they're doing. The challenge is that they're all so clever and hardworking that they end up working in different directions.

I have to read so much literature to be around the ideas that they're interested in to help curate and shape and, hopefully, make their science more profound and robust. I find that enjoyable, but it's the biggest challenge. It's just digesting the amount of information required to do integrative neuroscience, I think it's just a career challenge. If anyone has really good ChatGPTs that are trained up, let me know.

Paul Middlebrooks

That is a career challenge. Thanks for sharing it. I know you have your hands full. Is there something research-wise or your next project-wise, like, if I only had this, if there was only a couple more Michaels in the world, I could-- Something like that. Is there something like that? It's fine if there's not, because I know you're juggling a lot.

Mac Shine

Recordings of all the different cell types in the cerebellum during complex behavior, that's the one that's really fascinating to me at the moment. The cerebellum is beautiful, neuronal machine, but for reasons that are quite practical, it's really difficult to actually measure cerebellar neurons and actually know exactly which part of the complex circuit they come from.

The cerebellum is so convoluted and sort of smooshed down into this tiny space that the elegant patterns that you see when you blow up the cerebellum like a balloon that are there are all hidden deep into the structure. If you wanted to understand the algorithmic basis of cerebellum function, you need to both record from it, which is really hard because it's so deep in the parts that are in the brain stem of the circuit are deeply related to a lot of physiologically deeply important structures. It's really hard physically to do it. Even if you do it, it's very hard to untangle it.

There was actually a really awesome paper that came out in Seoul a few months ago, from a consortium group where they're using deep learning to, essentially-- they take optogenetics, they can turn on and off all these different cell types. They can then have a neuropixels probed down, and then they can learn what they actually did and how to infer that from the complex signal they get in the neuropixels probe so that now they can go through with a new data set that hasn't got any optogenetics and say, what was the likely cell type that I was actually recording this wave form from?

Now, what you can do is you can go in with a kind of guess ahead of time, a little bit like spike sorting, but using powers of deep learning. I think they used a variation autoencoder or something like that to essentially get a guess about what cerebellar type you're measuring. To me, that's deeply exciting, because the cerebellum's kind of terra incognito for me and I just love thinking about it.

Paul Middlebrooks
I was going to say, it's too bad the cerebellum is something that is easily abstracted away when building any useful models of game activity.

Mac Shine
[laughs]

Paul Middlebrooks
No, I'm just--

Mac Shine
We'll see, we'll see, we'll see.

Paul Middlebrooks
Michael, is there anything that is keeping you from going to sleep at night? Aside from getting the surf the next day?

Michael Breakspear
I did have a surf this morning, actually. You can catch the--

Paul Middlebrooks
I was wondering--

Michael Breakspear
I was a bit licky. No, look, there's process issues, similar to Mac, and funding and governance that I could complain about. Putting them aside, actually, I'm really excited about systems and theoretical neuroscience and what we're working on in my group and through collaborations at the moment. I think the work that we've done, that we've discussed here, the argon mode stuff, the neural field stuff, it's a real golden era, that I'm really excited about.

We've got some papers, one about to be posted, with Mac, on cortico-hippocampal neural field modeling. That's where my passion suddenly currently resides in, in this relationship between a relatively low-dimensional hippocampus and the relatively high-dimensional cortex and how that interaction is underlying so many of our complex cognitive behaviors, so I'm excited. I just need time to write it all up and continue to learn from my colleagues.

Paul Middlebrooks
That is like the best thing about what I do. It sounds like you guys agree. It's just that you can just continuously learn super interesting stuff. It doesn't become boring, which is so wonderful. All right, so this is a very high-level and kind of a race of a discussion here today. Hopefully, I'll fill in some of the details also in the introduction to sort of frame it. Thank you both for being here. I appreciate the time. Have a wonderful Monday.

Michael Breakspear
Okay. Thank you very much-

Mac Shine
Thanks, Paul.

Michael Breakspear
-Paul, and thanks, Mac.

[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 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.