THE TRANSMITTER | BRAIN INSPIRED

Ann Kennedy explains the theoretical neuroscience of survival behaviors

The Scripps neuroscientist calls for a broader theoretical neuroscience approach in her area of research, which focuses on how the subcortex bridges life and cognition.

24 SEPTEMBER 2025 | by PAUL MIDDLEBROOKS

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

[music]

Ann Kennedy
Once I got to this subcortical area, I realized that there really was a lot of opportunity for a theorist to make a contribution and connect some of these older models and ways of thinking about instincts and behavior from the behavior-first perspective, and connecting them to modern systems neuroscience methods.

I think that there's a clear importance of neuromodulators and neuropeptides in brain function, almost like a context signal. Something that reshapes the dynamics of a neural network to do something different depending on the condition that this signaling molecule is reflecting.

She doesn't feel pain, but she also reports that she's never felt angry or afraid. She's just super chill. It seems like it's very much suppressed in her, maybe because this mutation is reducing the ability of her hypothalamus to rev up its activity and produce persistent firing and persistent motivational states.

[music]

Paul Middlebrooks
This is "Brain Inspired," powered by The Transmitter. Hi, I'm Paul. Welcome to "Brain Inspired." Ann Kennedy is associate professor at the Scripps Research Institute. She runs the Laboratory for Theoretical Neuroscience and Behavior. Among other things, Ann has been studying how processes that are important for life, like survival, threat response, motivation, pain, how those processes are mediated through subcortical brain areas like the hypothalamus.

She also pays attention to the time course those life processes require, which has led her to consider the expression of things like proteins that help shape neural activity throughout the brain so that we can behave appropriately in those different contexts, like threat and pain, and so on. You'll hear us talk about how this is still a pretty open field in theoretical neuroscience, unlike the historically heavy use of theory in popular brain areas throughout the cortex and the historically narrow focus on spikes or action potentials as the only game in town, basically, when it comes to neural computation.

We discussed that, and I link in the show notes to a commentary piece Ann wrote in which she argues for both top-down and bottom-up theoretical approaches. I also link to her papers about the early evolution of nervous systems, how heterogeneity, AKA diversity of neurons, is an advantage for neural computations. We discuss a Kaggle competition that she and her team developed to benchmark automated behavioral labels of behavioral organisms, the goal of which is to produce consistent behavioral labels across different labs which use different recording setups and systems, and serve as a common tool for different researchers to use and to aggregate data into bigger and better datasets.

You can find those show notes and learn how to support this podcast, to get full episodes, the full archive, and more at braininspired.co. Thank you to my Patreon supporters. Thank you, as always, to The Transmitter for your support. Here's Ann.

[transition]

Ann, I saw-- I've been aware of your work for some time now, but then I saw you give a talk a couple months ago at the Georgia Tech. What was it called? InterfaceNeuro workshop/conference. I think you started off your talk by saying how out of place you felt at a conference where they were talking about brain implants and modulating brain activity, et cetera. If I remember correctly, you said you felt a little out of place giving a talk in that environment. Why was that?

Ann Kennedy
It's not what my lab mainly focuses on. The organizers were some friends, and I guess I felt a little bit not exactly curing paralysis or human disease the way that they are. I think it's a fascinating community, but different background, different set of goals. I guess a common thread in thinking about how the brain modulates itself versus how we might modulate it as device developers.

Paul Middlebrooks
That's interesting. We'll get more into that. One of my takeaways, and I was-- I don't know why. Like I said, I was aware of your work, but maybe the one or a few of the threads of your work that most relate to my own work. I was surprised when it seemed to me that you expressed an interest in life processes, which was-- which is kind of the antithesis of the computational perspective that dominates neuroscience.

Now I think I had that wrong, that it is from a computational viewpoint. I was struggling to figure out, like what do I really want to know from you? That question has changed over the course of a couple of weeks in leading up to this. Are you interested in life processes or are you interested in looking at different levels of cognitive function at a systems view? How am I saying those things wrong? [chuckles]

Ann Kennedy
I think one of the nice things about theory is you can touch on a lot of different topics, but yes, life processes. I guess one of my guiding impulses has been to look for things that are outside the beaten path of theoretical neuroscience. Subcortical structures, areas that we haven't really modeled to the same extent that we have, say, cortical processes for motor control and sensory processing, so that led me to life processes. I guess survival behaviors, how the brain modulates our behavior in response to our needs.

Something I started into just because it felt underexplored by theorists and then discovered that there's actually a lot of people thinking about this. It's just they're not necessarily systems neuroscientists, and that it's only been fairly recently that the parts of the brain that are involved in survival behaviors, keeping us alive, have been accessible to the tools of systems neuroscience, to single-cell resolution recording, to stimulation and manipulation.

Once I got to this subcortical area, I realized that there really was a lot of opportunity for a theorist to make a contribution and connect some of these older models and ways of thinking about instincts and behavior from the behavior-first perspective, and connecting them to modern systems neuroscience methods.

Paul Middlebrooks
You ruined it because I was going to make a joke that what I've learned in my many years studying neuroscience is that the cortex is the only interesting part of the brain and [chuckling] everything else is so boring, right?

Ann Kennedy
It's decorative. It's just the protective outer shell that keeps the middle part. No.

Paul Middlebrooks
You did say that you're interested in these parts of the brain that are dedicated to more survival and motivation in these life processes, but the cortex is also. It's just-

Ann Kennedy
Oh, for sure.

Paul Middlebrooks
-not that that's how it's been traditionally studied, I suppose.

Ann Kennedy
That's true. Cortex is evolved alongside these deep subcortical structures, and it interacts with them and provides information to them, and enriches the things that they can do. I'd say you can do a lot without a cortex. My favorite study that I like to mention is work done by a friend who was a grad student in Markus Meister's lab when I was a postdoc at Caltech.

The Meister lab got hold of this mouse line that ex-- the driver line expressed in all of the neural progenitors that went on to form the cells of cortex and the cells of hippocampus. You could express a toxin in these progenitor cells and just completely remove the cortex and hippocampus from the developing brain. The mice came out just fine. They got a little smaller-looking head. They're a little more anxious and more aggressive, but they can do a lot of mouse stuff, and they've got nothing of those more recent structures.

Paul Middlebrooks
Like what kind of mouse stuff-

Ann Kennedy
Whereas if we could--

Paul Middlebrooks
-can they-- go ahead. Sorry.

Ann Kennedy
Yes, they can--

Paul Middlebrooks
Grooming?

Ann Kennedy
They can groom, they can mate, they can fight, they can escape from threats. They don't learn especially well, but a lot of the core things that you have to get right to survive and reproduce they get right because that's the stuff that's most ancient to our behavior, and that's the stuff that is in these subcortical structures. Cortex informs that. It gives you plasticity to those behaviors and context dependence, but you need those things to survive.

Paul Middlebrooks
Since you mentioned that they don't learn. I was going to say these subcortical structures are often what we associate with these more innate behaviors. When you say innate, that tends-- you tend to think of you come out of the womb or the egg pre-programmed with these things. However, there are people who say that, well, there are people who study learning. There are people who claim, based on their own research, that these innate capabilities actually are learned, and there is learning in these subcortical structures.

Ann Kennedy
Oh, there absolutely is, yes. I like the language of canalisation, that you've built a brain that is predisposed to producing certain outputs. It still has to refine its activity via not necessarily supervised learning, like I'm teaching you how to eat or how to fight, but there's some experience-dependent and activity-dependent wiring up of these structures, but they tend to all converge on the same wiring. Not at the cellular level, but the same repertoire of behaviors, the same drives, the same sensory motor transforms across animals.

Paul Middlebrooks
Those are deep canals, deep attractors-

Ann Kennedy
Yes.

Paul Middlebrooks
-in the dynamical systems viewpoint.

Ann Kennedy
Yes. This comes out of Waddington's landscapes. It's funny. He wrote this book, The Strategy of the Genes, in 1957. He's a biologist, he's not a mathematician, so he calls these basins of attraction creodes, not attractors. He describes it entirely through this intuitive way of thinking that genes can push cells along particular developmental trajectories, and if you perturb them away from those trajectories, they'll recover to them.

It's an attractor. He just doesn't-- didn't know that language in that area of math. I think that idea is very much present both in development of cells and also in wiring of the brain, that you have attractor states of circuits that predispose animals to forming particular patterns of behavior.

Paul Middlebrooks
You come from cortical work largely. I know that Larry Abbott does a lot of different work with a lot of different brain structures. Coming out of that vein of research, is that--

Ann Kennedy
My graduate work was in electric fish, so it's kind of impossible to get scooped in electric fish work.

Paul Middlebrooks
See, that's what you were going for with the-

Ann Kennedy
Yes.

Paul Middlebrooks
-subcortical structures, right? The anti-scoop career. [chuckles]

Ann Kennedy
Yes.

Paul Middlebrooks
Right?

Ann Kennedy
I had worked in electric fish. You can know the whole literature. You can know all the labs. I had done a little bit of work on Drosophila mushroom

body, and I was like, no, this is too crowded. There's too many people working on the same problems. I saw David Anderson give a talk at COSIGN, this is 2011, 2012, on Dayu Lin's work, optogenetically evoking aggression behavior in-- by stimulating neurons in VMHvl. It just felt so wildly different from anything anybody else was talking about at COSIGN.

This is like, you're hitting a part of the brain with an optogenetic hammer. This is very artificial, but you're getting a complex visually guided behavior. It just felt like everything else, but it was missing this very clear thing the brains are supposed to be doing, producing these complex behaviors. At the time, I was like, "Oh, nobody must be working on that. It's going to be a fun, open area of exploration." Turns out people are working on that. A good number of people.

Paul Middlebrooks
Yes, probably for a long time, too.

Ann Kennedy
Yes.

Paul Middlebrooks
One thing I was going to ask you, because you've written about evolution and studying early nervous systems and pre-nervous systems. Did that interest come from the early graduate work that you were doing? How did you come to that?

Ann Kennedy
I guess before grad school, I was briefly a technician in a stem cell lab. I did a lot of co-transfection of transcription factors into a cell line, and looking at how one transcription factor affected expression of genes involved in oligodendrocyte progenitor cell differentiation.

Paul Middlebrooks
Of course.

Ann Kennedy
Totally random. I hadn't realized until then, this is going to sound dumb, that different cells in the body express different genes, and that development is a thing. I really fell in love before I ever worked on neuroscience with developmental biology, and from developmental biology to evo-devo. This study of development from an evolutionary perspective.

I never really worked in the space, but just the set of problems that you think about in those fields of how you form an organism from a single cell, how you evolve a new behavior or evolve a new function. The ways that evolution has taken parts of a body plan and fine-tuned them and adapted them over time. That was just something, going back to high school, that I was really in love with. It was in the background for a while through grad school, but it's something that's really stayed with me.

Paul Middlebrooks
Then maybe I had it backwards. Maybe the thing to ask you is how you went from that to then studying the cognition in the dynamical systems regime in naturalistically behaving and otherwise organisms.

Ann Kennedy
When I was applying for grad school, I also thought neuroscience was cool. I think a lot of people just have this general interest in how their brains work. When I was applying for grad schools, I interviewed at Columbia, just because one of my professors mentioned it on a list of schools I should look into. Larry Abbott was meeting with the grad students, and he showed them what eventually was FORCE, this work with David Sussillo, training recurrent neural networks to produce complex behaviors. Complex time series.

The example he had was this-- they had motion capture data from a person who was walking, and they trained the model to reproduce that walking behavior. That was just mind-blowing and felt like I wanted to know how that worked and how you could have these distributed networks of neurons do these complex things.

Paul Middlebrooks
I promise we're going to get to your specific work. A recurrent neural network is all about dynamical systems because there's recurrence, and there's dynamics, and that's how they compute. Whether it's like a liquid state machine, or you're using the FORCE algorithm to train an RNN. When you saw that, for example, Larry Abbott giving that demonstration, did you immediately connect that dynamical systems view to the cellular and subcellular levels and to things like subcortical processes because you're finding in these subcortical areas that they have these same kinds of dynamical systems, shapes, and attractors, et cetera?

Ann Kennedy
Oh, yes. In grad school, not at all. I didn't know anything about subcortical structures until I started my post-doc. I entered the lab-

Paul Middlebrooks
That's right. You were dumb. I forgot. You were dumb, right? [chuckles]

Ann Kennedy
I was dumb. It took me a while to even get VMHvl down right and say the right acronym after I joined the Anderson Lab. I had learned about dynamical systems in undergrad and saw this as a cool application, but I think the narrative that was in theory at the time was, individual neurons have really short membrane time constants. They can't do things that require working memory and complex motor control, and things that require you to really integrate information over long periods of time. The rest of our computing was the way past that.

That was really compelling and exciting to me when I started grad school. For the first year or two, I was working a lot on just making-- trying to learn more about what a recurrent neural network can do and how it does it. It wasn't very connected to data, so I was always struggling with the, "Is what I'm doing really telling us anything about the brain, or am I just studying this dynamical system that is nice and easy to study?"

Paul Middlebrooks
What you're trying to do is like a little tic-tac-toe problem, or like some sort of binary on-off thing to train it to-

Ann Kennedy
Yes, it's a very-

Paul Middlebrooks
-do some do some kind of stress sort of work.

Ann Kennedy
-reduced thing. I was barely even training. I was looking at memory capacities of these things. Like how long into the past can you reconstruct the input to these networks? I think what I eventually settled on some years later was that, yes, you can get long time scales out of recurrence of a neural-- of a firing rate neural network, but I think biology probably solves the problem all the time.

Most of the time, in other ways, it has a variety of time scales that it has access to through the time constants of molecular signaling processes, through circulating hormones in the blood, through gene expression time scales, that maybe, by focusing on reservoir computing-- it was really cool, but maybe, by focusing on understanding this cool thing, we were missing how biology was actually solving the problems.

Paul Middlebrooks
Was it cool, and yet somewhat narrow and brittle because you had to tweak it just right to get the right-- to get it to function right? My understanding of reservoir computing is that it actually has a pretty high capacity for expression, and that's the whole point. Reservoir computing, you randomly assign weights to all-- between all the parameters, and if you have enough units, you can then train it with a linear readout to do lots of what-

Ann Kennedy
Lots of things.

Paul Middlebrooks
-some people think are interes-- lots of things, yes. Whether they're interesting or not. Lots of things. People spend a lot of time training these things and trying to figure out what parameters worked. Am I wrong? Are they not highly expressive, or are they highly expressive but if you set them up just the right way, and then if you tweak them, then they aren't?

Ann Kennedy
I think we've gotten better at training them stably. I have a post-doc who does this a lot. He's fitting some RNN models to some data right now. With the right constraints and tricks, I think that you can get past this fiddly feeling of trying to train neural networks to do things. I think that methods have just improved a lot over the past 10, 15 years for this kind of model fitting.

It bothered me that I wasn't studying brains, I was studying this toy model system. That's really just something that everybody has to figure out when they're in grad school, right, is, do you care about methods? Do you care about understanding math and dynamics? Do you care about machine learning? For me, it was really that being able to point back to the biology and point to a system that I was helping to understand that made me feel like I was contributing to the field.

Paul Middlebrooks
I was going to ask you this later, but since you brought up machine learning. It sounds like you're interested in the actual biological brain and processes. Where are you on the neuro-AI hype train these days?

Ann Kennedy
Oh, goodness. [chuckling] I always have to stop and try to think about how that's defined.

Paul Middlebrooks
It's not really well, but nothing is.

Ann Kennedy
What is it? Like, improving AI by looking for additional principles of computation in the brain?

Paul Middlebrooks
That and the reverse, right? It's the beautiful-- the virtuous circle of AI and neuroscience.

Ann Kennedy
I think most people who say they work on neuro-AI, it's not like I'm using AI tools to understand my neural data. It's also studying the brain will help us build better networks.

Paul Middlebrooks
Sure.

Ann Kennedy
I think the family of networks and methods that people in AI use is pretty different from the way the brain solves problems. I think you can make a compelling argument for, maybe, it's good to try to understand how brains solve problems, not just so that we can help fix people when brains break down, when things aren't working properly, but also to adapt those principles for our own computational tools.

Nervous systems have this capacity to do things locally instead of having centralized descending hierarchy of control that I think is interesting and compelling, and that we don't really fully know how to leverage in machine learning and AI systems. I think that there's, obviously, an energy-efficient argument-- efficiency argument to be made for understanding how brains compute things and how that's different from the way that we compute things in a massive language model running on our GPUs.

I think there's a good argument to be made for studying the brain as a way to push forward our engineering of artificial systems. At least I've tried to make that argument. For me, understanding the brain is really the goal, but I do think that there's-- there've got to be tricks that the brain is using that these networks aren't, and whether that's just a substrate thing the brain uses. Neurons communicate via a big array of different signaling molecules that have intrinsic time constants. These might be things that we can't really implement in a computer, but I think they're interesting to understand.

Paul Middlebrooks
What is it? Are you going to put money on that?

Ann Kennedy
Ha. No, [chuckling] I don't--

Paul Middlebrooks
I'll ask you another time. Maybe when you've done a little more research. On the other hand, it's not like dynamical systems came out of nowhere, but you mentioned David Sussillo who was working with Larry Abbott doing the reservoir computing and the FORCE algorithm that you mentioned earlier. His modeling work with MONTE is one of the early works that people point to as a proof of principle that if you study machine learning or artificial neural networks, it actually gives rise to principles that then you can transfer over to help understand brains, the dynamics in this-

Ann Kennedy
For sure.

Paul Middlebrooks
-recurrent neural network trained to do a simple context-dependent decision-making task. You studied the dynamics of the artificial network, and you actually learned something, potentially, about how brains do it. Even though brains are doing all sorts of different timescale signaling and using all sorts of different molecules and scales, et cetera. You must appreciate that aspect of it, I suppose.

Ann Kennedy
Yes. Training a network on a task and then looking at how it's learned to solve that task, and trying to use that as a way to interpret neural activity.

Paul Middlebrooks
Sure.

Ann Kennedy
Yes. No, I think that is a big aspect of how we interpret neural activ-- like when we have recordings of neurons from some complex behavior, we need some in to how we're going to interpret those recordings. I think the dynamical systems framework is a really powerful one, and that it has helped us to deal with the fact that a lot of our handcrafted ways of interpreting neurons, things like tuning curves, just don't capture much of the variance of neural firing rates. I definitely agree that a dynamical systems framework is a good way to make sense of neural data.

Paul Middlebrooks
You-- go ahead. I'm sorry.

Ann Kennedy
I was going to say, I don't know how far you can push it because you're not necessarily going to find in your artificial network a solution to a problem that is the same solution the brain uses. It's a thing you can try and sometimes it works, but sometimes it doesn't. Maybe that's because you didn't have the right representation of the task. Maybe the way that information is reaching a network, how it's processed by the time you get to the computation you're trying to understand.

In the case of subcortical structures, we're very much abstracted away from the visual field of the animal, we think, and the rich sensory environment. We're dealing more with slow motivational states in things that are hard to predict. If I took a neural network, and I trained it to forage when there's a predator around at night, attack it, I feel like it's going to take some effort for a virtual agent trained on these kinds of survival-related tasks to really have beha-- have dynamics to a state that you can map onto brain structures.

Paul Middlebrooks
You just said that those subcortical structures have slower dynamics, but under threat, you have to have fast dynamics, right? If your flight kicks in, it's fast, and that's subcortical as well, right?

Ann Kennedy
Right. For threat specifically, you need to be able to respond quickly to it, but you also need persistence. If you see a predator, and you run, and now you don't see it anymore, you need to remember that you just saw a predator and keep hiding for a while. There's not necessarily slow, but there's persistence to the dynamics of these internal states that are important in a lot of cases.

Paul Middlebrooks
All right, so let's rewind. Not rewind, but let's go back before brains existed, which you've written about recently. Then maybe we'll work up to the early evolutionary versions of brains, which is the subcortical processes, right, which then became more elaborated, and then we grew this cortex, which you referred to as-- what? As the extra stuff? Outer shell?

Ann Kennedy
The refinement,-

Paul Middlebrooks
The refinement?

Ann Kennedy
-I'll say. Yes.

Paul Middlebrooks
It's like a turtle shell. It's there for protection, perhaps.

Ann Kennedy
Yes, exactly. Probably unnecessary.

Paul Middlebrooks
[chuckles] Yes. This is a quote, actually, a very short quote, but it's nice, from the paper that you wrote about the evolution of nervous systems. Let me see. I have the title right here of the paper. Dynamics of neural activity in early nervous system evolution. "Much of what we think about as neural arose before the nervous system." This is what I was getting at in terms of your thinking and where it came from, as viewing everything at this systems dynamical level because a lot of people who get into theoretical neuroscience, there's a population doctrine now, right?

We're beyond single neuron tuning curves, and so we're going more and more abstract. People resist studying ion channels and neurochemical signaling, but you're going back down to that level but viewing it from that systems perspective. I think that that's what I was trying to get at earlier. This is along that same line.

Ann Kennedy
Yes. There's a growing few of us that think that-- that are interested in these molecular and subcortical aspects to neurocomputation. A big part of the inspiration for that line in that paper was Romain Brette has this great paper on the behavioral repertoire of paramecium, which he calls a swimming neuron. It has sensory systems, it has motor effectors. It can produce complex behaviors, exploratory behaviors, can bump into things and adjust its course.

It does all of these things, and it's a single cell. Within a single cell, everything is chemical signaling between your sensors and your effectors, but the parts that are there within the cell are the same as-- the same parts that are present in a larger organism. The motor effectors have changed. We don't have flagella, we've moved to having contractile tissue, but the sensory molecules are touch sensors and photosensors. Those have been around for a long time.

Paul Middlebrooks
The same molecules or the same functions with different molecules?

Ann Kennedy
To a large extent, all--

Paul Middlebrooks
Not that it really matters, but I'm curious.

Ann Kennedy
A lot of the same channels and molecules predate nervous systems. A lot of the ion channels are there. Mechanosensors and photosensors are pretty similar between-- what is it? Doesn't channelrhodopsin come from cyanobacteria?

Paul Middlebrooks
That sounds right.

Ann Kennedy
Man, I should have googled some of this before coming to this podcast.

Paul Middlebrooks
Say it, and then [chuckles] if it's wrong, I'll-- they'll take it out.

Ann Kennedy
Oh, man. Okay. Channelrhodopsin clearly came from cyanobacteria.

Paul Middlebrooks
In a way.

Ann Kennedy
I think Piezo receptors for mechanosensation are found in-- they're plants, they're found in single-cell organisms. The pieces are all there because they served a function before they were used in neurons.

Paul Middlebrooks
Yes, it's not like evolution was like, hey, oh, yes, we need a-

Ann Kennedy
I got neurons, now I need sensors, yes.

Paul Middlebrooks
-nervous systems. Now I need to invent some new stuff. [chuckles]

Ann Kennedy
Yes. There were these parts that were found in cells, single-celled organisms, that became co-opted to do some-- I guess, the same thing in a different structure once you got to the point of having a nervous system. There's a couple folks who really think about this. What is the evolutionary process of going from a single cell to a multicellular organism? How do you go from every cell being identical and being in a colony to having more specialization, and from specialization, how do you develop organs and nervous systems and contractile tissue?

It's really a fascinating area to explore and something where we have a pretty decent understanding of what that process could have looked like. There's a couple of different theories about nervous system origin, whether it was controlling ciliated motion or contractile tissue, or whether it was primarily secretory and modulating activity of other cells.

Paul Middlebrooks
Those are two ways, but you expound upon a third way, which is the spontaneous endogenous activity geared towards maintaining an internal state. Am I saying that right?

Ann Kennedy
Yes. That was something that we pushed in that paper because we were looking at this through the lens of jellyfish, which have a lot of endogenous activity. Endogenous periodic activity is pretty easy to generate from a lot of dynamical systems. All you need is some coupled-- some feedback coupling with a separation of timescales of your interacting components. Like one that's fast and one that's slow.

Paul Middlebrooks
Does one need to be excitatory, one inhibitory as well? Like for the--

Ann Kennedy
I think you can do it with a purely inhibitory system. I think that's how it works in the pyloric rhythm.

Paul Middlebrooks
I guess. Well, probably. I guess it depends on the timescales of the signaling.

Ann Kennedy
Yes. Having separate timescales of your interacting components, that's most important. It's very easy to get and, possibly, a useful thing to have. Periodic activity shows up in a lot of parts of the body. It's involved in swallowing and digestion, swimming if you're a jellyfish, circadian rhythms, possibly periodic signals in release of hormones, and signals like vasoconstriction. There's interesting bursting of insulin-releasing cells for a variety of systems that really have these oscillators in them.

One of the arguments for that is that if you need to convey a signal from one area to another, you can do it in the amplitude modulation way. You can send a fixed level of a subs-- of a ligand to your downstream receptor, or you can do it in a frequency-modulating way. If you do amplitude modulation and you have any noise in your signal, it's impossible to distinguish that noise from the amplitude you're trying to convey, but if you're communicating with pulses, having little white noise on top of that isn't going to interfere with your sensing of the frequency of those pulses. Oscillators were, possibly, a feature of early nervous systems and are useful for a variety of biological functions.

Paul Middlebrooks
There's all this talk about what are brains for. They're for movement, they're for perception, they're for survival.

Ann Kennedy
Control of development?

Paul Middlebrooks
Control of development.

Ann Kennedy
Something that controls the timing of hormone release, something that drives you through--

Paul Middlebrooks
Well, you're talking about the oscillations themselves, right?

Ann Kennedy
No, just like a metamorphosis. When you decide to pupate and emerge into your adult form. These are things that are controlled by the nervous system. Puberty is controlled by the nervous system. It does a lot of things beyond behavior.

Paul Middlebrooks
Right. Well, right. Okay, that's the anti brains are for behavior. All of these are just, so, stories, but there's an argument to be made that brains are for controlling like the development and internal milieu and signaling. The thing that you point to is that it doesn't take much for oscillations to occur, and that oscillations are a great way to signal these things. Just from that, oscillations are important, endogenous spontaneous activity, important before nervous systems.

They have, these days, people like Earl Miller saying cognition is rhythm. It's all oscillations, like on, up, right? There is this recent paper looking at networks, cortical, of course, cortical networks- I don't know if subcortical stuff was in there too -using fMRI, finding that there are these transitions from different subnetworks, and that it's rhythm, and it's rhythmic. That it flows through these transition states, common transition states, at a-- on a cycle, in oscillation. Is it oscillations all the way up? Is cognition rhythm? Is that your bet or is that just-

Ann Kennedy
Yes.

Paul Middlebrooks
-one perspective?

Ann Kennedy
I've never worked in that space of looking at oscillations in human data or in vertebrate data. I often struggle with how do you ground the oscillations that you're pulling out of your EEG signal and relate that to what neurons are doing. I do think that it's worth thinking about the frequency component of neural signals. For example, I was talking to Lin Qian about her imaging of serotonergic signaling recently.

She develops these optic sensors for extracellular, neuromodulatory, and peptidergic signals, and in this case, serotonin. She was looking at different frequency components of the serotonin signal. You really do see that there's these high-frequency fluctuations and concentration of

serotonin, and then there's intermediate and low. If you look at the families of receptors for serotonin, there are some that have a certain low-pass filtering of their response. They need to be able to bind to serotonin for a certain amount of time to really kick off a downstream signaling cascade.

It's something my post-doc and I are talking about is like, is separation of signals in frequency space a thing that we should be thinking about in terms of how it can influence communication between barren areas? Also, maybe, development of neural networks. If you have plasticity rules that function on certain time scales, they let certain frequency ranges of activity shape synaptic weights, and they filter out other ranges of activity.

I don't know if different brain areas use certain frequency bands to communicate with each other. It's just not an area where I've ever worked. This temporal component of neural activity is a thing to keep in mind.

Paul Middlebrooks
Not to bring us back to AI, but this is a glaring omission in any AI system, is the temporal aspect of there's no timing in an AI system. Right?

Ann Kennedy
Right. Unless you're doing robotics and closed-loop control is not something that most of these systems that we think about have to deal with.

Paul Middlebrooks
Yes, right, but in biological systems, it is timing all the way up and all the way down. It's crucial. We're just listening. The development. You have to time the fight or flight. It's all timing. I'm on a kick. I'm betraying this a little bit. I keep coming back like, God, what are the differences between artificial systems and natural biological systems? What are the differences that make a difference?

Dynamics, writ large, is one. To say that "you can point to dynamics" in like recurrent neural networks that are artificial machine learning systems, but what I mean by dynamics more so is timing of things. Everything is very sensitive in biological processes. Whether it's oscillations or not, everything is very sensitive to timing. That's what I was getting at.

Ann Kennedy
Yes, the timing is something you can't really separate from biological computation the way you can from computation in a deep neural network. Yes, absolutely. There was this feel for a while of asynchronous computing, computing without a clock signal, where all of your components have to do things on their own timeframes without any kind of central controller that's keeping track of the progress of computation, and how you could make algorithms that work in such an asynchronous system.

I think it's-- that's a problem that the brain has to deal with, this fact that everything is done locally. You don't have global teaching signals for learning rules, necessarily. You don't have global clocks that are deciding when to query information from one area and when to pass information to another area. Time is part of the embodiment of biological neural networks. This is changing track a little bit, but-

Paul Middlebrooks
Sure.

Ann Kennedy
-the other thing that's very clearly essential in biological neur-- to the function of biological neural networks and something I'm very hung up on is just the diversity of signaling molecules and signaling pathways between neurons.

Paul Middlebrooks
Oh, okay. Timing was one thing. The other thing that you focus on in that paper about nervous system evolution is the endogenous nature of the spontaneous activity, oscillations or timing or not, that it is like an endogenous-- endogenously produced. I just wanted to highlight that. I don't know if you want to comment on that as well.

Ann Kennedy
That was largely in response to a lot of the literature on nervous system evolution. When it thought about behavior, it was very much stimulus-response as the framework as opposed to ongoing dynamics of a neural system and more of a continuous control problem. We're trying to just speak to that part of the literature and say it's not just something bumps into you and you respond to it, but there's constant action of the nervous system on the body that's taking place in organisms, and we should take that into account when we're thinking about what the early nervous systems looked like.

Paul Middlebrooks
You brought up the heterogeneity of the parts. There's the famous, in my little complexity discussion group, but we're going to get to this More is different paper. I had to be careful when I say that because your calling card is different is more. You don't seem that pleased with that little turn of phrase, but it's good advertising, right?

Ann Kennedy
Yes. It just popped into my head when I was writing. Yes, I think this goes back to this dissatisfaction I had when I was first learning about reservoir computing. That we were relying on firing rates to do everything, and that this didn't feel like what biological systems were doing. I remember

when I was a grad student early on asking some of the experimentalists at Columbia what neuromodulators were, what they did, and how many of them there were.

Paul Middlebrooks
What'd they say? Elevate. They just raised their demand?

Ann Kennedy
It was just this big unknown when I was in grad school. I knew there were some of them out there. I didn't know what they did.

Paul Middlebrooks
They seem to-- have seemed inconvenient in my-- just from my growing up in neuroscience. It's like, well, we'll get to them one day, but the important thing is the neurons and how often they're firing.

Ann Kennedy
Right, and the spikes. I think it's a bit of a drunk and the lamppost thing, that the spikes is what you can measure, and so let's describe everything in terms of spikes and call it a day. There's been this push for foundation models of the brain lately where you just record all the spikes and all the conditions and then you're done. The spike is telling you that someone is talking. It's not telling you what they're saying.

It's something that I started to become aware of when I was a post-doc and trying to get the hang of hypothalamus as pretty different from how I'd grown up thinking about neurons. The more I read about it, the more fascinating opportunity for theory I think there is there. That there's so many ways for neurons to communicate with each other. You really have not one weight matrix but many based on these different communication channels.

I think that there's a clear importance of neuromodulators and neuropeptides in brain function that we've found interesting effects of these in providing a-- almost like a context signal. Something that reshapes the dynamics of a neural network to do something different depending on the condition that this signaling molecule is reflecting. I'm thinking here, for example, I've done some work recently on neuropeptide Y, which is released by neurons in the arcuate nucleus when animals are food-restricted.

These arcuate nucleus neurons send projections to different parts of the brain and start to release NPY when food-restricted. There's a projection to the medial preoptic area that reduces fertility and delays onset of puberty. This is sort of saying, if you're starving, it's not a good idea to get pregnant, and so maybe wait for a little bit.

Paul Middlebrooks
This is in mice that we're talking?

Ann Kennedy
This is in mice. When mice are pregnant, they-- I forget if it's double or triple their calorie intake. Then, when they're lactating, they lose 30% of the calcium in their bones just to produce milk.

Paul Middlebrooks
This is a neuropeptide Y, related?

Ann Kennedy
This is pregnancy and a child—a pup really, that your body undergoes these huge physical demands. If you're starving, you're not going to do so hot if you get pregnant. NPY is trying to reduce the likelihood of that. Then my collaborator, Nick Betley, found that there's a separate set of NPY-releasing neurons that project to their brachial nucleus and shut out a lot of chronic pain signals. Inflammatory pain, long-lasting nerve pain.

Animals that are food-restricted don't respond to these pain signals the way they do if they're fed. You can give a mouse a formalin injection in its hind paw and get this big inflammatory response in the paw, but if the mouse is starving, it just completely ignores it and behaves as if it's not experiencing the pain.

Paul Middlebrooks
The story there would be that it needs to find food and other things are important.

Ann Kennedy
Yes. If you're in pain, you go back to your nest and lick your wounds and recuperate, but if you're starving and you do that, you're going to die, so you need to turn down this thing and go out and do this other thing that's more important for survival. NPY seems to play this role in just preferentially biasing the animal away from some actions and towards other actions.

There's a boatload of these things. My collaborator, Moriel Zelikowsky, showed that chronic social isolation drives this brain and body-wide upregulation of the neuropeptide tachykinin 2. Mice have signs of Tac2 overexpression in inhibitory interneurons and cortex. Cortex does

something. In subcortical regions in the spleen and the gonads. These isolated mice are also a lot more anxious and a lot more aggressive towards other mice.

If she takes a group-housed mouse and she overexpresses Tac2 in different areas, she can get just the anxiety phenotype or just the aggression phenotype. If she blocks Tac2 overexpression in an isolated mouse, it behaves as if it's group-housed. This is sort of this, I don't know if you want to say inductive bias or contingency plan that's built into the genome that's saying, if you encounter certain circumstances, turn on production of this thing. Turning on production of that thing somehow mana-- magically manages to change the dynamics of fast timescale computation to alter behavior.

That is really fascinating to me, is how you can make a brain that effectively modulates itself so that it doesn't just do one thing, it can do many different things without, at the same time, breaking. You don't suddenly forget how to recognize objects when you're socially isolated. A lot of stuff is fine, but then some things are very different.

Paul Middlebrooks
If you get your leg bit off by a shark when you're surfing, you can still see the surfboard to try to maintain your buoyancy. You're not too worried about your leg being gone. You're focused on getting the hell out of there.

Ann Kennedy
You're, maybe, not feeling the pain that you-

Paul Middlebrooks
Not feeling the pain,-

Ann Kennedy
-do when you get to-

Paul Middlebrooks
-but you're very worried about it.

Ann Kennedy
-the beach. You're in panic mode. You've got adrenaline pumping, and you're acting, but you're not prioritizing, like, obsessing over the pain.

Paul Middlebrooks
What I want to ask you about is causality. It sounds like you're naming these peptides, and when they're generated, they have these massive effects on our cognition while still leaving some parts of our cognition intact. Pain is absent when NPY is expressed, at a certain amount. I'm not sure if I'm mapping the names onto the function.

Ann Kennedy
Yeah.

That sounds like it's-- the system is fragile, even though you just said it's robust. How do you think of it causally? Do we want to say like that is a mechanistic account of function? You used the word that the peptides had a role in. Do we need to bring it back from strong billiard ball causation to more of a context-dependent, softer causality story? How do you think about causality in that way?

Ann Kennedy
Yes, causality is tricky. Maybe I've been a bit billiard ball in the way of framing it and thinking about it. We do have results from perturbation experiments. For example, with the NPY story, if you just-- if you take the neurons in parabrachial nucleus, they express the receptor for NPY, and you just inhibit those neurons artificially, you get the same effect as if the animal is food-restricted.

You can go in and apply a perturbation where, if you block NPY, you block the hunger-induced suppression of pain responses. We can manipulate these systems, and we can ask, do they really produce the effect that we think they're doing? That's a lot of the hard part of experimental neuroscience, is really convincing yourself that this correlation between NPY release and a change in pain-coping behavior is a causal link.

Paul Middlebrooks
This makes me also think our-- neuroscience has turn toward naturalistic behaviors. What you're describing are naturalistic-type behaviors, and yet, in a truly ecologically fully naturalistic setting, the mouse might be pregnant and hungry and running from a hawk. There's lots of different contexts going on. Whereas in the lab, you're like, well, we're going to-- it's in a box.

Ann Kennedy
Do one thing at a time.

Paul Middlebrooks
We're going to do one thing at a time, so it is reductive in that respect, I suppose.

Ann Kennedy
Yes. That was something that I think I mentioned in my commentary recently, is that in one sense, you're studying a behavior that's high-dimensional. The mouse can do anything. It's behaving spontaneously. In another sense, in one of these essays, if you're studying aggression, you have a mouse in its home cage , you plunk another mouse in there, you watch them for 10 minutes, and then you stop. In a sense, that's a very low-dimensional thing to do.

You're looking at the animal in a particular state with a single type of stimulus in a particular environment, so the dimensionality of the thing that you're studying really depends on where in the brain you're studying it. If I was studying motor control, I'd see a lot of different types of actions during that resident intruder essay. If you're studying motivational states, you've got a pretty simple single axis of variation that you're looking at. I've started to see work looking at how these motivational states interact with each other. We have a--

Paul Middlebrooks
Individually? Within an individually, within an individual brain interaction with each other?

Ann Kennedy
Yes. I guess the work with Nick would be an interaction of hunger and pain. Pain, really, is-- it's not nociception. It really is a motivational and affective thing. We have a follow-up project where we're looking at interactions between hunger and predator threat, which is something that other folks have looked at. Interactions between hunger and aggression.

People have started to look at these pairwise interactions to really go to the full umwelt of an animal. All of the states that it can be in really would require-

Paul Middlebrooks
That's impossible.

Ann Kennedy
-care-- yes. It would take careful experimental design to collect a dataset that's really going to be usable.

Paul Middlebrooks
To answer those-- that question in-

Ann Kennedy
To study those things.

Paul Middlebrooks
-the totally ecologically valid sense. At the same time, it's like if you have the perfect model of a cat is a cat.

Ann Kennedy
Preferably the same one.

Paul Middlebrooks
We're always abstracting. What's that?

Ann Kennedy
Preferably the same cat.

Paul Middlebrooks
Preferably the same cat. With models or with theory, you're always abstracting so there's never going to be the perfect-- you don't just want to reproduce everything. You have to abstract to say something about the process that you're interested in. I'm not sure that you would want to reproduce the full umwelt of the affordances for the animal.

Ann Kennedy
I think in practice, you move one system at a time, one project at a time, and try to explain it and understand it, and then make some predictions about how it might interact with some other system. It's an iterative process.

Paul Middlebrooks
When you were describing the action of the neuropeptides and their downstream effects- that word causality is linked with effects, and I just said effects -anyway, you started talking about-- you mentioned in there how these control signals and the subcortical processes are context-dependent. It can sculpt or shape the way that neural activity happens in other brain areas.

One way of thinking about that, so there's the whole manifold push, everything's a manifold these days in low-dimensional population studies of neural activity, and the word subspace has become more popular now, so therefore, I think of everything, when you can have one population of neurons and a bunch of different subspaces. Which means that different pools of neurons are acting in different ways depending on what kind of internal dynamics they have and what kind of input they're receiving.

Do you think of it in terms of subspaces of orthogonal population activity, different manifolds in the state of possible manifolds, or do you think of it as just like a game, increasing the firing rates of everything, decreasing? That sort of thing.

Ann Kennedy
I think this is where there's room for theory, is just figuring out what is the space of things that you can do when you're modulating activity of a pool of neurons. If you're dumping on a neuropeptide, are you just-- there's a set of cells that's responsive now that wasn't responsive previously, so you have just recruited new cells into your computing population? Are you changing the excitability of everybody?

It's something that I think theory can really contribute to. We've done a little bit of work in the space thinking about heterogeneous neural populations and how adjusting the heterogeneity of the excitability of neurons can change the way they compute as a reservoir. My post-doc, Richard Gast, worked on this. We added a pretty simple form of heterogeneity to neurons, which is just we said we have a population of spiking cells, and rather than all having the same threshold, we have a distribution of spike thresholds. Some cells are a little bit easier to excite, some cells are a little bit harder to excite.

Then we asked, as you change the width of this distribution, how does it change the dynamics of the population? We initially did this thinking about just variation within a cell type and the fact that if you look at a given population of cells, they're not all the same. There's some variance to them. We talked to Luca Mazzucato, who's been thinking about a very similar model for understanding the effects of acetylcholine on cortical dynamics. Luca pointed out, which I agreed completely, that this distribution of spike thresholds could really be changing on the fly depending on availability of some ligands.

If you dump acetylcholine onto cortex, some cells get depolarized and some cells get hyperpolarized. The more you dump on, the more this distribution broadens. You can really change this-- the width of the spike threshold distribution on the fly using a neuromodulatory signal. In Richard's work, what he did is he took the spiking network of heterogeneous neurons and derived a set of-- a mean field model that described the firing rate of the population and its behavior.

With this reduced mean field model, he can perform bifurcation analysis and look at the different computational regimes that you can push the system into depending on how much input you're feeding it and depending on the degree of spike threshold heterogeneity of the population. It was really fascinating. You can produce pretty profound effects on just the transfer function of your neural population by adjusting spike threshold heterogeneity across the pool of neurons.

In a pool of excitatory cells, you can linearize their responses. This makes activity higher-dimensional and better for things like function generation, but it reduces this nice bistable regime that a homogeneous network has that's really useful for things like working memory.

Paul Middlebrooks
In a super high heterogeneous regime, you push it toward a linear regime?

Ann Kennedy
Yes, a linear input-output transfer function.

Paul Middlebrooks
Yes, transfer function.

Ann Kennedy
It becomes higher dimensional, can be entrained by a wider range of frequencies. It has richer repertoire of-- Oh, yes, break that apart.

Paul Middlebrooks
For capacity, maybe, right?

Ann Kennedy
Yes.

Paul Middlebrooks
Isn't it like if you push it too far, then you can't do anything because it's too high-dimensional? Am I reading that right?

Ann Kennedy
Yes, if you push it too far, you have neurons that are spiking all the time and neurons that are silent, so we get into this extreme regime where your

spike thresholds are so crazy that you just lose cells in your effective computational population. There's a pretty decent range of distribution widths that can give you very different behaviors.

Paul Middlebrooks
It's not one sweet spot, it's a sweet range. [chuckles]

Ann Kennedy
[chuckles] Yes. It's another knob that you can turn to change what your system does. If you do this in the inhibitory cells, it's very different. If your inhibitory neurons are all the same spike threshold, are all carbon copies of each other, you very often fall into these bursting regimes of the network where you get epileptiform activity, making the inhibitory neurons more heterogeneous, unmasks the bifurcation structure of the excitatory cells.

If you're doing this on the fly, you can do maybe a gating of information, where by changing how variable your inhibitory cells are, if you make them more heterogeneous, now the transfer function of your excitatory cells shines through. Then you switch to make them more homogeneous, and you'd lose that transfer function, and you fall back into the bursting regime. It's these means of controlling computation in neural population without fine-tuning the synaptic weights and supervised learning. It's just this thing that you could do out of the box.

Paul Middlebrooks
I was thinking about subcortical processes as dials, knobs, things like that. This reminds me of Mac Shine's work from a few years ago as well, where he essentially, he and Michael Breakspear, whom will be actually tonight, I release that episode, who were on my podcast recently. They controlled the gain, essentially, I guess the spiking threshold, essentially, in their model of populations of neurons, and found there are these scale-free dynamics-

Ann Kennedy
Oh, nice.

Paul Middlebrooks
-and there's these sweet spots. It reminds me a little bit of your work. In his case, he was mapping it onto the ascending arousal system. Was it acetylcho--? No, it's noradrenaline. If you turn up or down the spiking threshold of neurons with the noradrenaline projections or whatever, you get that same sort of sweet range. Not the same sort of, but a sweet range within which there's lots that you can do, high capacity. You can think of the ascending arousal system as a control knob in that sense.

Are all these subcortical processes, do you think of them as control knobs? How do you think of them? Is this a cybernetics kind of view of subcortical processes?

Ann Kennedy
Yes, I think it's an open question of how high-dimensional these control knobs are.

Paul Middlebrooks
How high-dimensional are they?

Ann Kennedy
[chuckles] We've got a lot of channels to communicate with, although, a fairly small number of ways that those channels can influence postsynaptic cells. We have hundreds of G protein-coupled receptors, but then what happens within the cell is fairly restricted. It's this interesting-- I think that the reason for this is that it gives you specificity. If you--

Paul Middlebrooks
Can you elaborate on that? You have a bunch of different types of receptors on the membrane of a cell that's receiving input from some sort of subcortical structure, G protein-coupled receptors. That seems like a high diversity, but then what you're saying is internally, there's a bottleneck sort of, of what can happen?

Ann Kennedy
Yes. Inside the cell, G protein-coupled receptor, you have the receptor part, which is specific to its ligand or not. Well, has a set of ligands that it can bind. Then inside the cell, you have the G alpha, G beta, G gamma subunits. When the receptor gets bound, these subunits separate and go off and do their thing inside the cell. There's way more of the receptor types than there are subunit types. Once you activate the subunits, mostly, they all converge on modulating the amount of cyclic AMP inside of the cell. Possibly--

Paul Middlebrooks
Why do you want that dimensionality reduction? I'm sorry to interrupt.

Ann Kennedy
Sorry? What was that?

Paul Middlebrooks
I'm sorry to interrupt. Why do you want that dimensionality reduction?

Ann Kennedy
It's a good question. It's still an open question how much of a dimensionality reduction it is. Is cyclic AMP like a pool within the cell, or are there compartments where you have cyclic AMP in one part versus another and they can function independently of each other? It's the same thing with everything in biology communicates through calcium. How do you use calcium to modulate action potentials, and transcription, and everything else? Are there compartments?

That's always bothered me that you have this capacity, that everything just runs through this bottleneck of calcium ions inside the cells.

In terms of the diversity of receptors, I think that-- There's this concept in evolutionary biology called Ohno's dilemma, which is if you have a gene that does X and you need it to do something else. Need is implying some agency that's not there,-

Paul Middlebrooks
Sure, yes.

Ann Kennedy
-but you have a gene that does X, you'd like it to also do Y. It's hard to mutate it to do Y without breaking X, and so what you do is you make a copy of the gene and you mutate the copy. Maybe you express it in a different part of the brain. Maybe you change its binding affinity for a ligand or change it to prefer a different ligand entirely.

Paul Middlebrooks
Evolution does this, you're saying, or within lifetime?

Ann Kennedy
Evolution.

Paul Middlebrooks
Evolution, okay.

Ann Kennedy
It is a way to take something that is useful, like a cell surface receptor, and increase the range of things you can hook that receptor up to. You're recruiting the same downstream signaling pathway, but now you can do it in a different subset of cells by turning on this receptor. You can do it with a different threshold for activation or a different time scale of activation.

We see dozens of receptors for serotonin or acetylcholine, or at least handfuls of receptors for a lot of neuropeptides. That could be a way of achieving specificity, of making sure that the right subset of cells is listening to those signals. Then what you do once you get the signal is pretty consistent and conserved. That's my hope, is that it feels like a big biological mess, but then the range of things that you're doing to modulate your cells is fairly small.

Paul Middlebrooks
It is a big biological mess. It's mind-spinning, so I'm glad people like you are working on it. My mind can only contain three acronyms. After that, I'm done. The world of neuropeptides is I wouldn't even know how to begin.

Ann Kennedy
The hope is it doesn't matter. Really, what it is, is just which entries of your weight matrix are non-zero. Then this cell can communicate with this subset of cells. Once it communicates, the thing those cells do is fairly constrained. It's, again, it would be a way of baking in connectivity to the brain without needing to learn it through experience.

If I have one set of cells that releases a certain molecule and another set of cells that can listen to that molecule, I can form interaction between those two populations, without needing to guide an axon to a dendrite in a very specific way. If I make the expression of the signaling molecule sensitive to animal state, to food restriction, or social isolation, or stress, then I could have some sort of-- The language is, in evolutionary biology, is phenotypic plasticity. You have phenotypes that are contingency plans, like circuits and patterns of interaction in the brain that only become available when you need them.

Paul Middlebrooks
Yes, on the fly. Are these types of things that we're talking about, this is what you're partially pointing to in the commentary that you wrote about theoretical neuroscience having room to grow. Again, just going back, traditionally, theoretical neuroscience has conceived of brains as these sort of symbolic computational processes doing tasks, but you're really bringing it into, A, the naturalistic world now. Like the rest of the field is pushing it into these ecologically valid kinds of behaviors, et cetera. Oh, my microphone just fell down.

Ann Kennedy
Oh, no.

Paul Middlebrooks
Hold on one sec.

Ann Kennedy
[laughs]

Paul Middlebrooks
Also, you're looking at across levels and down into the subcellular processes, and the way that these subcellular networks of communication and signaling. Is that where you see theoretical neuroscience has room to grow?

Ann Kennedy
I think it's one of the places, yes, that there's just a lot of room for theory and hypothesis generation here. How do you want these signaling systems to be organized? The arcuate nucleus seems like it has this kind of hub-and-spoke deal going on where you send different projections to different targets that you want to modulate by when you're hungry.

Paul Middlebrooks
What's the arcuate nucleus?

Ann Kennedy
That's where your hunger-activated neurons are. The ones that release neuropeptide Y. It's one part in the brain that broadcasts the signal to many areas. "Is that how you always want to do it?" If you look at the Tachykinin-2 system, it seems very different. The cells that express Tac2 in isolated animals are all over the place. All of your inhibitory interneurons are producing it, so there's no one hub that's sending out a signal. It's genes that turn on in different sets of cells, maybe under different circumstances.

There's room to just think about what principles could be shaping these systems. How you would best use them if you wanted to use them to make your neural networks more flexible and expressive. Also, how do you use them in a way that doesn't break what the networks are supposed to be doing in the first place?

Paul Middlebrooks
You mentioned the word principles. Are they going to be the same sets of--? I hesitate to return to dynamical systems view because I talk about it a lot, right? You can apply the dynamical systems to any network of interacting parts. That applies to traditionally what we do in the neurosciences, like the spiking neural networks. It also applies to molecular networks, to subcellular protein interaction networks, et cetera. Are the same principles going to apply in those systems, or will there be a new set of principles, depending on the level that we're looking at?

Ann Kennedy
Yes, I think if we take these very molecular, subcellular things, we can still think about principles of using them on a circuit scale. For example, in my work that's in press with Nick Betley, this hunger suppression of pain--

Paul Middlebrooks
Isn't it accepted now? Is that-

Ann Kennedy
Yes.

Paul Middlebrooks
-you mentioned? Congratulations, officially,-

Ann Kennedy
Thank you.

Paul Middlebrooks
-by the way. Yes.

Ann Kennedy
I'm excited. You have this release of neuropeptide Y when you're hungry and it's blocking pain. You can ask computationally, what is happening here? Should we think about this as--? The first way that we approached this was like, oh, when you're responding to pain, you're expending effort to, in the case of the mice, just freeze or lick your wounds.

You're doing that instead of doing other things, like foraging. If you're hungry, maybe you have a part of your brain that is comparing how hungry you are to how much pain you have and making a decision based on these relative levels. When you're hungry, the hunger bar is higher. That could be one way of thinking about this signal is it's going into some comparator that's deciding what to do.

Paul Middlebrooks
Do you think it's that simple?

Ann Kennedy
That was what we were hoping we'd find going into it, but you can imagine it working otherwise, too. Like you have some motor control part of your brain that's saying, "I'm in pain, I need to lick my wounds and stop foraging." Hunger could be just shutting off the output of that. It could be saying, "Other stuff is more important. Don't activate these motor programs right now, you've got other stuff to do." Also seems like it makes sense.

Paul Middlebrooks
Like a stop signal.

Ann Kennedy
Yes. Then the other thing that we thought about is you could be blocking the pain signal from getting to the brain in the first place. NPY could be shutting out the pain signal, so that if you have your policy, this abstract distributed the thing, the process by which the brain decides what to do next, if you just don't let the brain know that the pain is there in the first place, then it does other stuff. We built a little toy RL agent that had a pain state, and an effort state, and a policy. We said like, "Which of these things should hunger be modulating?"

The only thing that really matched this reduction in the pain coping behavior was blocking the input to the system. The other thing is because the pain input was so persistent, if you raised your threshold to start responding to it, you'd wait another couple of seconds, and then you'd start responding. Or, if you increased the cost or blocked the action, you just start spamming that action until you got one through.

Gating the input was a much easier way to control what our system did compared to controlling the policy or the output. It also makes different predictions about what the animal is feeling because if we took a model and we blocked the motor output, it was actually in more pain than in the sated case and was in--

Paul Middlebrooks
As assessed by behavioral observation, right?

Ann Kennedy
Not by behavior. The behavior was the same, but if you look at the internal variables of pain and efforts that were used in making a decision, your pain would rise to a higher point before you're producing the actions. The same was true if we messed with the policy itself. We're basically saying, "You have to be in more pain before you act." If you see an action, it's actually communicating more underlying pain. Whereas, if you're gating out input, then your underlying pain state is actually lower.

When Nick looked at parabrachial nucleus neurons, there were these cells that in a sated mouse had this long-lasting persistent activity during the inflammatory pain state that were just gone in the food-restricted animals. It really seems like this is the percept of pain.

Paul Middlebrooks
Wow.

Ann Kennedy
Maybe.

Paul Middlebrooks
[laughs]

Ann Kennedy
Yes. We didn't write that in the paper.

Paul Middlebrooks
I thought it was C-fibers. Isn't it supposed to be C-fibers? Isn't that--?

Ann Kennedy
That's nociception.

Paul Middlebrooks
Oh. Oh, sorry.

Ann Kennedy
Nociception and pain are different.

Paul Middlebrooks
Yes, okay.

Ann Kennedy
Then you could say, "Hey, I'm going to develop a new drug, and I'm going to look at this population of neurons to see how much pain my animal's in. I don't have to watch it or I don't have to prick it in the paw to see if it's in pain or not." It actually has huge implications for the field of chronic pain treatments because now you have a neural basis of the state that before we had to really do something to the animal to read out. A person who's feeling chronic pain, it's not that you pinprick them and they say, "Ouch," they're feeling it, even if they're not communicating it.

Paul Middlebrooks
This is the point of the podcast where we take a break for a commercial and I sell NPY to the techbro manosphere.

Ann Kennedy
Yes.

Paul Middlebrooks
The other thing that reminds me of is that that's similar, cutting it off at the pass in the psychological literature of habit formation and changing habits. The way to effectively change habits is, so you have the cue, you have your response to the cue, and then you have the reward, which is the high that you feel or whatever, if it's drug-related, et cetera. There's that loop or whatever. If you take the cue away, that's a super effective way. Don't go under the bridge for the heroin or whatever. Just stay away from the bridge. That's the same sort of principle, I suppose, at the neural level.

Ann Kennedy
Yes. It seems like it's possibly a general thing. There's other cases where we see a change in animal behavior that's been linked to a change in sensory processing. For example, Catherine Dulac has this beautiful work on pup-directed aggression in male mice. Male mice are typically infanticidal. If they encounter a mouse pup, they'll attack it. If they encounter that mouse pup at a time that is one mouse pregnancy after they've mated, so that it could be their pup, there's a shift in their behavior, and they show more parenting behavior, as opposed to just attacking the pup.

The shift in behavior is linked to a change in the processing of sensory cues from the pup, pheromonal cues. When I went into this, I was thinking, "Oh, the hypothalamus is this big policy. It's weighing these things off against each other and making a decision about what to do at a given moment. There's magic and fanciness happening there." It could be that really the best way to change your behavior isn't to change the interaction of these drives or these need states, it's to change the way you're sensing your environment.

Paul Middlebrooks
Elaborate on that, just so it's clear in my head. It's not the interactions of the drive states within the hypothalamus, it is the way that you're reacting to the sensory states.

Ann Kennedy
Yes, at least that's how it worked out for this hunger and pain story, is that it was easier to change the behavior of our little RL agent by changing its input, as opposed to trying to manipulate the policy or the output.

Paul Middlebrooks
Okay. Wait, actually, before we move on, I didn't ask you this, and the listeners need to know. What is the hypothalamus known for? Traditionally, what do we think that the hypothalamus does?

Ann Kennedy
Right. There's a whole bunch of things. There's people who study it for its role in growth and development in puberty and control of hormone release. There are people who look at it for its role in feeding and metabolism. Then there are parts of hypothalamus that are involved in things like predator defense, reproductive behavior, aggression behavior. In general, it's very deep in the brain. It's this set of interconnected nuclei. They express all sorts of different signaling molecules and receptors, and they seem to be involved in regulation and control of survival behaviors.

Paul Middlebrooks
That's what I was going to get at is I want to get your thoughts on the connection between these subcortical processes, the things that you study, these survival behaviors, pain, motivation, hunger, these things that are traditionally thought of as these basal level cognition. Almost not even cognition. It's almost like this will take over if you need it.

What's beautiful about humans and some mammals is that we can ponder. That's what our cortex is for. We have the beautiful thinking processes, and we can imagine scenarios and simulate in our minds, and we have models of the world, et cetera. Yet, these lower subcortical processes, subcortical structures are so powerful to our behavior, and they're affecting on our ongoing behavior.

Ann Kennedy
Yes.

Paul Middlebrooks
If I get my leg bit off by a shark, I will see the surfboard that my visual perception is still working, but I probably won't be writing poetry in my head in that instance. These higher cognitive functions aren't really happening in general. How do you think of that kind of cognition across the spectrum of what we consider cognition? Sorry, that was long-winded.

Ann Kennedy
It's almost a sort of working memory for your needs. It keeps track of how hungry am I? How thirsty am I? How stressed am I? Have I been in a fight recently? Are there other mice people around that I am stressed out about, or want to meet with, or want to interact with?

Paul Middlebrooks
Mice people? Did you say mice people?

Ann Kennedy
Mice or people.

Paul Middlebrooks
Mice or people. [chuckles]

Ann Kennedy
I don't know. It informs our actions, but it doesn't command them, right?

Paul Middlebrooks
You can override them, you mean?

Ann Kennedy
Yes. There's this really cool case. There's this woman in Scotland who, she's one of these people who doesn't feel pain.

Paul Middlebrooks
Oh, God.

Ann Kennedy
She has a mutation in the gene for fatty acid amine hydroxylase, something.

Paul Middlebrooks
Well done.

Ann Kennedy
She doesn't feel pain, but she also reports that she's never felt angry or afraid. She's just super chill. This is a gene that's expressed in the hypothalamus, and so she's still a person, but she doesn't have this sort of persistence of these drives, these aggressive and anxiety-related drives, the way that you or I do. It seems like it's very much suppressed in her, maybe because this mutation is reducing the ability of her hypothalamus to rev up its activity and produce persistent firing and persistent motivational states. That's my hypothesis.

Paul Middlebrooks
Otherwise, she seems like completely--? She's not some excellent poet, right, or some--? Is she like otherwise, normal IQ, et cetera, normal cognitive? Yes.

Ann Kennedy
Yes, just very chill. She gave birth and was like, "That was weird."

Paul Middlebrooks
[laughs]

Ann Kennedy
She can recognize when people are in pain or are distressed, but it's like a cognitive thing for her. There's no gut feeling. Yes, it's emotions that are affected as opposed to thinking and actions.

Paul Middlebrooks
The whole embodied cognition movement, the four Es, I can never remember which-- Naming the Es, I can never remember all of their names at once, but we're going to call them the four Es. The push is that our cognition is connected to our bodies, right? We shouldn't think of our brains as

separate from the bodies. We shouldn't think of our brains as manifesting our cognition because everything we do is part of all the same system connected to our bodies.

The things that you study seem to be close to what's going on in our bodies, which are signaling these brain areas to spit out these peptides that affect our ongoing behavior and cognition. Where are you in the brain-body dichotomy, non-dichotomy?

Ann Kennedy
It's a dichotomy? Do I have to pick one?

Paul Middlebrooks
No, no, that's what I'm saying is maybe it's a non-dichotomy. The four E people would say that there's no separation, right?

Ann Kennedy
Yes. I'll say the hypothalamus is in a privileged position. I forget the name of the thing, but there's a part of the brain where the blood-brain barrier is reduced and the brain is able to sense signals in the blood. The hypothalamus is sitting right there on that.

Paul Middlebrooks
Oh, I didn't know that.

Ann Kennedy
It picks up a lot of signals from the blood that don't make it to other parts of the brain. I think it's involved in, especially in things like hunger and thirst and sensing your nutritional need state, but also, lets you sense circulating hormones and other things.

Yes, brain-body interactions is a very hot topic right now. I think it's a really fascinating one that's, again, not a place where we've had a big neuro theory presence in the past. Yes, going back to what I was saying earlier about reservoir computing, the body is an excellent way to have memories, to have long-time scales that can influence neural computation. You have circulating signals that represents the state of the body, of available nutrients, signals about past physical activity that can signal from the body to the brain. Certain need states.

I think the role of hypothalamus, to some extent, is sensing those signals, and inferring need states, and using that to then broadcast to other parts of the brain what I should be prioritizing right now. Should I be paying attention to food smells? Should I be paying attention to other conspecifics around me?

Paul Middlebrooks
Part of what you've been doing also, so just changing topics as we get close to the end here, is the modeling of behavior. Tell me more about what you're doing in that vein. You're releasing a dataset, if I understand correctly, as well?

Ann Kennedy
Yes, pretty soon. I first got involved in this just because my postdoctoral lab wanted to automate scoring of mouse behavior because when you're studying social behavior to interpret neural activity, you need to know what the mice are doing. When I started, that was done by very patient postdocs or technicians going through and watching videos of mice interact and just manually labeling frame by frame what they're up to. I think that's still the case in a lot of labs.

Initially, it was just making life easier for the study of behavior. We worked for a while on computer vision systems for pose estimation and supervised behavior classification. It's still a thing that we're interested from that making life for scientists easier perspective. I think there's also just things to be learned from studying the dynamics of animal behavior in a more quantitative computational way. That's obviously a field that's really blown up over the past 10 years.

Paul Middlebrooks
Yes, there's all sorts of automated annotation tools right now.

Ann Kennedy
Yes. We, a couple of years back, started running into this problem where everybody who was publishing a paper on behavior classification or unsupervised behavior segmentation, they make their own in-house dataset, and they show that their method works really well on that dataset, and nobody-

Paul Middlebrooks
That's what we do.

Ann Kennedy
-can really evaluate how well it's doing.

Paul Middlebrooks
Sure.

Ann Kennedy
My collaborator at Caltech, Pietro Perona, had this long history of putting out benchmark datasets for the fields of computer vision and machine learning. His former trainee, Fei-Fei Li, put out ImageNet, which clearly made a huge difference in the field of computer vision. We wanted to do something similar, put out benchmark datasets for the field of behavioral neuroscience. We've run a series of multi-agent behavior challenges where teams would be tasked with solving certain computational problems in mouse behavioral datasets.

The first of these was just the same thing that we did in our mouse behavior paper. We published the tracking dataset of interacting mice, and we said, "Make us classifiers that can detect when animals are sniffing, mating, or fighting." When you have enough training data, you can do about as well as another human being at this point.

We released a second dataset, maybe '22, which was aimed at unsupervised behavior analysis and representation learning. It's hard to evaluate how good an unsupervised model is. You get some segmentation of your behavior trajectories, and you're like, "Well, cool." What we said--

Paul Middlebrooks
Without hand scoring the whole thing, you mean, or--?

Ann Kennedy
Yes. Is this representation of behavior useful? If you did a bad job on unsupervised segmentation, you might get clusters which correspond to where is the mouse in the cage because you didn't control her position.

Paul Middlebrooks
Which you don't care about. You care about what it's doing, no matter where it is.

Ann Kennedy
Yes, exactly.

Paul Middlebrooks
At least in some experimental setups, sometimes you do care.

Ann Kennedy
Yes, yes. What is useful is tough to pin down. We put together this dataset where we had tracking data from mice, and flies, and this ant beetle interaction dataset. We had a gauntlet of tasks that you could do on that data, like detect when flies are receiving optogenetic stimulation, classify the sex of the flies, detect when the flies are doing wing threats.

We kept those tasks hidden, and we said to the teams, "Give us a representation of what the animals are doing." Then, our scoring of that representation was if we train a simple linear classifier for our hidden tasks, how well does it do given your representation? That was maybe '22. That dataset is now out there, the posed datasets, and then also just raw videos of animals.

This year, we're trying to tackle this problem of generalizability of behavior classifiers. If you're a mouse social neuroscience lab and you're studying sniffing behavior, you train a classifier to detect when two mice are sniffing each other. If you give that classifier to another lab and they run it in their setup, it's not going to work. It's going to completely break because the frame rate of the camera might be different, the placement of the camera is different. It's just out of--

Paul Middlebrooks
The lighting is different.

Ann Kennedy
Yes. It's out of distribution. We emailed a bunch of different social neuroscience labs, and 15 of them contributed datasets of top-view videos of socially interacting mice, most of them with posed estimates, some of them we tracked ourselves, and then all of them with manual frame-by-frame annotation for some behavior of interest. Across the whole set of labs, I think we have like 38 different behaviors that people care about that they've annotated.

Paul Middlebrooks
Then you've used this to train a model, and therefore, generalize across labs?

Ann Kennedy
The goal of the competition is to say, how well can you train a model and generalize across labs? Because in practice, my lab has focused more on neural dynamics and less on machine learning tools, but we want to make these datasets available. This will launch on Kaggle in a week or two. You'll have the--

Paul Middlebrooks
It will have launched. You should just send me the link. It will have launched by the time this comes out, then, yes.

Ann Kennedy
Hopefully. It's been a slow process. We'll see.

Paul Middlebrooks
Sure. [chuckles]

Ann Kennedy
We're hoping sometime this month of September. Anyone who's interested can work on their own method to detect each lab's behaviors of interest from their closed data. We're hoping, not sure, but probably, if you train a model that can fuse data from different labs and learn the internal dynamics of mouse behavior in a way that's invariant to how a particular lab's camera is positioned and which body parts they trunk, that you'll get better performance in recognizing these social behaviors that we want to detect.

We put together this dataset for a pretty cut-and-dry supervised behavior classification problem. We're also really interested in if you can really build lab invariant representations of dynamics of mouse behavior. What can you learn about how animals make decisions and what information informs their decisions during social interactions?

If you have pose estimation, you have a decent proxy for the animal's sensory environment. It's not perfect, but you know where their head is. You can predict what they're seeing. You can use this to say like, given what my animal is experiencing, its location and its social context, can I predict what it's going to do next? We've done this in single animal datasets. We're looking to move into this, not published yet, but some work in progress.

We're really interested in this problem of forecasting, and using forecasting as a way of understanding the sensory motor transformations that are shaping our social behaviors and the context of our histories that shape our social actions.

Paul Middlebrooks
What is forecasting in this case? What does that mean?

Ann Kennedy
Forecasting is, given my animal, given its recent history, predict just where it's going to go next or what it's going to do next.

Paul Middlebrooks
In behavior world, that's called forecasting?

Ann Kennedy
Yes. It comes up all over the place in behavior world. It's one of the big limiting factors for self-driving cars, is being able to forecast what pedestrians and other cars are going to be doing. It's useful for, say, human-robot interactions. If you have a robot that's interacting with a person that needs to be able to predict how a person is going to approach it and interact with it.

It exists in a broad range of applications in machine learning and robotics. We're interested in forecasting because we think that building a not necessarily biological but a generative model that predicts how an animal is going to behave, will tell us something about the sensory information that informs those decisions. Also, on what timescales the experiences and histories of the animal matter for predicting what it's going to do next.

Paul Middlebrooks
You think that this is going to be possible based on the slower timescales of the states from subcortical processes or just from purely thinking behaviorally? Why do you think this is going to work?

Ann Kennedy
I guess the question is how well it will work and where we fall short. It's just one thing to ask, how deterministic is mouse behavior? If you see where a mouse is, can you predict what it's going to do next? From talking to experimentalists in my old lab, it seems like you can tell when a mouse is getting ready to attack. They become a lot more dirty, they're rattling their tail, their actions are changing.

You can probably infer state from moment-to-moment actions, but also if you know the history of an animal, that might influence your prediction of what it's going to do next. History, both what it's done in the past couple of minutes and history in terms of the sex and the strain and the experimental history of the mouse, which is information metadata that we have from the labs that provided these videos.

Paul Middlebrooks
That'll be available, that dataset. The Kaggle competition will have been released by the time this comes out.

Ann Kennedy
Yes.

Paul Middlebrooks
I'll point people to that. It seems like you're firing on all cylinders, obviously. What's holding you back? Is there something on the theoretical side, studying? Are there tools that's holding--? What is holding you back on the theory side? What's holding you specifically back on the theory side in the subcortical/subcellular signaling domain?

Ann Kennedy
I think it's maybe more data to constrain our thinking.

Paul Middlebrooks
More data. We need more data. This is the big data day.

Ann Kennedy
That's the cop-out answer. Yes.

Paul Middlebrooks
No, I wasn't judging. I was surprised.

Ann Kennedy
We see what a given population of cells in a given brain area is doing in whatever context the experimenter decided to study them, but we don't know what those cells are doing in every other context. It's very early days still for charting what these cells respond to. There's been, in cortical neuroscience, as mesoscale imaging has become possible, we've come to realize that these functions that we think are localized to specific brain areas, they actually have maybe not computations everywhere, but signatures of those functions are present everywhere.

Choice, you can decode from many parts of cortex. Actions, you can decode from many parts of cortex. Being able to look at the cortex as a whole during a task has really changed the way we think about local cortical computation. Subcortically, you can't do mesoscale imaging. You have to look at one population at a time, and you have to look at one behavior at a time. I'm really curious. I would love to see a broader range of contexts in which a given set of cells is recorded. Basically, higher-dimensional data.

Paul Middlebrooks
Even higher dimensional? Yes, okay. That's a good way to say it.

Ann Kennedy
Not number of neurons, but the-

Paul Middlebrooks
Different modalities.

Ann Kennedy
-state of the animal.

Paul Middlebrooks
Yes. Okay, got you.

Ann Kennedy
Because there's just a lot of unknowns.

Paul Middlebrooks
Yes. All right. There are a lot of unknowns, and we started talking about how one of the reasons you're studying what you're studying is because you saw it as an opportunity where the field didn't feel crowded.

Different people thrive in different environments, research-type question environments, but if you were starting over these days, so what I'm going to ask you is what advice you would give to someone who is beginning their career, what they should get into? In some sense, in the theoretical neuroscience commentary that you've written, you're sort of defining that space, and saying, "Here's a good place to go to," but I don't want to answer for you or put words in your mouth. What would you do if you were thinking about getting into theoretical neuroscience these days?

Ann Kennedy
Yes. I think there's which problems you work on, and then there's also which skills you develop. I think the skills have changed a lot from when I was a graduate student. Fitting models is no longer like a manual adjusting parameters until your model does what you want it to do and call it a day. We have tools for simulation-based inference that let you really see the whole space of performances of a model, and explore the capacity of a system for degeneracy and redundancy in solutions. I think just building your skillset is the most important thing you can do as a graduate student.

Paul Middlebrooks
Do you mean the breadth of skills? What do you mean skillset? In what?

Ann Kennedy
Building your technical skills. Yes. You don't have to do everything. I think it's easy to fall into a mindset where you just have to know more math than everyone around you and learn more math than everyone around you, that you never get to the applications. You have to strike a compromise. If you specialize in the biology too early, you might miss out on some of the computational skills you need to address biological questions that interest you. I think, yes, it's hard to give general advice because it really depends on where someone's coming from.

Paul Middlebrooks
Sure. Okay. Ann, thank you so much for spending so much time with me. I'm happy for you that you have found what seems to be a huge space of possibilities for your future endeavors. I appreciate your time.

Ann Kennedy
Yes, thank you so much.

[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]
