Brain Inspired Interview – Podcast by Paul Middlebrooks

BI 034 Tony Zador: How DNA and Evolution Can Inform AI

May 8, 2019

https://braininspired.co/podcast/34/

Paul M.:
Tony, welcome to the show and thanks for being here.

Tony Zador:
It's a pleasure to be on.

Paul M.:
If I'm right, this is your 20th year at Cold Spring Harbor Laboratory, where you're the program chair in neuroscience. Is that right? And if so, have you celebrated? Or, will you celebrate?

Tony Zador:
Indeed, it is just about to be my 20th year here at Cold Spring Harbor Lab. I did recently step down as program chair of neuroscience, but I was program chair for quite a while. They do have a celebration for those of us who make it to the 20th year. I have been invited and I do hope to attend.

Paul M.:
Oh, good for you.

Tony Zador:
So, I'm in the rarefied company of my boss, Bruce Stillman, and a handful of others, including Jim Watson.

Paul M.:
Yeah. So, you've been there since 1999. For those of you who are listening to this podcast, years and years from now. It is 2019 now. But, Tony, your research really has run the gamut in terms of levels of inquiry. From down at the level of molecules, to the synaptic level, to cells, to neural circuits, up to behavior. And you're also perfect for this show, because you really entered your career, you originally worked on the theory side of neuroscience and did work with neural networks, and we'll talk about your early participation at the NIPS conferences, or the NeurIPS conferences now, and how you started COSYNE out of your involvement in those early NIPS days.

Paul M.:
And through your career, you've incorporated more of the experimental side. Working with slice neurophysiology, studying decision making in primates, eventually, pioneering the study of decision making in rodents. And in your case, auditory decision making. And we can talk more about this. I think some of those obstacles that you faced in that work with the rodents, has led you to your connectomics work developing a DNA barcoding method, to figure out the wiring diagram of mouse brains. And so, I know that's an embarrassingly incomplete list there.

Tony Zador:
No, no, I was going to say, that's a fantastic summary of my career. And I think we can just call this entire conversation to an end, because I don't know what is left for me to say. I'll try to fill in a couple of the details.

Paul M.:
Okay. Well, good. I was going to ask you how you would characterize your journey thus far. But, if that's good enough, we can just proceed.

Tony Zador:

That's great. I think the overarching theme is that I have a short attention span.

Paul M.:
Sometimes I ask this question toward the end of a show. But, in this case, I'd like to ask you now, because you've had a definitive path, and everyone has a unique path. If you were going to go back in the early days of graduate school, now, let's say. Let's say you're starting off graduate school, or somewhere around there, what pathway would you lay out for yourself? Or, what would you start off in?

Tony Zador:
That's an interesting question. I'm pretty happy with the path that I took, even though, I did a lot of wandering. I mean, going back even further, I started off in physics, then wandered over to linguistics, then went to medical school, and then started in theoretical neuroscience. And all of those things and in different ways, contributed to what I ended up doing. And at each stage, I guess, I was driven by this tension between these fundamental questions that I was excited about answering, which are the big questions that many of us go into neuroscience to answer, what is consciousness? How does our mind work? Things like that. And my attraction to simple, rigorous answers.

Tony Zador:
And the problem is, there's a real tension there, that the bigger the question, often, the less satisfying the answer. And so, it's been a back and forth. I loved physics for the beauty of a simple equation that explains everything that you ask. I loved calculating how fast that ball goes down the inclined plane. But then, in the end, I didn't care how fast the ball goes down on inclined plane. And on the flip side, I would spend my evenings wondering, "Well, what is consciousness?" And I don't think we're that much closer to understanding that, than we were when I started thinking about that. So, it's that tension between interesting questions and satisfying answers.

Paul M.:
Yeah. I wonder if a better question is just how to approach the areas of study, when you go into it. Because, it takes a craftsman's mindset, if you will, because you have to be... Okay, there are two ways to approach things, right? One is, you follow your passions, and then, see where that leads you. And another is, you have an interesting question, you might not necessarily be passionate about linguistics, let's say, and maybe you were. But, you will still approach it with the mindset of, "Okay, I'm going to figure this out." And then, you can learn and become passionate along the way. Is that a useful distinction at all for approaching?

[Minute 10]

Tony Zador:
It fascinates me the differences in how different the many different paths there are to success, to successfully addressing scientific problems. What drives different people, turns out, as far as I can tell, to be very different. In my case, I'm not very good at doing things that I'm not passionate about. I didn't always turn in my homework on time. So, if I don't understand why I'm doing something, I just can't do it. If I don't understand why I'm memorizing something, I can't do it. So, in my case, I don't really have a choice. I get super excited by a question and I dive into it. And then, if I get disillusioned by it, then, I end up moving on. But, other people are much more disciplined. And they say, "Here, I see that mountain top that I'd like to be on. I understand that to get there, I'll need to do a bunch of things that I'm not that excited about doing, but I will do them."

Tony Zador:
And many of those people end up getting to the tops of some very hard-to-climb mountains.

Paul M.:
Yeah, that's true. That's a good analogy. Well, let's talk about history for a minute. I've had a few people on the show, who have lived through the early days of AI, and seeing that transition from symbolic AI to statistical based neural net type AI. Terry Sejnowski was on the show, and he documented his experiences in his book, The Deep Learning Revolution. I've had Jay McClellan on the show, and he shared some stories about the general antagonism of the symbolic AI pioneers toward the statistical parallel distributed processing and neural network folks. And you are well traversed as well. So, you were in attendance, and participated in the early NIPS conferences, Neural Information Processing conferences, which started in '87. I think, maybe your first abstract was in '90? '91?

Tony Zador:

Yeah. Actually, I'm a slightly younger academic generation than they are.

Paul M.:
I was going to mention that, you didn't have to chime in with that. I would have mentioned it.

Tony Zador:
Yeah. They were already established around the time that I started graduate school in 1986. So, the very first conference that I ever went to as a graduate student, was actually the precursor to what is now NeurIPS, which is the Snowbird conference on neural networks, held at Snowbird. I think, the one that I went to. So, that started in the early to mid 80s, from a bunch of neural network people at Bell Labs, including John Hopfield, and Terry was among those people. And that was a small invitation only meeting, and I was lucky that my PI was invited and I came along for the conference and the scheme.

Tony Zador:
And then, I started going to the NIPS meetings, that became my go-to conference, even more than SFN, more than the Society for Neuroscience meeting. Just because, back then, there were very few people doing what I was doing, which is, work at the interface between computational neuroscience and neural networks. And in fact, those two fields hadn't even fully diverged at that time. And so, that was, I think, to my mind, it was the best place to see computational neuroscience, that was the NIPS conference.

Paul M.:
And that lasted for a few years, right? Where, their computational neuroscience and neural network work, were really in a discourse. So, can you give the flavor of what it was like then, and maybe how it's changed since then?

Tony Zador:
Yeah. To my mind, back then, the question I was interested in then, and to some extent, it's the question that has been what has driven a lot of my work, is, what have the conventional neural network models failed to capture about how the brain works? And so, that question really has two flavors. Can you use neural network like models to model the brain? And can you look at the brain and use them to build better neural network type models that compute more effectively? And now it's easy to articulate those as two separate questions. But, at the time, I certainly didn't have quite as much clarity, that those were two questions. And I think, a lot of the people in the field considered them to be similar questions. And many of the same people worked on both, including Terry.

Tony Zador:
I mean, Terry's lab is one of the preeminent computational neuroscience labs, and he is one of the people who contributed in a lot of ways to basic neural network theory. So, back then, those were not fully differentiated. And it was only later that the NIPS community started pushing the argument that, "Look, yes, birds..." Everybody would trot out the birds analogy.

Paul M.:
We're going to talk about that later, too. Yeah.

Tony Zador:
Yeah. So, just for your listeners, birds fly but we don't necessarily look to birds to build better airplanes. And that became the prevailing wisdom at NIPS. And at that point, the neuroscience community, the computational neuroscience community at NIPS shrank and shriveled up for the most part.

Paul M.:
Well, let's just interject, because I want to talk about this birds' thing real quick. Because, I swear to you, this happened to me before I even read the paper that we're going to talk about today. Because I've had multiple guests on the show use the birds' analogy, and I've nodded my head, "Yeah, that's a good analogy. You're right." And then, I was on a run the other day, and I was thinking, it's a really a poor analogy. Because, sure, if you want to talk about how airplanes "fly," we build wings and use propellers, and it can fly.

Tony Zador:
Nowadays, we even use jet sometimes.

Paul M.:

Or, turbine engines, jets. But, if you compare that to the diversity of how flies, insects and other birds fly, it's not even a comparison. But anyway, I thought of this, and then, I read your paper, and in the last paragraph of your paper, you say that you articulate these same ideas. And I thought, "Oh, that guy, he must be a genius, because he thinks like me." That's what I thought.

Tony Zador:
Exactly. Yeah. No, I mean, the birds analogy is flawed. And actually, I think, another one of your guests, Matt Bottvinick, provided some history I didn't know. That, apparently, some of the early aeronautics engineers actually did take specific inspiration from birds, beyond the fact that they provide an existence proof for flight. Right? I don't really know the history of aviation well enough to comment on that. But, I was gratified to hear that the birds' analogy is flawed in other ways as well.

Paul M.:
That's right. And many of those early people who took inspiration, probably did die. But, that's for the history books.

Tony Zador:
While flying.

Paul M.:
Well, while attempting to fly.

Tony Zador:
Oh, attempting to fly.

Paul M.:
Yeah. Okay, anyway, sorry for the diversion. But, so, then, through the bird analogy, the NIPS people shrewd away all the computational neuroscientists.

Tony Zador:
For the most part, yeah. I mean, I don't think it was an explicit decision, but the community there began to shrink to the point where it was not... I mean, I think there was always a neuroscience track at NIPS. And honestly, there might have been people who continued to go there for that. But, I found it less appealing and so do the most of the people I knew.

Paul M.:
You found it so unappealing that you started your own conference, the Neural Information and Coding, or NIC's conference.

Tony Zador:
That's right.

Paul M.:
That was an invitee only affair as well, correct?

Tony Zador:
That's right, that's right. So, inspired actually by the Snowbird conferences, I started a series of... when I was still a postdoc, I started a series of invitation only conferences on computational neuroscience. But, that differed from some of the other computational neuroscience meetings, in that, we really tried to always invite a mix of theorists and experimentalists. And those went on for, I guess, the first one was '96, at Jackson Hole. And then, we started to have a series of them at Snowbird and a few other... Coincidentally, they were all ski resorts in the West. And, they were typically held in February. But, that was just pure coincidence. But, I think they were very, really influential meetings for me and a group of people who went back year after year.

Tony Zador:
And they continued to go on, even after I stopped being the primary organizer at some point, several other people organized it. And they continued till, I think, around 2002. And then, at that point, we decided that the demand in the community for these meetings exceeded what you could do with a small invitation only meeting. And then, we rebooted as, what is now the COSYNE meeting, which is not invitation only, it's open to everyone. And it has abstracted its model directly after NIPS.

Paul M.:
And it's growing like NIPS, I think, right?

Tony Zador:
It's growing like NIPS was growing before the AI explosion, the modern machine learning explosion. We started off at around, I guess, 300. And this year, we topped, I think, 1,100.

[Minute 20]

Paul M.:
Yeah.

Tony Zador:
Yeah, our goal was not to grow, but we are growing. And so, we're figuring out how to handle that exactly.

Paul M.:
I mean, I guess, that's a good problem to have, in some respects.

Tony Zador:
Yeah, it changes the flavor of a meeting. But, that's a whole other discussion. It reflects the fact that there are a lot of people interested in computational and systems neuroscience. So, that's what the COSYNE. C-O, as in, computational. S-Y, as in, systems. N-E, as in, neuroscience.

Paul M.:
Like you said, I know that the NIPS conferences these days, because of the AI explosion, sells out in minutes, right?

Tony Zador:
Yeah, that's my understanding. I haven't been to NIPS literally in 10 years. So, from what I understand, it's hard to get tickets to that than to a very popular concert.

Paul M.:
Both of which I'd probably stay away from. Okay, last question about some history and the current state of affairs. So, it's not NIPS anymore, it's NeurIPS, right?

Tony Zador:
That's right.

Paul M.:
And even in Terry's book, he refers to it as NIPS. That's how recently the change happened, I guess. When I heard, "Oh, they're changing, I guess, acronym to NeurIPS," I honestly, I'm so naive, I couldn't figure out why, until I heard that it was because NIPS can be short for nipples, and that's the reason, right? Do you have a take on the necessity of changing it? Or, was there some chauvinistic culture that really shadowed it? Or, what? Do you know?

Tony Zador:
I mean, I only, I read a couple articles about it, the same ones that are publicly available. Yeah. I think it's no secret that, especially the tech world is prone to juvenile jokes, and is not always welcoming to a diverse group of people. And so, I think this was by popular demand, in order to help change the culture. I have no idea whether it's successful, but it's certainly worth trying.

Paul M.:
Yeah. Well, and so, we advance. So, let's talk about evolution here. Actually, the paper we're going to talk about was sent to me via Twitter. Tim, on Twitter, suggested that I cover this paper and have you on the show. So, thanks, Tim.

Tony Zador:

Thank you, Tim.

Paul M.:
Something actually useful came from Twitter, imagine that. Okay. So, A Critique of Pure Learning: What Artificial Neural Networks can Learn from Animal Brains. Is it a biorxiv?

Tony Zador:
Yeah, I put it on biorxiv.

Paul M.:
biorxiv, for now?

Tony Zador:
For now, it's under review. So, hopefully, it'll be in a referee journal at some point.

Paul M.:
Okay, great. So, let me see if I have the take home right, and then, you can correct me. Over evolution, DNA provides a massive amount of inherent learning that manifests as the wiring principles of our brains, and thus, sets up a massively pre-trained system, essentially, that we can then learn on top of during our lifetimes. Whereas, any current deep learning or AI system, supervised or unsupervised, suffers from having to learn everything from scratch. So, in broad strokes, is that accurate? What am I missing there?

Tony Zador:
In broad strokes, that's right. Part of the goal of the paper was really to point out to people who work in machine learning, that, much of what they call learning, would not, to a biologist, seem like learning. A lot of it came from miscommunications that seem to arise because people in different fields use words in different ways. And it was my attempt to lay out cleanly, what I think most biologists would agree is happening, most neuroscientists would agree is happening in real organisms. Namely, that, most organisms come with a great deal of innate capacity to perform well in the environment. And how much that's in contrast with how machine learning people see what the problem is, the intellectual problem.

Tony Zador:
So, they usually, the problem formulation, at least, traditionally, starts with a tabula rasa network. And from an evolutionary point of view, it would be insane to not build in structure, if it were possible to do so. Right? If you have two organisms, one of which comes out of the box, born with the capacity to navigate the environment effectively, and the other has to learn everything from scratch each time. If it were possible to build in stuff innately, that organism would be selected for. And in fact, I would argue that, that is what happens. So, that's really the core idea behind the paper is trying to clarify for both communities, how they formulate the problem, and to recognize that a lot of what is called learning, in the machine learning community, would not be considered learning in the neuroscience community.

Paul M.:
This is an ongoing issue, I think. I mean, even calling artificial units, neurons, is problematic just in a conversational way. It's difficult.

Tony Zador:
Indeed, indeed. In fact, my thesis, my PhD thesis was on trying to understand the ways in which actual biological neurons differ from the summation units of neural networks. My thesis was on dendritic processing in real neurons.

Paul M.:
Yeah, well, okay. I mean, there's a lot to sort out moving forward in both fields, or the crosstalk between the fields. Okay, Tony, I mean this in the best way possible, you're a rodent guy, meaning that you've studied auditory decision making in rodents and sensory processes. And in the paper, you argue that, if we figure out mouse level AI, that essentially we'll be right super close to human level AI. How is that? Do you want to just expound on that idea for a moment?

Tony Zador:
Why do I believe that?

Paul M.:
Yeah.

Tony Zador:
You're asking why I believe that. Why do I believe that, if we could just get to mouse level AI, we'd be a short step away from human level AI?

Paul M.:
Is it for grant funding? Is my question.

Tony Zador:
No. There's plenty of stuff that I do for grant funding. But, this actually is a fundamental belief that is why I study decision making in rodents. And for example, not in spiders, but on the other hand, not in people or monkeys. It's why I chose rodents. And so, basically, my argument is an evolutionary argument. That, animals have been on this planet for of order 500 million years or so, depending on exactly where you draw the line for animals, five, six, seven hundred million years. Vertebrates have been around for 400, 500. Mammals have been around for about 100 million years. And over that time, there have been something like Avogadro's number of individual animals, that evolution has operated on.

Paul M.:
99.9% of which are gone.

Tony Zador:
99.9% of which are gone. But, if you formulate the problem right, even the ones that didn't make it, have to be considered when you think about how the ones that did survive, what they have extracted. How much of the space has been sampled? The space of all possible networks. So, really, there's been a tremendous amount of evolution to get to vertebrates and to get to mammals. And the reason I'm focusing on mammals, is that, mammals all share what I consider to be the fundamental advance that enabled human intelligence, which is the neocortex. So, we've had the neocortex for about 100 million years. And our cortex, as far as we can tell, the basic principles are not very different between a mouse, a primate, and in particular, a human.

[30 minutes]

Tony Zador:
The difference between our primate ancestors and our modern rodent, probably, we have to go back more than 10, 20 million years, to see really big differences. And so, the things that really make humans special, in terms of their apparent intellectual prowess, those things came in the last few million years. And I started off in linguistics, so, I'm particularly impressed by our ability to do language. And that, depending on whether you think Neanderthal had language or not, probably emerged in the last couple of hundred thousand or million years. We know that our closest primate relatives, apparently, can't master language. So, whatever that evolutionary jump was really very recent. And the primate population size is pretty small.

Tony Zador:
So, we're talking not very much evolution to get us from a neocortex to a creature capable of human type intellectual functioning. And so, that's why, I believe, that the hard step is to get to the general purpose cortical operation that rodents are capable of. I say, rodents, I mean, we could easily have done cats or dogs, it's just that rodents are the common lab species, tractable lab species for studying these things.

Paul M.:
You're saying that, I can't wait... If it is that recent of a step and that small of a step to get from our symbolic like processing, or to our symbolic like processing, I can't wait to see what's next for us evolutionarily, right? What our cortices are capable of.

Tony Zador:
Yeah. Although, it's not clear where these, any particular strategy asymptotes, right? So, there's this, those who believe in the singularity, think that, "Aha, as soon as we get to human level intelligence, we'll quickly surpass it." But, it's also possible

that, this kind of strategy tops out after something a little bit higher than human intelligence. You need a fundamentally different strategy that we can get to, using these types of approaches.

Paul M.:
Proof that we're at the tip of the evolutionary arrow, right? We're the cream of the crop. Supervised learning, the main type used in deep learning these days, that's associated with the AI explosion, is not enough to explain how we learn from early childhood. So, for example, kids don't need a million examples to learn what a car is, for instance, right?

Tony Zador:
Exactly.

Paul M.:
And one suggestion, as you note in the paper, has been, maybe that kids are built in with this powerful unsupervised learning system, really early on, and that can then give way to a supervised learning type system later on. You suggest that an early unsupervised learning system, early in the developmental process, is not enough to explain the behavioral skills of young animals and humans. And you give a lot of examples about how many of our learning skills are innate. Right? So, there's a possibility that the maximum absence of evidence is not evidence of absence, could be in play in this case. Or, do you think that it's just too wide of a chasm to imagine an unsupervised learning algorithm, that is powerful enough, that could get kids on the right track, animal kids and children kids on the right track early enough?

Tony Zador:
Well, in some sense, if you divide learning into supervised and unsupervised, and we agree that learning that is required for a child to learn language is not supervised, then, that only leaves unsupervised. So, in some sense, it must be the case that unsupervised learning is responsible for learning things like language or categories. The question isn't whether it's unsupervised learning, the question is whether, to my mind, is whether it's a general purpose unsupervised algorithm, that could learn anything in a comparable amount of time. And my belief is, that it's not a general purpose unsupervised algorithm, it's some unsupervised learning on top of a great deal of innate structure that we are born with, because of evolutionary selection.

Tony Zador:
So, the hardest place for us to see that at work, I think, is in humans, in children. Right? That, it's easy to see why one might believe that humans require or rely on a tabula rasa unsupervised learning approach. Because, humans are so helpless at birth, and then, they take years to get up to speed. Right? But, we are, I don't believe, because I'm a neuroscientist and a biologist, who believes that we are just an incremental step away from the creatures from which we evolved, I think the idea, my approach, anyway, is to go back and look at these earlier creatures and say, "Well, could that really explain what's going on with these other creatures, like rodents, or spiders?"

Tony Zador:
So, if we go back to insects, insects are born basically ready to go. And fish are born with a great deal of innate structure, ready to go. That's not to say, they function absolutely perfectly at birth. No, there's a little fine-tuning that needs to go on. But, for the most part, they're not spending three years getting up to speed with their environment, before they're ready to reproduce. They've got, hours, days, or usually, at most, weeks, before they're let loose in the environment. And so, to my mind, that is the argument that there has to be a great deal of innate structure. There's other arguments too. There are some cases where creatures do things for which they have zero examples, not one example, but zero examples. So, there are great number of innate behaviors that are really innate. At least, portions of them.

Paul M.:
Do I have this right? Did you use the squirrel jumping from branch to branch in the paper?

Tony Zador:
Squirrel jumping, I haven't studied squirrel jumping, but I imagine that they take small jumps, and then, larger jumps. No, but, for example, burrowing behavior.

Paul M.:
Burrowing, yeah, there you go.

Tony Zador:
So, burrowing behavior, there's this fantastic work by Hopi Hoekstra at Harvard, who studies different species of a rodent called Peromyscus, that, depending on the exact species, they build very different nests, very different boroughs. And some are long and complicated and have a bunch of side pockets, and some are shorter, and they all have characteristic shapes. And these are built in to the hardware. These are innate, in the sense that, if you take a pup from one strain that builds a complicated borrow, and let it be reared by a mom from the other one, that pup will grow up and build a complex, not a simple borrow. In other words, it builds the one that it's genetics endowed it with, not the one that it's mom would build.

Tony Zador:
And so, how do you encode stuff like that in a gene? Well, you can imagine how you might do that. But that's just a potentially experimentally tractable example of a huge number of innate behaviors. And so, to my mind, if you believe that there are things that can be inscribed into your behavior by genetics, then, I think there's a strong argument to be made that evolution would select for as many of those as are useful. And if you can build into my nervous system, a way of learning languages more quickly than my competitors, or more effectively, I will outperform. There'll be selection pressure for my offspring to thrive. And so, really, the question to my mind isn't whether it's possible to put these things into the genome, but rather, why is it that in humans, some of them, some of that innate stuff was left out, or so much of that innate stuff was left out?

Tony Zador:
And I think that the answer, of course, is that, there's a trade-off. The more innate behavior you have, the less flexibility you have to deal with new environmental situations, new environments. And so, we as humans, are probably as far away from maxing out on what's the need, as any creature in the animal kingdom. But, that said, we're still mostly innate.

Paul M.:
That's an interesting dynamic of losing the innateness to grant other abilities. But, we'll hold off on that for now. So, you distinguish between supervised learning, in the paper, for artificial neural networks, and what you call supervised evolution. I mean, this is along the same lines of what you were just talking about. So, what do you mean by the term, supervised evolution?

Tony Zador:
Yeah, that was just an attempt to be a little bit provocative. To say that, to think about what goes on in a supervised learning network, as supervised learning is no more accurate than it would be to call it supervised evolution. This supervised process is a really useful way for us, as people building technology, to get the network to do what we want. To find the network structure, the network weights that achieve some objective function, and that's fantastic. It's way more efficient than what I would say is likely to be the dominant way that evolution comes across it, which is basically a random walk. Right? So, evolution has the advantage of being able to operate on Avogadro's number of organisms over many, many years.

Paul M.:
Like a genetic algorithm.

Tony Zador:
Yeah, exactly. For the most part, people can quibble, but, in most cases, genetically algorithms are not as efficient as gradient descent. But, apparently, for a variety of reasons, evolution didn't use supervised learning. The problem formulation is not well adapted to the real challenges, to the problems that biological organisms face. So, it's fine that we use supervised learning, it's fantastic that people have figured out how to accelerate the process of finding the right weights in a network to solve a problem. That's fantastic, right? But that doesn't correspond, in my opinion, to learning. It's just a trick for finding interesting network structures that solve problems.

Tony Zador:
And so, we could apply that same approach to "evolution." Evolution in the lab, and we could call that, supervised evolution. And had the history of the field been different, had it been evolutionary biologists who figured out how to do what we now call machine learning, we might very well call it supervised evolution.

Paul M.:
I can't wait for the first supervised evolution lab by a future PI. So, could or should DNA, through evolution, then, be considered a mechanism of information compression?

Tony Zador:
Exactly.

Paul M.:
Or, like a bottleneck of information to accommodate such vast amounts of data?

Tony Zador:
Yeah. The way I see it, what's going on is, biological organisms have to compress everything they know, and where I use know in quotes, about the structure of the network that they're going to build into the genome. They have to pass network structure through, what I would call, a genomic bottleneck. And that's just a fact. Because, every organism arises from a single cell with one genome, everything about that organism's network structure has to come out of that genome. And so, that genome encodes a series of rules for wiring up a network. And those rules, which could, by the way, be use-dependent, some of those rules can be use-dependent, are the rules that provide us with the brains that we then walk around with.

Tony Zador:
So, I think that that passing the network structure through a genomic bottleneck, actually access, what in machine learning might be called a regularizer, it requires boiling down the most useful network motifs, into a relatively short description length. And then, that rigor, that constraint actually requires that, really, it's the most interesting network motifs that get passed along. And I think, really, it took a long time, but, eventually, one of the most fruitful network motif was the cortex. So, the neocortex is, I think, the brilliant result of many, many hundreds of millions of years of evolution.

Paul M.:
It's interesting, I had Federico Turkheimer on the show, this past episode, which I hadn't aired yet, since I've talked to you. But, one of the points that he makes in his work is how the brain, both functionally and anatomically, is fractal in nature, right? It looks the same on multiple scales. And it just struck me that, thinking about DNA as a bottleneck of information, I cannot remember the name, but there has been recent, fairly recent work on deep layered networks, like the kind that are used to mimic layers in the brain, write a visual processing, and how these can be thought of, mathematically, as information bottlenecks from one layer to the other, to the next in compressing the information.

Paul M.:
It escapes me, who that work was done by, I'll look it up after the show. Anyway, just an interesting side note that, could the DNA algorithm toward cortex and all the layers in between, be fractal in nature?

Tony Zador:
Whether it's fractal, yeah. I mean, there's the information bottleneck idea, which goes back to Naftali Tishby and Bill Bialek, I think, in a series of papers, probably starting around the late 90s. And the idea that, by compressing information down... So, it's a generalization of the idea of the autoencoder, that, by compressing information down, if you can send it through a compressor, what comes out the other side, if you play your cards right, it is only what's useful.

Paul M.:
Yeah, it's forcing what's useful. Yeah, interesting.

Tony Zador:
Exactly.

Paul M.:
A principle of the universe, perhaps.

Tony Zador:
Yes. No, I think that that's exactly the principle that evolution exploits. As an interesting side point, you might think, well, given that you have to pack everything you're going to have about the brain network into a genome, you might want to pack more. And the way to do that, would be to expand your genome. And turns out there are organisms that have genomes that are 100 times larger than the human genome. And so, there's no biological constraint that says we're limited by the actual size of our genome. And so, interestingly, I would say that the fact that our genome isn't 100 times larger than it is, is an argument that, it's advantageous to compress down the wiring diagram into even a tiny fraction of our existing genome.

Tony Zador:
Because, if it were useful to use a bigger genome, that would apparently not be a fundamental problem for biology to solve.

Paul M.:
For biology. But, let's say that we're somehow able to decode the pre-trained information that is in the DNA, in a way that gives deep learning the same start that an animal has upon birth, right? In some sort of network. I mean, could we actually be limiting ourselves, and thus, our AI systems, by the constraints that are compressed encoded into biological DNA, into the innate learning? Does that make sense?

Tony Zador:
Well, I guess the way I would say it is this, if we consider the set of all possible networks, that's a pretty big space to wander through. So, in the modern era of deep networks, we have already constrained ourselves to a subset of all those, right? So, even if you take a deep network that has millions of parameters, they're not organized willy-nilly, they're organized in a particular structured way. Right? I mean, hence the term, deep networks. Layer after layer and the layers of characteristic property. So, we're already imposing some structure. And there's, I think, a lot of theoretical work going on right now, trying to understand why that particular structure is a useful one. And people have some ideas, and I think they don't fully understand why.

Tony Zador:
But, it turns out that somehow, those kinds of networks are well suited to learn the kinds of things that you might want a network to learn. And they're particularly trainable in small numbers of... are relatively small. Smaller numbers of samples than networks of other structure. So, it's quite possible that there exists... So far, in the absence of theoretical results, to the contrary, it seems possible that there exists, somewhere in the space of all possible networks, ones that are even better. But I would argue that, rather than looking through the space of all possible networks, why don't we take our inspiration from the networks that we already know can solve the problems of interest?

Paul M.:
Spoken like a true neuroscientist. So, as you note, in the paper, transfer learning, or the ability for a neural network to be trained on one task, and then, transfer its learned skills, let's say, to a new task, that's an active area of work in artificial intelligence these days. But, this is very different from the way that DNA gets transferred from one generation to the next through time. So, can you just comment on that and talk about how-

Tony Zador:
Yeah, I think the big difference there, is that, in transfer learning, at least, as I know the field, the amount of information that one might use to transfer to the new network is not constrained. So, you can use as much information you could. Basically, the naive way to do transfer learning is you just take the weights that you trained it up, strip off the last few layers, and then you initialize the network to the weights of the first end layers of the network. And that's a lot of information. And so, I suspect that that will work for a smaller class of problems than the more general solution. Which is, again, to boil down the key aspects of that wiring diagram, through something that is a bottleneck, the genome.

Tony Zador:
So, I would say that, conceptually, that you might not want to relearn everything from scratch, every generation, but I think biology actually has turned what, at first, might seem to be a constraint into a feature, which is that it has to abstract the key aspects of that wiring diagram and fit it into a small bit of a genome. And, in fact, it has to do so, what comes out of the genome, of course, at least, in most creatures, is not a detailed wiring diagram of the form, wire neuron one to neuron three, seven, nine and to billion, but rather, a set of rules. And so, by forcing it through a genomic bottleneck, it sort of asks us to look at what kinds of wiring rules would be effective for that wiring.

[50 minutes]

Paul M.:

Right. Like the nematode, as you point out, its connections are actually completely instructed by the DNA. Whereas, maybe humans and mammals, this could be the DNA evolution's test on, let's take it from a complete wiring diagram to just some rules and some algorithms, suggestions.

Tony Zador:
Exactly.

Paul M.:
And so, we're seeing how that test is coming along. So, here I am again, Tony, I've asked this question multiple times in the past to many guests here. But, I've asked it with neuronal properties in mind, and anatomy in mind, and biophysical properties in mind. And now, I'm going to ask it with the DNA in mind. How biophysically accurate, how biologically accurate, in the details, are we going to eventually need to be, to make an AI that will work on the level of, in your case, let's say, a mouse?

Tony Zador:
So, my day job is being a neuroscientist. So, as part of my day job, I care deeply about all those details. As I mentioned, as a graduate student, I made models of processing in single neurons, dendritic processing. As a postdoc, I looked at the nuances of synaptic transmission. And now, as a PI, I study neural circuits. So, I care about those details intimately. But, I don't believe that there's only one solution. So, I don't believe that the solution that will allow us to build an intelligent machine, will necessarily involve single units that have dendrites and particular compliments of calcium and sodium channels arrayed in just the right way. The reason I study the details, is so that eventually, I can get to the underlying principles, abstract them away, and then, say, "Ah, what's important about all those details, is X, Y, and Z."

Tony Zador:
And so, my belief is that, almost none of the biophysical details will be important. The problem is, that without, at least, the only way I know how to figure out which details are important, is to understand in actual neuroscience settings, how those details work together to enable behavior. So, once we understand what's going on in a biological system, then we can abstract that stuff away, and say, "No, no, no, no, what's important is not the calcium channels, et cetera, what's important is X, Y, or Z." So, over the course of thinking about this, I've had different opinions about what really is important. And as a graduate student, I thought what was most important was that the simple processing units that are used in artificial neural networks are too simple.

Tony Zador:
And by the end of my graduate work, although I had a lot of fun learning how individual neurons worked, I concluded that, okay, really, it's not a qualitative difference. Sure, an individual neuron does something more complicated than a processing unit. But, by that, you could really just replace one individual unit with... or, one individual neuron with maybe a dozen units, and you're done. So, in that sense, there's no deep thing missing, it's just maybe, it'll take 10 times as many simple units as neurons. Okay. What I ended up concluding, was that, what is fundamentally different, is that, actual neural networks in the brain, real networks of real neurons, have highly structured connectivity. It's sparse, and the neurons are wired up in very specific ways. And that, really, it's just what endows a brain with its capacity is the wiring diagram.

Tony Zador:
And so, that is how I got from studying pure decision making to trying to understand at the level of the actual wiring diagram, how those decisions are made. At the level of, which neurons are connected to which other neurons, to produce this neural activity?

Paul M.:
I mean, there're so many different ways we can go from here, because you're doing some really great work with connectomics and just what we're talking about. Maybe, before we get there, and I can't take all of your time for the day. But, just thinking about levels, these levels of understanding and what it takes, I had John Krakauer on the show, maybe a month or two ago. And I don't remember if we talked about it on the show, or I had the opportunity to have beers with him beforehand. And I don't remember if it was over beers or on the show. But, he made the point, his outlook is that, you can use brain circuitry and brain activity, neuronal activity even, as a check on the validity of our understanding of higher cognitive functions.

Paul M.:
And there are hypotheses for how these higher cognitive functions work, let's say. But, you don't need to really understand

these lower level phenomena, you don't really need to understand the circuitry, the wiring diagram, the spike rates of all the neurons, for instance, to understand the higher cognitive functions themselves. So, do you think that understanding the constraints imposed by DNA, the Y, and then, the connectomics, the connectome, the wiring diagram, will improve our understanding of the mental or psychological higher cognitive, I'm doing air quotes right now, functions that we call our cognition?

Tony Zador:
Yeah. Obviously, I wouldn't be doing what I do, if I didn't believe-

Paul M.:
Well, no. But, you could be doing it because it's interesting in its own right.

Tony Zador:
You're right, I could be. You're right, it's not obvious. In my case, I think if we went back to the starts of our careers, we would probably find that we were interested in understanding the same fundamental questions. How do humans think? I just concluded that the way to get there, was to understand the circuits at a really low level, because those are tractable. It's, again, the trade-off between really interesting questions, and what, to my mind, are satisfying answers. So, there was just no question that you could get satisfying answers to questions about, how one neuron drives another neuron to fire? How do synapses work?

Tony Zador:
And potentially, I can tell you what the form of the answer is, for what a wiring diagram looks like. That's well defined, and I can tell you, in principle, whether or not we got the right answer. It's hard, I might not get there. But, at least, it's really well defined. And so, I would say that the argument that we can just look at the top level of cognition, I push back on that, because there is a name for the field where people look at that sort of top level, and traditionally, that is psychology, right? The very top level, where you look at human behavior, or you look even at animal behavior, without trying to pry open the black box, that's psychology. And that is, I believe, you can only get so far, without opening up the black box.

Tony Zador:
I don't think that people today are that much smarter than they were 50 or 100 years ago, I don't think we're that much better at designing experiments, so that we can expect without something new to make progress on that. And this, actually, he wrote a very provocative and controversial piece, I think, it was in neuron. John Krakauer did, a couple years ago, maybe two, three years ago.

Paul M.:
This is what we talked about when he was on the show. Yeah.

Tony Zador:
Yeah. And it made a lot of points, and I was sympathetic to some, specifically that one can't forget about behavior. But, the idea that one can only look at behavior, and maybe use neuroscience as a check afterwards, makes no sense to me.

Paul M.:
Yeah. That's not what they argued in the paper, really. And I'm speaking for him, and I shouldn't be, obviously. But, I think their point was just that, it needs more behavior to be more inclusive of behavior. But, yeah, we don't need that.

Tony Zador:
As someone whose labs spent the first maybe five years, along with my close colleague, Zach Mainen, figuring out how to train rats to perform the kinds of tasks that previously, people had only trained non-human primates to perform, I'm very sympathetic with the idea that one has to study behavior very, very closely. But, my belief, which I think is actually at odds with, at least, how I read that paper, is that, that's not nearly enough. And that, the way that, in general, science and biology, in particular, and neuroscience, even more in particular, move forward, is with new technologies. And so, having a well controlled behavior provides an interesting object to study in the right way.

Tony Zador:
But the tools with which we study it, those have to, or it would be foolhardy to ignore the fact that we've made tremendous advances in how we can study, single neurons, populations of neurons. How we can manipulate them, and channel rhodopsin,

the whole, up to genetic revolution. To not use those tools would be a mistake. And there's a whole other discussion, but I'll just make an argument, I'll throw out this argument that, actually, the history of science is really the interaction between scientific questions and technologies. So, if you go back to almost every major discovery that I can think of in neuroscience, where I know the history pretty well, it was enabled by a technological advance that preceded it.

Tony Zador:
So, Hodgkin, my favorite paper in all of neuroscience, the Hodgkin Huxley papers involved the voltage clamp, the development of the voltage clamp, which was developed by Casey Cole, and taught either Hodgkin or Huxley, a few years after he developed it. And that, in turn, required the development of tubes that were capable of enabling a feedback amplifier, right? So, you can trace the direct line from something from electrical engineering and electronics, to the Hodgkin Huxley equations. You can do a similar thing with single channel recording, low noise amplifiers. That were sufficiently low noise that you could resolve single channel, and you can go on and on, to photon, et cetera, et cetera.

Tony Zador:
So, all these advances that enabled us to study things we couldn't previously study, came from technology. But, that's a whole other-

Paul M.:
Yeah, it is. I'll just have to have you back on the show. So, in my experience, and you've talked a little bit about this, and I fought against this. My experience in science, in academia, is that, people often come in with the big questions, what is consciousness and things? But then, knowledge ruins it. We learn that we don't even understand the processes that are underlying the processes that we think are associated with the big questions.

Tony Zador:
That's right.

Paul M.:
So, the idea that you develop in the paper, it's not exactly reductions. So, my point is, people in their careers tend to end up studying lower level processes, and then, lower level. And that tends to be the trajectory. Yeah. And you're interested in it, too, because you think, oh, it undergirds the higher level process. But, then, 20 years later, you're studying electrons or whatever.

Tony Zador:
That's right, yeah.

Paul M.:
So, the idea that you espouse in the paper with the DNA, it's not exactly reductionist, because it gives DNA, essentially, a higher order role in intelligence, right? But, one could go off the deep end here and say that, well, DNA is made of nucleotides, and it's under the rule of, principles of entropy and metabolism and so on. And so, we could look at that level and say, actually, you have to consider not just the whole of life before us, which evolution encompasses, but the whole of existence as the innate intelligence that we start off with, maybe from the Big Bang. Is that just too crazy of an idea? Why is DNA the right level to zero in on here?

Tony Zador:
I'm not saying that... Okay. So, at one level, what you're saying is right, and you can sit there stoned in your dorm room, wondering.

Paul M.:
Is it okay, if I light up real quick while we talk?

Tony Zador:
Yeah, exactly. You can sit there stoned, in your dorm room, wondering about the deep questions. And that can paralyze you, or, at least, entertain you for as long as you like. But, operationally, there is, to my mind, a reason for stopping at the level of DNA, of recognizing that the genome actually conceptually, plays an important role in this whole process by acting as a bottleneck to the transfer of information about the circuit from one generation to the next. To my mind, that's all we need to know at the abstract level, to move forward on this, as a theoretical question.

Tony Zador:
Now, as a practical question, there are developmental neuroscientists who study the process of going from DNA to a wiring diagram. And I happen not to be one of them, I think it's an interesting field. But, right yet. But, sure, that is, in the actual experimental science, that's what one would study. One would choose to study, how you go from DNA to a wiring diagram, that's developmental neuroscience. And there are certainly some insights from that, that might guide one's thinking, at the theoretical level, for how to build a network through, pass through some kind of a bottleneck. But, the way I look at this, is, you look to the biology for inspiration, and then you take your best shot, right? And that's actually what the field of neural networks did, right?

Tony Zador:
If you look at what happened, as you made the transition from symbolic AI to neural networks, to the first generation of neural networks, symbolic AI, basically involved people thinking, introspecting and thinking about psychology, without thinking about all those nasty details that are inside that black box. I don't blame them. If I'd come on the scene in the 1950s, there's no way I would have been suckered into working on neurons, because that was clearly going to be a long haul, right? So, if there was this promise that you could just, you've got a computer, I've got a brain, you've got a brain, I can look at your brain. You know what? I can just figure it out and just write down a program. Sure, right? Why not? That's clearly, that was an obvious thing to try.

Tony Zador:
And, at least, I suspect that, had I been around back then, I would have kept trying it, and retired after not having figured out how the brain work. But, by the time I got to the scene, it seemed pretty clear to me that the symbolic approach wasn't going to work. And so, there was this generation of one simple level of abstraction, of how the brain works, that was built into neural networks, right? A simple processing unit, that takes inputs from a bunch of other neurons, and produces an output. And I said that, that's, given what people knew about how neurons work back then, it wasn't bad, was an abstract abstraction. And wasn't bad for an abstraction, but it missed some stuff.

Tony Zador:
And now, 50 years later, we're talking about, well, out of all the many, many things that missed, which of those were important? Or, were any of them important? Was that enough? And so, my own belief is that you should keep going back to the neuroscience, and ask, "Okay, given what we know today, is there anything that we can use to guide the next generation?" And there are some clear successes there. So, convolutional neural networks were explicitly inspired by looking at Hubel and Wiesel type receptive fields. Yann Lecun was inspired by that. And he thought about it, and then, abstracted it really nicely. He didn't slavishly build centers around receptive, or edge detecting receptive fields. But, he was inspired by that. And if you look at the work and reinforcement learning, again, people were inspired by the biology of reinforcement learning.

Tony Zador:
That's been a great example of a field, the people who study reinforcement learning in neuroscience are talking to the people who are exploiting reinforcement type learning ideas to build better algorithms. And those people are talking to each other all the time. In many cases, it's some of the same people.

Paul M.:
Well, that's one of the examples where the biology and the artificial neural AI, really map onto each other well.

Tony Zador:
That's right, yeah. Yeah, that's been a huge success.

Paul M.:
Any reactions that you've had to the manuscript that stand out, that have been either really positive or have kept you up at night?

Tony Zador:
Well, I would say that I circulated that manuscript around to a bunch of friends and colleagues. And so, the early drafts of the manuscript, elicited some really strong, mostly negative responses. But, from friends, who-

Paul M.:

That's the best kind, yeah.

Tony Zador:
Yeah, there was fantastic, from a bunch of colleagues, mostly from the machine learning community, who thought I was saying things that I wasn't saying or didn't mean to say, and who also didn't necessarily share my intuitions. And a lot of what this paper was, was laying out the basis for where my intuitions come from.

Paul M.:
I see.

Tony Zador:
Examples of innate learning in animals, for example. So, I would say that a lot of the feedback that I might have gotten after I released the paper, were encountered earlier, anticipated, and then, caused major rewrites. So, I've rarely rewritten a paper as extensively as this one in light of feedback, just because, I was really trying to reach out to a community, and it was useful to figure out why they rejected what I was saying. So, I don't think I convinced anyone. But, at least, I think they understand that what I'm saying.

Paul M.:
Well, that's called progress.

Tony Zador:
Yeah. So, at least, we're using words, more or less, the same way. Or, at least, I'm explaining how I would like people to use words. And explaining the differences between, how I would say a neuroscientist uses some words, and a person in machine learning uses those same words.

Paul M.:
Well, I hope the result, when it comes out in a journal, isn't just, it doesn't lead to a battle of semantics or something. Because, the ideas in the paper are awesome, it's just really fun to read. And, of course, I'll link to it in the show notes, the biorxiv version, anyway. So, thanks for writing it, and thanks for making me think more deeply about our genetic code, as that sort of compressed information. So, I'll stop using words, because I don't want to use the wrong words now.

Tony Zador:
Words mean whatever people want them to mean. It's just that, if you're a large community of people using words one way, and a different community using that same word a different way, it's just a setup for misunderstanding.

Paul M.:
Yeah. Well, I know that you've got to go pretty soon. And so, we didn't really get into your fascinating work on DNA barcoding. I will talk about it a little bit in the introduction and point people to it. And then, maybe we'll have you on again at a later time to talk about it, because it's really cool work that you're doing. So, just a few general questions before you go, then. I heard you say in the talk, that, you had, I believe, as a graduate student, who, you suggested that he or she not pursue a certain path. And then, they eventually just didn't listen to you and pursued it, and it turned out to work. So, do you recommend that advisees not listen to their advisors?

Tony Zador:
More or less, yeah. So, that's not exact. For me, the ideal interaction with anyone in my lab, or anyone I work with, is that, we have conversations until we figure out why we disagree about something. And so, there are a couple reasons you can disagree. One is, as simple as, you're using words differently. Of course, it's even possible that, one of us is just plain wrong about a fact. Or, maybe one of us has the logic wrong. But, the most interesting disagreements come from having intuitions that are different. And intuitions arise from this nebulous cloud of some of all the experience you've had, all the papers you've read, all the experiments you've done, et cetera. And your guess as to how an experiment might turn out.

Tony Zador:
And so, in particular, I was talking about my work with a brilliant former student in my lab, Peter Znamenskiy. And when he joined the lab, we agreed on the overall scope of the project, and we agreed that we would have to look at one output pathway of the auditory cortex. And then, since these experiments were really hard, we had to pick which output pathway, we had to commit to a pathway. And I was advocating by analogy with work in the non-human primate, that he look at some cortical,

cortical pathway, let's say, from the auditory cortex, to the posterior parietal cortex. And he kept suggesting that we look at the striatum. And I didn't know anything about the striatum, other than that it was a complicated place and I didn't know anything about it. And he kept saying, "So, what do you think about the striatum?" "Oh, no, don't do that." And, at some point, I think he recognized that I had no logical basis for rejecting it, it was just complete naivete and ignorance. And so, he did exactly what he should have done, which is, after having considered my arguments and rejected them as unfounded, did exactly what he thought was best. And he came back six months later with fantastic results that subsequently transformed my lab's research program.

Tony Zador:
That is exactly the best interaction, to my mind, between two people who work together. He's the one doing the experiment, he's the one who's ultimately responsible for the success of his project. And so, I am somebody who, I think, could provide, at worst, the sounding board, and perhaps, even some insight. But, in the end, he's got to take responsibility for it. So, he does what he thinks is best. And this is a conversation I have with everybody who enters my lab, that ultimately, they are responsible for making their project work. And I'll do my best to help them, advise them, discuss with them. But, in the end, it's their decision what they do.

Paul M.:
You have a lot of expertise in a lot of different areas. If you could wave a magic wand right now, and just, poof, be an expert in something, what would it be?

Tony Zador:
Oh, wow. That's a great question. I'm not sure I even would. For me, the fun is learning something that I'm not an expert in. So, if I could just wave that wand, then I would lose interest in it.

Paul M.:
That's a really good answer.

Tony Zador:
Right now, I'm trying to learn developmental neuroscience. And it's super exciting, because I don't know anything about it. But, it turns out, it's really interesting and important.

Paul M.:
When you figure it out, sum it up and just tell me. I don't need to learn it.

Tony Zador:
Well, yeah. But, that's what you think, that's what you think. And that's what I thought, too. And the fact is, as you said, that anything you dive into, almost anything you dive into, that smart people have studied, is worth studying. Otherwise, they wouldn't have studied it.

Paul M.:
What is something that you used to believe, that you consider naive now?

Tony Zador:
Oh, that's easy enough. I mean, what I did my PhD thesis on, which was the idea that the key thing missing in artificial neural networks was single neuron complexity. I think, single neuron complexity is really interesting and important to biology, but I don't think it's the fundamental thing that's missing. I think that what I'm studying now is the fundamental thing that's missing, namely, circuits. Of course, I'll probably look back in five years and think that was naive. But, right now, I'm convinced that that's true.

Paul M.:
Tony, thank you so much for your time. I know that you have some fun to get to, with your kids in the city. So, I appreciate your time here and continue the good work, man.

Tony Zador:
Thanks a lot. This was really fun.