



Her Key to Modeling Brains: Ignore the Right Details

Being able to think like a physicist helps Carina Curto, a mathematician-turned-neuroscientist, pull insights about the human brain out of theoretical models.

By Siobhan Roberts



[Hatnim Lee](#) for Quanta Magazine

Carina Curto's love of mathematics first steered her toward theoretical physics. But today, this associate professor of mathematics at Pennsylvania State University applies her expertise to theoretical and computational neuroscience instead.

In May 2013, the mathematician [Carina Curto](#) attended a workshop in Arlington, Virginia, on "[Physical and Mathematical Principles of Brain Structure and Function](#)" — a brainstorming session about the brain, essentially. The month before, President Obama had issued one of his "Grand Challenges" to the scientific community in announcing the [BRAIN Initiative](#) (Brain Research

through Advancing Innovative Neurotechnologies), aimed at spurring a long-overdue revolution in understanding our three-pound organ upstairs. In advance of the workshop, the hundred or so attendees each contributed to a white paper addressing the question of what they felt was [the most significant obstacle to progress in brain science](#). Answers ran the gamut — some probed more generally, citing the brain’s “utter complexity,” while others delved into details about the experimental technology.

Curto, an associate professor at Pennsylvania State University, took a different approach in her entry, offering an overview of the mathematical and theoretical technology:

A major obstacle impeding progress in brain science is the lack of beautiful models. Let me explain. ... Many will agree that the existing (and impending) deluge of data in neuroscience needs to be accompanied by advances in computational and theoretical approaches — for how else are we to “make sense” of these data? What such advances should look like, however, is very much up to debate. ... How much detail should we be including in our models? ... How well can we defend the biological realism of our theories?

The trend in Europe — for instance, with the [Blue Brain Project](#) and the [Human Brain Project](#) — is to run complicated models on supercomputers, with the goal of creating a digital reconstruction and simulation of the human brain. Curto sees the value in such enterprises, but she argues that these models are often ugly, cumbersome, difficult to understand, and problematic to build upon:

Beautiful models are simple to describe and can be mathematically analyzed. They contain important kernels of “truth” in them, even if they are not messy enough to be completely accurate. The general principles behind these models are often easy to understand, making them attractive candidates to be tested experimentally. Equally important is the fact that beautiful models can synchronize a theoretical community around a coherent line of research, leading to work that can be understood and built upon by future generations.

Currently, Curto is working with a beautiful model called the combinatorial threshold-linear networks model, which she is developing in collaboration with Katherine Morrison of the University of Northern Colorado, with funding from a BRAIN Initiative grant. Incorporating combinatorics, linear algebra, differential equations and dynamical systems, the model facilitates her study of how the precise connectivity within the brain shapes the emergent dynamics of neural activity.

That connectivity-dynamics interplay is still very poorly understood, in part because of a dearth of data and technology until fairly recently. On that front, the BRAIN Initiative has supported two corrective measures. One is a push to scale up the technology for recording thousands of neurons firing simultaneously (rather than, say, recording only a couple of hundred, which is where the standard had stagnated). The other push is to advance “connectomics,” the study of connections within the brain, by identifying, characterizing and mapping the links between neurons.

“Mathematically, the interplay between connectivity and dynamics is a really hard problem,” Curto commented. Even if scientists have full knowledge of and control over the circuitry, it’s still difficult to predict the dynamics. “And because of that, the theory has been kind of stuck. It’s hard, theoretically, to make this connection between structure and dynamics.”

Yet the only laboratory Curto really needs is her blackboard (and occasionally just pencil and paper) — though sometimes she uses computer simulations to check the model and generate ideas. “It’s really fun and pretty work. But it also seems genuinely relevant to big, fundamental neuroscience questions,” she said. “I think I’ve found the model that’s in the sweet spot.”

Quanta spoke with Curto in May, just before and after she presented her latest work at a [conference on “connectomes”](#) (maps of neural connections) at the Janelia Research Campus at the Howard Hughes Medical Institute in Virginia. The interview has been condensed and edited for clarity.

You’re a mathematical neuroscientist, but you set out to be a theoretical physicist, and your Ph.D. at Duke University explored algebraic geometry questions in string theory. What happened?

So yeah, “What happened?” A lot of people ask me this. I loved the math that I learned in grad school, but I became a bit disillusioned with string theory, just from a scientific perspective. It’s still a very cool area of physics for inspiring interesting math and making connections to all kinds of mathematics that aren’t typically considered applied. But from the perspective of someone who always wanted to be at the interface of science and mathematics, it became less and less satisfying.

Halfway through grad school I started looking for something else. I sat in on an economics course, which was totally boring, because they spent all their time reviewing vector calculus, and I was not interested in learning that for the sixth time in my life. And then I discovered neuroscience kind of randomly. I sat in on a course and went to some talks, including one by Larry Abbott, who is a very famous theoretical neuroscientist at Columbia University and a former physicist. He talked about the visual system, but he had this way of approaching the neuroscience questions as a theoretical physicist would. And I said to myself, “Oh, this feels like what I thought theoretical physics would feel like.”

Approaching problems as a theoretical physicist would — what does that entail?

There is a certain way of looking at problems whereby you try to simplify them in ways that biologists often don’t like. Biologists are often committed to the details of what they study. They think those details are important, and for many questions they are important. But depending on the question, some details may not be. A theoretical physicist simplifies a problem in a way that makes it amenable to a mathematical analysis, while still preserving the essence of what is complicated, the essence of the phenomenon one is trying to understand. Of course, the simplification has to depend on the question you’re asking.

To me it was eye-opening that this kind of training or way of thinking could be applied to neuroscience. I had no idea at that point in my life that neuroscience — or any field of biology, for that matter — could be tackled productively using the approaches or thinking styles that a physicist or a mathematician might have. I had never really studied biology. Since high school I had always avoided it, because I had a terrible biology course in ninth grade that was just memorizing parts of things and I’m terrible at memorizing.

That was one of the things that drove me to physics, because physicists are very proud of being able to derive things from first principles, and everything logically fits together. A lot of people have a wrong impression of math because of a bad high school class or a bad experience with math early on. That happened to me with biology.

As a grad student, I finally sat in on a neuroscience course, and I started learning more and more. And then I made the decision in my fourth year of grad school that when I finished my Ph.D., I was going to switch to neuroscience. I wasn't planning to abandon my math and physics background, but I decided to apply it to neuroscience rather than string theory.

Toward the end of my fourth year, there was a professor, Ken Harris, then at Rutgers University, who was spamming math departments across the country. He sent an email to everyone in the Duke University math department about how he was trying to recruit people with math and physics backgrounds to his neuroscience lab. I went to meet Harris and he offered me a job as a postdoc, essentially on the spot. After getting my Ph.D., I spent three years in his lab, doing mostly data analysis and traditional computational neuroscience — learning a completely new set of tools and a lot of neuroscience.

Do you have any regrets?

No, not really. Early on when I made this transition, the only thing that I had regrets about was that I didn't know how much math I would get to use. I plunged straight in. I was like, "I'm going to learn neuroscience, I'm not going to worry about applying specific mathematical tools, I'm just going to be open-minded and immerse myself." It was almost like an exchange program, when you go to another country and live with a family and immerse yourself in the new culture and language. That's what I did with neuroscience. I just went full in and didn't try to force my own background onto it.

There were times when I thought, "OK, I'm never going to do real math again. I'm never going to prove another theorem, because all I'm doing is data analysis and computational work, and it's interesting, but it's not really math." And so I did have that feeling at times, with some regret.

But then it was kind of amazing. After a few years, I started having my own ideas. I started asking questions that were meaningful to neuroscientists, and having ideas about how I could tackle some neuroscience problems that were quite mathematical — bringing in tools even from areas of math like topology and commutative algebra and combinatorics that were not traditionally being used. That was very exciting. I got to use much more sophisticated math but still address questions that were meaningful to the neuroscience.

In what direction is your research heading currently?

One research direction I've gone into, and the reason I was invited to the Janelia conference, is related to the larger project of trying to understand the relationship between a network's connectivity structure and its dynamics — the dynamics being the patterns of neural activity we observe in neural recordings.

For instance, you may have some set of neurons that fire in a regular sequence. You could be recording, say, 10 neurons, and you could see over and over again the sequence of: 3, 5, 6, 7, 9 ... 3, 5, 6, 7, 9 ... and it repeats. The question would be: Why are you seeing this pattern of activity? What does it tell you about the underlying connections between the neurons, and what do the connections tell you about the activity pattern? How does the structure of the network affect the more ephemeral dynamics? How are these patterns of activity being used to encode and transmit information in the brain?

And you developed your "network songs" to help explain this interplay?

Oh, you found the network songs! That was kind of a gimmick; I used the songs in talks a few years

ago and people loved them. The idea was to show the rhythmic activity that these networks generate. Sometimes it's hard to see that just by looking at plots.

So I set the activity to music: You have individual neurons and their firing rates. Their activity levels go up and down, and as a collective population, they generate rhythms. I did something very simple: I assigned a piano key note to each neuron, and then I used the neuron's firing rate to modulate the amplitude. When it's high-firing, that note is playing loudly, and when it's low-firing, it's not playing at all. It's another way of representing the very same thing that's on the plot, just a solution to a differential equation. But you can hear it, and somehow the ear picks up recurring patterns better than the eye. It's really just a way of turning the same information into sound so that the rhythmicity can be more apparent. It makes the point nicely that these networks are very rhythmic.

Ultimately, those dynamics are what we believe is leading to behavior, and leading to perceptual experiences, memory recall and so on. And the work that I'm most excited about right now is on a very nice model of a neural network, the combinatorial threshold-linear networks model. My collaborator Katherine Morrison and I have been playing around with the model and proving theorems as well — we can actually prove theorems that allow us to tie the structure of the graph to important features of the dynamics. That's been very exciting.

What does it mean to be proving theorems about the brain, so to speak, in this context?

Of course, you still have to do experiments. Proving something about your abstract model is not proof that you are right about the actual science. But it gives you insights into things that can then be checked in a more traditional, scientific-method-type way.

In a sense, a mathematical model is a simplified model in the same way that a worm, like *Caenorhabditis elegans*, is a model organism: There is a lot of Nobel Prize work in neuroscience that was done on worms, very small creatures with very simple nervous systems, because these simplified models really do give insight into the human brain. In a similar way, we can prove theorems in simple mathematical models and gain real insight into how the model is working. That insight then gives us things to look for in the real system, the real brain.



[Hatnim Lee](#) for Quanta Magazine

“Having good models, beautiful models, is a really nice way of generating hypotheses in a controlled and rigorous fashion,” Curto said.

When you can prove a theorem, it basically captures infinitely many examples that you could never individually check on your own — even with the most powerful computers. And often, the essence of what is needed for a proof to work gives insight into the key features that make the phenomenon tick. You could discover something in simulations by playing around with the parameters, and then make the guess, “I think this feature is what’s causing the phenomenon to happen.” But if you can really prove it, even if the proof requires a simpler setting, then you can nail down the key features that lead the phenomenon to work. Being able to do real math in relation to a model like this helps a lot in gaining insight into the mechanisms driving the phenomenon.

For example, there are certain periodic patterns of activity that we see in real neural circuits, and we always wonder, “OK, what’s causing these rhythms of neural activity to emerge in the network?” And in the context of a simple model one can actually prove, to some extent, which features are causing that pattern to emerge.

This gives you something to look for. If I have my connectomics data, if I have some structural information about the network, I can check: Do I see the same structural pattern in the network that is giving rise to this phenomenon in my model? So the mathematics gives a way of generating hypotheses.

The one bad thing about the traditional scientific method is that it doesn’t really tell you how to generate hypotheses in the first place. When you get to complex systems, there is this combinatorial explosion of possible hypotheses, and you cannot check them all. Having good models, beautiful models, is a really nice way of generating hypotheses in a controlled and rigorous fashion.

Overall, what do you find most fascinating about neuroscience?

The field is evolving very rapidly right now. New types of neurons are discovered all the time, and new coding properties of neurons, too. We’re just now figuring out how networks are structured in the brain.

I remember hearing about the 1960s as this great age for particle physics: A new particle was being discovered every week, and theorists were busily trying to figure out how to organize the particles, and what kind of mathematical structures should be involved, and what kinds of symmetries were present, and so on. There were lots of interesting connections to mathematics. That’s the sort of exciting thing I wanted to be a part of.

When I got into neuroscience, it was kind of like that. Instead of a new particle discovered every week, it was a new type of neuron, or new rules for how neurons are connected, or new experiments showing interesting and surprising results. It felt like — and still feels like — a very exciting playground for a theorist.