A conversation with Professor Eve Marder, May/June 2020

Participants

- Professor Eve Marder - University Professor and Victor and Gwendolyn Beinfield Professor of Neuroscience, Brandeis University
- Joseph Carlsmith - Research Analyst, Open Philanthropy

Note: These notes were compiled by Open Philanthropy and give an overview of the major points made by Prof. Marder.

Summary

Open Philanthropy spoke with Prof. Eve Marder of Brandeis University as part of its investigation of what we can learn from the brain about the computational power ("compute") sufficient to match human-level task performance. The conversation focused on computational models of different biophysical processes in the stomatogastric ganglion.

The stomatogastric ganglion

Some of the work done by Prof. Marder and her collaborators focuses on the stomatogastric ganglion, a circuit consisting of about 30 neurons in the stomach of decapod crustaceans. This circuit contains two central pattern generations (the pyloric and the gastric mill central pattern generators), which create rhythmic patterns of neural firing that control muscle movements.

You can see maintaining these rhythms as the high-level function that the circuit is performing at a given time (transitions between modes of operation are discussed below). Neuroscientists had a wiring diagram for the pyloric rhythm in 1980, and there was a fairly good first-principles idea of how it worked back then. It is not too difficult to model tri-phasic rhythm.

However, we didn't understand the biological implementation back in 1980, and Prof. Marder does not think the project of modeling the pyloric rhythm particularly interesting in itself. Rather, such models are interesting insofar as they help you answer questions that you don't understand, that are relevant to many neural systems. For example: how does the crab maintain phase constancy when every single ion channel has different time constants? What tricks does it use? Why does it have thirty neuromodulators, as opposed to e.g. three?

Transitions between modes of operation

Biology has found a series of mechanisms that allow the system to transition smoothly between different modes of operation. For example, you can walk slowly or quickly. Although eventually you will change gait. Prof. Marder believes that such smooth transitions are centrally important to understanding brains, especially big brains. The mechanisms involved allow brains to avoid having to fine-tune or find singular solutions.

However, most computational models don't capture these transitions. For example, if you want to capture the behavior of an eight channel neuron with a three channel model, you'll hit nasty bifurcations. Indeed, one hypothesis is that neurons have many ion channels with overlapping functions because this facilitates smooth transitions between states.

However, the mechanisms involved in these transitions are not all well-understood biologically, partly because there isn't enough work on the topic. Many people want to average biological complexity away, but Prof. Marder thinks that you can't do this without losing those smooth operational dynamics. You don't want your model to fall down, or stop breathing, as it goes from a walk to a run.

Timescales

"Timescale," in this context, refers to the time between a specified beginning to a process, and a specified effect. For example, a neuromodulator binds to a receptor, activates a G-protein, the G-protein does something, and then 200 milliseconds later, some channels are phosphorylated and so the neuron behaves differently.

Both experimentalists and theorists sometimes act as though there's a mechanistic wall between short-term, middle-term, and long-term changes in neural systems. This is partly because you have to come up with experiments that will occur over a given timeframe (two hours, two days, two weeks). But that doesn't mean the time constants of these processes are two hours, two days, two weeks, etc: it's just that you designed an experimental protocol that allows you to see the difference between these periods of time.

Historically, limitations on computational resources have also played a role in popularizing such separations. In the old days, people were limited by how much they could compute by the timesteps and integrators they were using, so there was tremendous pressure to

separate timescales: no one wants to integrate over very long times at the rates you'd need to in order to capture fast dynamics.

Thus, for example, people will take a model with eight or ten currents, and try to reduce it by separating timescales. If you're clever, you can retain various essential features, but it's hard to know if you've got them all. Whether or not such separations between timescales are biologically reasonable, though, they were computationally necessary, and they have resulted in ingrained beliefs in the field.

In reality, the nervous system has an incredible ability to move seamlessly between timescales ranging from milliseconds to years, and the relevant processes interact. That is, short time-scale processes influence long time-scale processes, and vice versa. And unlike digital computers, the brain integrates over very long timescales at very fast speeds easily and seamlessly.

Neuromodulation

Timescales of neuromodulation

Mr. Carlsmith asked Prof. Marder's opinion of the argument that because neuromodulation works on comparatively slow timescales (e.g., effects that proceed via G-protein receptors involve delays of hundreds of milliseconds), it won't be very computationally intensive to capture, relative to faster processes in neural systems.

This depends on how accurate you want your model to be. Some people will say: here's my starting state, now I'm going to put a modulator on and change the state of my circuit, and then I'm going to compute again only in the fast domain, not worrying about whether there are smooth transitions in between.

However, Prof. Marder is confident that's not how the brain does it. In reality, there are many loops and state changes, involving processes that go from fast, millisecond timescales, to hundred millisecond timescales, to minutes, to hours, and slower processes are changing the availability of what's present in the fastest timescale.

You could argue that all we need to do is understand what these slow processes are doing to the state variables of the fast processes, and then we can just compute in the fast timescale domain. At some level that works, but only if the network is ever stationary (you need to pick the window over which you'll treat it as stationary), and it's never really stationary in the slower domain.

Spatial resolution of neuromodulation

Mr. Carlsmith asked Prof. Marder's opinion of the idea that because neuromodulation is a global, broadcast message (as opposed to, e.g., a differentiated signal directed at a specific neuron), you can model it fairly efficiently: e.g., using a single variable, or, if you have many neuromodulators, a vector.

Prof. Marder noted that you also need to capture the role of different neuromodulator concentrations. The more neuromodulator you release, the further it will travel, so the spatial scale is not constant.

Epistemic barriers to understanding circuits

It's been hard to make progress in understanding neural circuits, because in order to know what details matter, you have to know what the circuit is doing, and in most parts of the brain, we don't know this.

This is why Prof. Marder continues to work on a small invertebrate motor system: she knows exactly what it's doing, together with when it's working or not working, and she can measure this very well (being able to identify the neurons involved in a circuit is also crucial). The retina affords similar advantages: we basically know what the retina does (though there's still more complexity to understand), and that's why it's been so useful. There are also some circuits in leeches, *C. elegans*, flies, and electric fish that are relatively well-characterized.

It's not that you can't make simplifying assumptions. It's that absent knowledge of what a piece of nervous system needs to be able to do, you have no way of assessing whether you've lost something fundamental or not. This is a serious barrier to progress in the field, despite the availability of many impressive new technical capabilities. And scientists aren't always good at maintaining awareness of what they do and don't understand about a biological system.

Computational requirements

Prof. Marder and her collaborators have used single-compartment conductance models to replicate the rhythms in the stomatogastric ganglion. However, the computational power necessary to run e.g. a full Hodgkin-Huxley model depends a lot on implementation: e.g.,

what platform you use, what language you're using, what method of integration, and what time-step for integration (all of your compute time goes to integrations).

Prof. Marder does not think about biology in terms of the "computational power" of e.g. a neuron in the stomatogastric ganglion. And in general, the level of modeling detail you need depends on the question you're interested in. For example, in the context of the retina, it will be substantially easier to create a first-pass model that captures a lot of what it does, than to create a model that replicates the full range of the human retina's functionality.

All Open Philanthropy conversations are available at
http://www.openphilanthropy.org/research/conversations