A conversation with Professor Chris Eliasmith, September 25, 2019

Participants

- Prof. Chris Eliasmith -- Professor of Philosophy and Systems Design Engineering, Canada Research Chair in Theoretical Neuroscience, and Director of the Centre for Theoretical Neuroscience at the University of Waterloo
- Joseph Carlsmith - Research Analyst, Open Philanthropy

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

Summary

Open Philanthropy spoke with Prof. Chris Eliasmith of the University of Waterloo 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 Prof. Eliasmith's experience building large-scale brain simulations, and on the compute required to model different processes in the brain.

Functional brain simulations

Prof. Eliasmith has built large-scale brain simulations that aim to reproduce different high-level aspects of brain function, while capturing certain kinds of neurophysiological data as well. One example is SPAUN (Semantic Pointer Architecture Unified Network), a simulation consisting of six and half million model neurons, which can perform various perception, pattern recognition, and motor control tasks.

Unlike SPAUN, standard machine learning models do not attempt to replicate physiological features of the brain like connectivity or number of layers. For example, some machine vision systems have hundreds of layers of depth, which doesn't seem likely to be what's going on inside the brain.

Complexity of neuron models

It is unclear what level of biological detail you need per cell in the brain in order to replicate biological function, and the answer depends in part on the function in question (e.g., capturing brain maintenance and growth is different from capturing performance on a cognitive task to within a given standard of error).

Prof. Eliasmith's general approach is to see what simple models are able to do, and to introduce additional complexity only when doing so becomes necessary. In his models, he has thus far been able to successfully replicate various types of high-level behavior, along with various types of neuro-physiological data, without recourse to highly complex neuron models -- a result that he thinks substantially less likely in worlds where the brain's performance on these tasks proceeds via biophysical mechanisms his models do not include.

However, this doesn't mean that we won't discover contexts in which greater complexity is necessary. And we are very far away from being able to test what is required to capture high-level behavior on the scale of the full human brain.

Replicating spike trains

There is no "magical answer" to the question of how accurate a model of neuron spiking needs to be. In experiments fitting neuron models to spike timing data, neuroscientists pick a metric, optimize their model according to that metric, and then evaluate the model according to that metric as well, leaving ongoing uncertainty about the importance of the aspects of neural activity that the relevant metric doesn't capture.

It is also an open question how much compute is required to capture the dynamics within neurons that lead to a decision about whether or not to fire. One possible source of evidence comes from neuron modeling contests, in which simple models seem to have performed better than complex models, because simple models have fewer parameters and hence avoid overfitting. Exponential integrate-and-fire neurons (a type of neuron model) seem to have done an unusually good job of capturing different types of biophysical data (specifically predicting spike train patterns given a input current).

Prof. Eliasmith also noted that with a linear filter and a non-linearity, you can get a surprisingly close match to what a neuron six (or more) layers into the cortex is doing, though such a model only gives first-order predictions that leave out subtleties in the data.

Spikes through synapses

In the types of models Prof. Eliasmith builds, a spike through a synapse would typically be modeled as follows: when a spike comes in, it is weighted by the synaptic weight, filtered by a synaptic filter, and then summed into the post-dendritic input that is going into the cell body.

You also need to update the cell body voltage itself, but because there are so many more connections than neurons, the compute for modeling synaptic input is hugely dominant (e.g., by three or four orders of magnitude).

Temporal precision

Prof. Eliasmith typically uses 1 ms time-steps in the simulations he builds.

Stochasticity

Pretty much everything Prof. Eliasmith does with his models works fine in a stochastic regime, but stochastic approaches require more synapses, so he does not bother with them. This decision is driven primarily by the availability of deterministic large-scale computational platforms. If there were cheap stochastic computers available, Prof. Eliasmith would probably use stochastic approaches.

Dendritic computation

Prof. Eliasmith believes that neurons probably have non-linearities in their dendrites. In attempting to construct models of attention, for example, he has found that he needs more model neurons than seem biologically realistic, and the neuron count would go way down if he had certain kinds of non-linearities in the dendrites.

Including these non-linearities would not drastically increase compute burdens (it might be equivalent to a 2X increase). A simple version would basically involve treating a single neuron as a two-layer neural network, in which dendrites collect inputs and then perform a non-linearity before passing the output to the soma. Prof. Eliasmith is sympathetic to the idea that the tree-structure of dendrites limits the additional complexity that dendritic computation could implicate in the context of such multi-layer networks (e.g., the tree-structure limits the outgoing connections of a dendritic sub-unit, and additional non-linearities in the neuron do not themselves add much compute in a regime where spikes through synapses are already the dominant compute burden). That said, there are many mechanisms in neurons that could in principle make everything more complicated.

There are also arguments that certain forms of active dendritic computation function to "linearize" the inputs -- e.g., to combat the attenuation of an input signal as it travels through the dendritic tree, such that the overall result looks more like direct injection into the soma.

Learning

We don't know which aspects of the brain's biophysical dynamics are important for learning and adaptation. Some neuroscientists suggest that the function of various molecular mechanisms is to embed long-term memories into a single synapse.

In the large scale brain simulations that Chris Eliasmith builds, he often uses an error-driven Hebbian rule, which computes updates to synaptic weights based on pre-synaptic activity, post-synaptic activity, and an error signal (which, in the brain, could proceed via a mechanism like dopamine modulation). This rule requires on the order of three to five operations per synapse (a couple of products, and then a weight update), though the total burden depends on how often you perform the updates.

Alternative signaling mechanisms

Prof. Eliasmith does not include glia, gap junctions, or non-standard forms of axon signaling in his models, because his approach is to see how far simple, textbook models can go before additional complexity is needed.

Adding gap junctions probably would not substantially increase the overall compute budget, because they are not very common. You can model a gap junction as a connection that updates every timestep, rather than every time a spike occurs.

Whether there is a "right" level of biophysical detail

There is no privileged model of the brain which can claim to be *the* model of how the brain performs tasks. You can't answer someone's question about how the brain works without knowing exactly what the question is. Nor is there a privileged level of biological detail that a model needs to include in order count as a brain model, as all models are wrong to some extent.

You can, though, specify a particular set of functions that a model needs to reproduce, with a particular degree of similarity to human behavior and anatomical and physiological data. Prof. Eliasmith's work is basically oriented towards building a brain model that satisfies constraints of this type.

Prof. Eliasmith thinks that neuron models at roughly the level of detail he uses in SPAUN (possibly including some non-linearities in the dendrites), if scaled up to the size of the brain as a whole, would be able not just to replicate cognitive performance, but also to reflect a functional profile similar to biological neurons. Models like SPAUN already reproduce various types of neurophysiological data, such as firing statistics, frequency profiles, and percentages of neurons responsive to different types of input stimuli. You can't do 1-1 neuron mapping, but you can't do that between individual humans either.

Other people to talk to

- Randall O'Reilly - Professor of Psychology, Computer Science, and the Center for Neuroscience at UC Davis.

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