A conversation with Professor Blake Richards, September 20, 2019

Participants

- Professor Blake Richards - Assistant Professor in the Montreal Neurological Institute and the School of Computer Science at McGill University
- Joseph Carlsmith - Research Analyst, Open Philanthropy Project

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

Summary

Open Philanthropy spoke with Prof. Blake Richards of McGill 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 the compute required to model learning in the brain.

Learning in the brain

Prof. Richards thinks it's reasonable to distinguish between the question of how much compute is required to replicate the behavior of a static snapshot of the brain, and the question of what it takes to get that system to learn in the way the brain does.

However, it is very difficult to say at this point exactly how much compute would be required to model learning in the brain, because there is a lot of disagreement in the field as to how sophisticated the learning algorithms in the brain are. This is partly because we don't have a good hold on how much human learning is truly general purpose, vs. constrained to particular tasks.

Types of learning algorithms

We can distinguish between at least three different types of learning algorithms, which vary in the scaling properties of the compute resources they require.

First-order gradient descent methods, like back-propagation, use the slope of the loss function to minimize the loss. Here, learning is basically a backwards pass through the network, so the compute required scales linearly with the number of neurons and synapses in the network, adding only a small constant factor.

More sophisticated learning algorithms, such as second-order gradient methods, take into account not just the slope of the loss function gradient but also its curvature. These require more compute (the compute per learning step scales as a polynomial with the number of neurons and synapses), which is why people don't use these techniques, even though they are arguably much better.

In the other direction, there are algorithms known as "weight-perturbation" or "node-perturbation" algorithms. These involve keeping/consolidating random changes to the network that result in reward, and getting rid of changes that result in punishment (a process akin to updating parameters based on simple signals of "hotter" and "colder"). These algorithms require less compute than first-order gradient descent methods, but they take longer to converge as the size of the network grows. In this sense, they involve trade-offs between compute and time.

Prof. Richards favors the hypothesis that the brain uses a learning method with compute scaling properties similar to backpropagation. This is partly because humans are capable of learning so many tasks that were not present in the evolutionary environment (and hence are unlikely to be hardwired into our brains), with comparatively little data (e.g., less than a weight-perturbation algorithm would require).

Biophysical complexity

Some neuroscientists are interested in the possibility that a lot of computation is occurring via molecular processes in the brain. For example, very complex interactions could be occurring in a structure known as the post-synaptic density, which involves molecular machinery that could in principle implicate many orders of magnitude of additional compute per synapse. We don't yet know what this molecular machinery is doing, because we aren't yet able to track the states of the synapses and molecules with adequate precision.

There is evidence that perturbing the molecular processes within the synapse alters the dynamics of synaptic plasticity, but this doesn't necessarily provide much evidence about whether these processes are playing a computational role. For example, their primary role might just be to maintain and control a single synaptic weight, which is itself a substantive task for a biological system.

Overall best guess

Based on Prof. Richard's best guess, it seems reasonable to him to budget an order of magnitude of compute for learning, on top of a budget of roughly one FLOP (possibly a bit more) per spike through synapse. However, it could also be higher or lower.

Other people to talk to

- Prof. Anthony Zador - Cold Spring Harbor Laboratory
- Prof. Cian O'Donnell - University of Bristol
- Dr. Timothy Lillicrap - Google DeepMind
- Dr. Adam Santoro - Google DeepMind

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