Brain building 101: advances in large-scale neural simulation

Chris Eliasmith

The construction of the world's largest functional brain model and large-scale, real-time hardware simulations rely on the same mathematical and neuromorphic methods.

Physicist Richard Feynman wrote: "That which I cannot create I do not understand." When it comes to the brain, we are only at the very beginning of attempting to understand how it underwrites the sophisticated, subtle, and flexible behavior we see in animals. Nevertheless, there are ongoing, well-funded attempts to create one. Perhaps most famously, the Blue Brain project has built simulations of millions of neurons.[1] IBM's Cognitive Computation project, too, has recently simulated one billion neurons.[2]

Two major challenges must be overcome for this work on creating brains to actually help us understand them. The first is that current simulations are far too slow. The Cognitive Computation project, which is much faster than the Blue Brain project, runs simulations about 1/400[th] of real time. The second, perhaps more daunting, problem is that these simulations do not exhibit any behavior. The simulated neurons generate action potentials or 'spikes'—the main means of communicating in brains—and influence one another. But they do not cause the simulated brain to recognize something, or learn something new, or control movement. In short, they do not control behavior, which is the main purpose of brains. Our recent work has focused on solving this second problem, as well as working closely with other groups to solve the first.

To build large-scale, behaving systems, we have taken a two-stage approach. First, in our earlier work, we developed a set of mathematical methods that allow us to take any nonlinear dynamical system and build a spiking neural network that can approximate it well.[3] We call this set of methods the Neural Engineering Framework (NEF), and we and others have demonstrated its successful application to navigation,[4] working memory,[5] human syntactic induction,[6] reinforcement learning,[7] and many other functions. Second, in our most recent work, we have identified a method of representation, a set of functions,
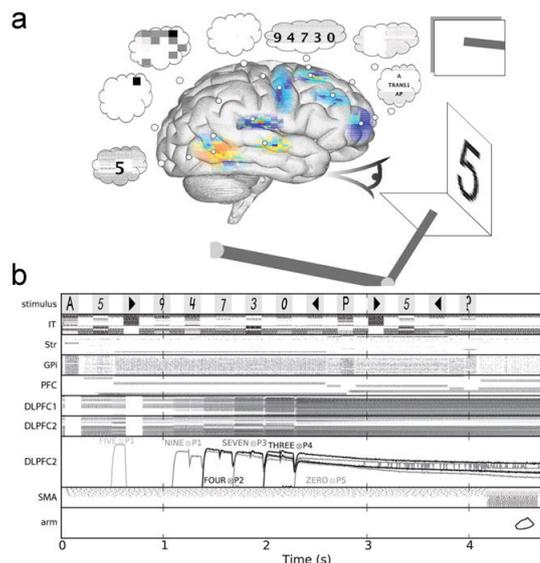
Figure 1. Semantic Pointer Architecture Unified Network performing a question-answering task. (a) A screen capture from the video of question answering (videos of all tasks are available).[3] This task consists of the model being shown a series of digits, and then being asked one of two questions about those digits. In this case it is asked a 'P' question, which means it must indicate what digit is at the queried position. Here, the queried position is '5,' so it should respond with the digit at the fifth position. It can also be asked 'K' questions, for which it must indicate where the digit of the indicated kind is located in the list. Several thought bubbles show spike trains in the background and decodings of those spike trains as digits. Colors on the brain surface indicate neural activity (low-pass-filtered spikes) in the relevant areas. Only some areas are shown. (b) The time course of the entire run, including neural spikes generated in various brain areas during the task. The input is shown in the top row, and the drawn response is shown in the bottom row. The solid lines for DLPFC2 indicate a decoding of the spike trains on the previous row, to aid visualization. IT: Inferotemporal cortex. Str: Striatum. Gpi: Globus pallidus internus. PFC: Prefrontal cortex. DLPFC: Dorsolateral prefrontal cortex. SMA: Supplemental motor area.

and a functional organization that we believe are sufficient for capturing many aspects of biologically based cognition.[8] We refer to this hypothesis as the Semantic Pointer Architecture (SPA).

At the same time as we have developed the SPA, we have been working closely with two neuromorphic hardware groups to address the first problem: i.e., how to run large brain simulations in real time. The first group is lead by Kwabena Boahen at Stanford University, CA, who has developed the Neurogrid chip, which is able to simulate millions of spiking neurons in real time with extremely low power. The power savings comes from the fact that each neuron is simulated by a small set of circuit elements directly on the chip. We have now implemented the principles of the NEF on this chip.[9] At the Neuromorphic Engineering workshop in Telluride, CO, this summer, we used this chip to control a simple robot arm.

The second hardware group is lead by Steven Furber at Manchester University, UK, who has developed the SpiNNaker system. This system takes the very different approach of digitally simulating neurons and building very large computer systems that will be able to simulate up to one billion neurons in real time. We have also implemented the NEF on this system,[10] and were able to realize the world's first untethered robot fully controlled by a spiking network at the Telluride workshop. This robot was able to navigate to a target and find its way back to the starting point, using neural responses reminiscent of 'place cells' identified in rodents performing the same task.

We have also developed a sophisticated but easy-to-use software environment called Nengo that implements the NEF.[11] With this software, we are able to build and test models before implementing them on hardware. In addition, we are able to test the SPA by building very large scale models. We have recently built the world's largest functional brain model, which we call the Semantic Pointer Architecture Unified Network (Spaun): see Figure 1. This model has one eye, one arm, 2.5 million spiking neurons, and realizes a wide array of biological functions, including perception, memory, learning, and motor control.[12] Spaun's only input is visual images (largely handwritten and computer-generated digits), and its output is the movement of a physically simulated arm (used to draw one or more digits). To demonstrate Spaun, we have it perform eight different tasks, in any order, with no modeler intervention. These tasks include object recognition, copy drawing, reinforcement learning, list memorization, and so on. Videos of the model performing any of these tasks can be found elsewhere.[13]

While we have shown that models, such as Spaun, are able to capture a wide variety of psychological, neuroanatomical, and neurophysiological data across species, we still have a long way to go to understand the brain. Obviously, many elements of the brain are missing from Spaun, and some of the details stand to be improved. One practical limitation is that Spaun takes about 2.5 hours to simulate one second of model time. We believe that our recent successes implementing the basic theory behind Spaun on Neurogrid and SpiNNaker bode well for large-scale, real-time neural simulations in the near future. Real-time simulation opens up a host of more interesting behaviors: anything that depends on rapid, dynamic interaction with the environment.

Author Information

Chris Eliasmith
University of Waterloo, Canada
Waterloo, Canada

Chris Eliasmith is a professor in the Philosophy and Systems Design Engineering departments at the University of Waterloo. He holds a Canada Research Chair in Theoretical Neuroscience, and is the director of the Centre for Theoretical Neuroscience at Waterloo.

References
1. H. Markram, *The blue brain project*, **Nat. Rev. Neurosci. 7** (2), pp. 153–160, 2006.
2. R. Ananthanarayanan and D. S. Modha, *Anatomy of a cortical simulator*, **Proc. ACM/IEEE Conf. Supercomput.**, 2007. Art. no. 3, doi:10.1145/1362622.1362627
3. C. Eliasmith and C. H. Anderson, **Neural Engineering: Computation, Representation and Dynamics in Neurobiological Systems**, MIT Press, Cambridge, MA, 2003.
4. J. Conklin and C. Eliasmith, *An attractor network model of path integration in the rat*, **J. Comput. Neurosci. 18**, pp. 183–203, 2005.
5. R. Singh and C. Eliasmith, *Higher-dimensional neurons explain the tuning and dynamics of working memory cells*, **J. Neurosci. 26**, pp. 3667–3678, 2006.

6. D. Rasmussen and C. Eliasmith, *A neural model of rule generation in inductive reasoning*, **Topics Cognitive Sci. 3** (1), pp. 140–153, 2011.

7. T. Stewart, T. Bekolay, and C. Eliasmith, *Learning to select actions with spiking neurons in the basal ganglia*, **Front. Neurosci. 6**, 2012.

8. C. Eliasmith, **How to Build a Brain: A Neural Architecture for Biological Cognition**, Oxford University Press, New York, NY, in press.

9. S. Choudhary, S. Sloan, S. Fok, A. Neckar, Eric, Trautmann, P. Gao, *et al.*, *Silicon neurons that compute*, **Int'l Conf. Artif. Neural Netw. Mach. Learn.**, pp. 121–128, 2012.

10. F. Galluppi, S. Davies, T. Stewart, C. Eliasmith, and S. Furber, *Real time on-chip implementation of dynamical systems with spiking neurons*, **Proc. IJCNN**, 2012.

11. http://nengo.ca Nengo software for simulating large-scale neural systems. Accessed 16 October 2012.

12. C. Eliasmith, T. C. Stewart, X. Choo, T. Bekolay, T. DeWolf, Y. Tang, and D. Rasmussen, *A large-scale model of the functioning brain*, **Science**, in press.

13. http://nengo.ca/build-a-brain/spaunvideos/ Videos for Spaun simulations. Accessed 16 October 2012.