Nengo and the Neural Engineering Framework: From Spikes to Cognition

Chris Eliasmith (celiasmith@uwaterloo.ca)
Terrence C. Stewart (tcstewar@uwaterloo.ca)
Center for Theoretical Neuroscience, University of Waterloo
200 University Ave West, Waterloo, ON, N2L 3G1, Canada

Keywords: cognitive modeling; neural engineering; representation; decision making; working memory; cognitive architecture; cognitive control

Tutorial Objectives

As we learn more about the neural activity underlying cognitive function, there is an increasing demand to explicitly and quantitatively connect cognitive theories to neurological details. Bridging these levels provides benefits in both directions; aspects of the cognitive theory can predict and be constrained by neurological details, and the neurological details can in turn identify important modifications to the overall cognitive theory.

This tutorial introduces the Neural Engineering Framework (NEF; Eliasmith and Anderson, 2003) and the associated open-source toolkit Nengo (<http://nengo.ca>), which offer a general method for implementing high-level cognitive theories using biologically realistic spiking neurons. This approach takes a high-level description of a cognitive theory (in terms of information being represented and transformed) and combines it with relevant anatomical and neurophysiological constraints, producing a detailed mechanistic model of how interacting neurons can efficiently produce the desired behaviour. The resulting models can be run to produce predictions of spike patterns, firing rates, fMRI time-courses, accuracy, reaction times, and overall behaviour. Complete details can be found in the book *How to Build a Brain* (Eliasmith, 2012; to be released by OUP at CogSci 2012).

These methods have been made more accessible by the construction of the software package Nengo, which provides a graphical interface suitable for network construction. This tutorial introduces the NEF theory explaining how high-level function can be systematically related to single cell activity, and provides extensive hands-on experience building these neural models using Nengo. Our central objective is to allow participants to leave the tutorial with a method for constructing cognitive models with spiking neurons, and experience using that method in an intuitive software environment.

Tutorial Structure

This full-day tutorial combines the theoretical bases of the Neural Engineering Framework with hands-on examples of practically applying these concepts using Nengo. For example, the presentation of the theory for how a scalar value can be represented by the spiking pattern in a group of neurons is paired with a tutorial on using Nengo to generate such a neural group and simulate its behavior over time. Participants are expected to bring a laptop to follow along with these tutorials (Windows, OS X, and Linux are all supported, and software is provided).

In particular, the tutorial covers using the NEF to represent scalars and vectors, perform linear and nonlinear transformations on these values, and store information over time. These are the basic mechanisms required for a wide range of algorithms, and form the basis for our models of sensorimotor systems, working memory, and cognitive control. This provides participants with basic building blocks for constructing novel neural implementations of a wide variety of cognitive models.

To supplement this, we more closely examine how cognitive theories can be expressed in terms of vectors and transformations. The basic approach of employing *semantic pointers* (vectors that combine the benefits of semantic similarity measures with the compositionality of symbol structures) is described. We show how this method provides a unified approach to many types of cognitive models, including perceptual, symbolic reasoning, and motor control models. For example, we show how to construct a non-classical symbol system, capable of performing the operations required for symbolic cognition. The result is a scalable and efficient neural cognitive architecture, constructed from the basic approaches described in the first half of the tutorial.

Finally, we explore recent results in building whole-brain models using the NEF. This involves a fully integrative model spanning vision, object recognition, working memory, cognitive control, and motor control to produce a neural cognitive architecture. This ~3 million neuron model is built in Nengo, uses images for input, draws digits using a 2-joint arm as its output, and is performs a variety of tasks, including list memory, mental addition, inductive reasoning over symbols, and reinforcement learning. The tutorial covers this model and its behavioural and neurobiological constraints, including the dopaminergic learning system.

Variants of this tutorial were presented at ICCM 2009, CogSci 2010, Telluride 2011, and CogSci 2011. An on-line tutorial is available at <http://nengo.ca>, and significant changes have been made in terms of scaling Nengo models up to larger neuron counts and more complex behaviour.

Tutorial Justification

The Neural Engineering Framework provides a method to bridge the gap between cognitive and neural theories. Its earlier applications have been to sensory and motor systems, including the barn owl auditory system, rodent navigation, swimming control in zebrafish, and the vestibular ocular reflex in monkeys. However, these same principles are now being applied to cognitive models. This includes models of serial-order recall (Choo & Eliasmith, 2010), action selection in the basal ganglia (Stewart, Choo, & Eliasmith, 2010), visual working memory (Singh & Eliasmith, 2006), deep belief networks for visual recognition (Tang & Eliasmith, 2010), the Wason card task (Eliasmith, 2005), the Tower of Hanoi task (Stewart & Eliasmith, 2011), and a model of inductive rule generation that received the computational modelling prize in higher-level cognition at CogSci 2010 (Rasmussen & Eliasmith, 2010).

While we find that the Neural Engineering Framework produces extensive new insights into the neural grounding of cognitive function, we also find that the underlying mathematics and a lack of familiarity with biologically realistic neural modeling have been a significant barrier to entry for new researchers. As a result, we feel that a full-day tutorial is most appropriate for introducing the necessary concepts from control theory, signal theory, and theoretical neuroscience. Feedback from previous tutorials has been extremely positive, with participants now using Nengo for their own research and in the classroom at the University of Manchester, Rensaeller, Yale, Franklin & Marshall College, and Stanford.

The NEF provides an exciting new tool for cognitive science, as it provides a technique for producing direct neural predictions from cognitive theory. Furthermore, it leads to important theoretical results as to the relationships between neural properties and the high-level algorithms they are capable of implementing (e.g. the relationship between neurotransmitter re-uptake rate and the 50ms cognitive cycle time; Stewart, Choo, & Eliasmith, 2010).

These consequences are also very general, as the NEF provides techniques that can be applied to a wide variety of cognitive theories. It provides a structure for organizing a high-level description such that it can be implemented by realistic spiking neurons, providing meaningful data in terms of the expected spike patterns, time course, and behavioural accuracy. We have made use of it in a variety of contexts, and have developed tools that support the creation and analysis of these models. Tutorial participants will gain hands-on experience with a tool that helps generate new models and can be applied to existing models. In both cases, these tools will help participants incorporate ever-more-abundant neural data into their research.

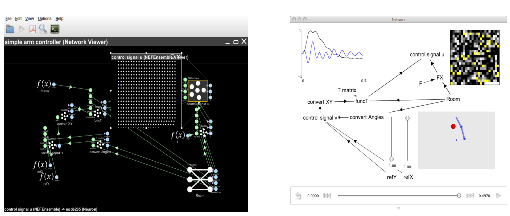

Figure 1: The Nengo interface. Network construction (left) is done either by point-and-click or by Python scripting. Visualization (right) provides on-the-fly control of inputs with plots of spiking activity, decoded representations, etc.

Audience

Participants are not expected to have any previous experience with neural modeling. All participants are encouraged to bring a laptop for installing Nengo (Linux, OS X, and Windows versions are provided), allowing for hands-on interactions with the models discussed.

Presenters

Chris Eliasmith holds a Canada Research Chair in Theoretical Neuroscience, and is director of the Centre for Theoretical Neuroscience at the University of Waterloo. He has over 50 publications spanning neuroscience, psychology, philosophy, computer science, and engineering, on topics including working memory, mental representation, population coding, neural dynamics, computation, automatic text classification, and cognitive architectures. His recent book, *How to Build a Brain*, and his earlier book, *Neural Engineering*, form the basis for this tutorial.

Terry Stewart is a postdoc in the Centre for Theoretical Neuroscience, and has developed large-scale models with the NEF, including the Tower of Hanoi task, focussing on problems of cognitive control.

References

Choo, F., & Eliasmith, C. (2010). A Spiking Neuron Model of Serial-Order Recall. *32nd Annual Conference of the Cognitive Science Society.*

Eliasmith, C. (2005). Cognition with neurons: A large-scale, biologically realistic model of the Wason task. *27th Annual Meeting of the Cognitive Science Society.*

Eliasmith, C. (2012). *How to build a brain: A neural architecture for biological cognition.* New York, NY: Oxford University Press.

Eliasmith, C., & Anderson, C. (2003). *Neural Engineering: Computation, Representation, and Dynamics in Neurobiological Systems.* Cambridge: MIT Press.

Rasmussen, D., & Eliasmith, C. (2010). A neural model of rule generation in inductive reasoning. *32nd Annual Conference of the Cognitive Science Society.*

Singh, R., & Eliasmith, C. (2006). Higher-dimensional neurons explain the tuning and dynamics of working memory cells. *Journal of Neuroscience*, *26*, 3667-3678.

Stewart, T.C., & Eliasmith, C. (2010). Neural symbolic decision making: A scalable and realistic foundation for cognitive architectures. *1st Annual Meeting of the Biologically Inspired Cognitive Architectures Society.*

Stewart, T.C., & Eliasmith, C. (2011). Neural Cognitive Modelling: A Biologically Constrained Spiking Neuron Model of the Tower of Hanoi Task. 33rd Annual Conference of the Cognitive Science Society.

Stewart, T.C., Choo, X., & Eliasmith, C. (2010). Dynamic Behaviour of a Spiking Model of Action Selection in the Basal Ganglia. *10th Int. Conf. on Cognitive Modeling.*

Tang, Y., Eliasmith, C. (2010). Deep networks for robust visual recognition. *International Conference on Machine Learning.*