Nengo and the Neural Engineering Framework: Connecting Cognitive Theory to Neuroscience

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

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 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. The NEF allows researchers to 1) provide a high-level description of a cognitive theory (in terms of information being represented and transformed) and 2) identify relevant neural constraints (anatomical, neurophysiological, and so on). It then produces a detailed model of neural activity, including predicted spike patterns, firing rates, connectivity, and overall behaviour.

These methods have been made more accessible by the construction of a 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

The tutorial is structured so as to combine the theoretical bases of the Neural Engineering Framework with hands-on examples of practically applying these concepts. To do this, we make use of Nengo <nengo.ca>, an open-source Java-based neural simulator that supports the NEF. 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 encouraged 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 the theory of Vector Symbolic Architectures can be implemented using the NEF. This involves using high-dimensional fixed-length vectors to represent symbols and symbol trees. The nonlinear operation of circular convolution is used to manipulate these symbol trees. This can be seen as 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, a variety of other uses of the NEF are provided. This includes learning rules for modifying synaptic connection weights (with examples for implementing an associative memory and reinforcement learning), a model of the Wason card task (symbol manipulation and generalization), and a model of the basal ganglia-thalamus-cortex loop which implements a basic production system. Together, these hands-on examples will introduce participants to many of the major components needed to address a wide variety of cognitive behaviour.

A previous version of this tutorial was presented at ICCM 2009 and CogSci 2010. Slides and step-by-step instructions are available at <http://ctn.uwaterloo.ca/~cnrglab/>. As a result of feedback from these tutorials, we have continued to improve Nengo's user interface, making common actions easier and developing new displays for observing and the ongoing neural activity and adjusting its inputs as the simulation runs (see Figure 1).

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, escape and swimming control in zebrafish, and the translational vestibular ocular reflex in monkeys. However, these same principles are now being applied to higher-level 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), a general-purpose cognitive architecture (Stewart & Eliasmith, 2010), 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 have also found 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 our tutorial last year was extremely positive, with at least two participants going on to teach Nengo to their students at their own universities.

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 time constant of neural transformations).

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 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 can help generate new models as well as be applied to existing models. In both cases, these tools will help participants incorporate ever-more-abundant neural data into their research.

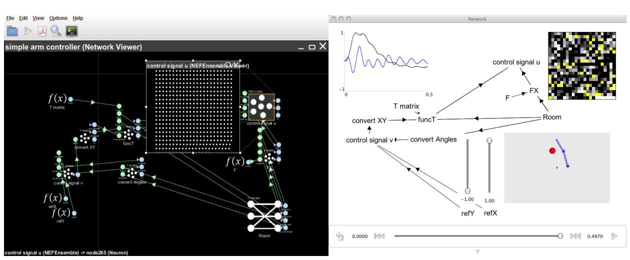

Figure 1: The Nengo interface. Network construction (left) is done either by point-and-click or by Python scripting. Visualization (right) provides on-the-fly control of inputs with plots of spiking activity, decoded representations, etc.

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 book, Neural Engineering, with Charles Anderson is now in paperback with MIT Press, and forms the basis for this tutorial.

Terry Stewart is a postdoc in the Centre for Theoretical Neuroscience, with a PhD in Cognitive Science examining the methodological issues surrounding the creation and evaluation of computational cognitive models. His current work applies the Neural Engineering Framework to develop a complete neural cognitive architecture.

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

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