

A unified theoretical approach for biological cognition and learning

Brent Komer and Chris Eliasmith



Large-scale neural models are needed in order to understand the biological underpinnings of complex cognitive behavior. Good methods for constructing such models should provide for: first, abstraction (analysis across levels of description); second, integration (incorporation of simpler models to build more complex ones); third, empirical contact (using and comparing to a wide variety of neural data); and fourth, account for the varieties of learning. In this review we evaluate three prominent recent methods for constructing neural models using these four criteria. Each of these methods is being actively developed and demonstrates clear strengths along some of these criteria.

Address

Centre for Theoretical Neuroscience, University of Waterloo, Canada

Corresponding author: Eliasmith, Chris (celiasmith@uwaterloo.ca)

Current Opinion in Behavioral Sciences 2016, 11:14–20

This review comes from a themed issue on **Computational modelling**

Edited by **Peter Dayan** and **Daniel Durstewitz**

<http://dx.doi.org/10.1016/j.cobeha.2016.03.006>

2352-1546/© 2016 Elsevier Ltd. All rights reserved

Introduction

There is general agreement that the vast majority of natural behavior is the result of the interactions between many brain areas, implicating huge numbers of neurons. Without constructing biologically detailed models that explore these complex interactions, we are unlikely to be able to understand how to help a distressed brain [1], or explain how low-level biological mechanisms give rise to cognitive behavior [2,3]. In short, without large-scale models, we cannot test large-scale hypotheses.

To understand biological cognition — that is, the link between biological mechanisms and cognitive behavior — we need to understand the complex interactions between individual neurons as well as their integration into a behaving, cognitive agent. This is a challenging and long-term goal, but it is nevertheless helpful to consider contemporary theoretical approaches in this context to determine promising avenues for research. Addressing cognitive behaviors in a neural model typically requires a

large-scale model: one simulating tens of thousands to several million neurons. This is due to the correlation between the complexity in a task and its likelihood of being deemed ‘cognitive.’ Complexity suggests that large numbers of neural resources, and often many different brain areas, are required to address the challenge embodied by the task.

Large-scale neural models are ineluctably complex. Consequently, in order for us to successfully understand and construct them, it is crucial to employ systematic methods. Ideally, good methods would provide [4,5^{*}]: first, smoothly moving between levels of detail depending on questions of interest (*Abstraction*); second, easy integration of simpler partial models (*Integration*); third, generation of empirical predictions while providing deeper insight into collected neural and behavioral data (*Empirical contact*) and fourth, methods for including adaptation and learning (*Learning*). We review recent work at the interface of biological and behavioral modeling along these four criteria. Specifically, we focus on methods from Denève *et al.*, efficiently implementing dynamical systems in spiking neurons [6^{**}], large-scale neural models proposed by Markram *et al.* [7^{**}], and our own work [5^{*}].

Dynamics in spiking networks

Abstraction

Sophie Denève and her collaborators have recently developed and explored a method for constructing spiking neuron implementations of dynamical systems [6^{**},8,9]. These researchers have shown the method can be used to build robust models of neural integrators, simple arm controllers, and optimal linear filters (e.g., the Kalman filter). Consequently, they have defined an effective means of abstracting from spiking neural networks to linear dynamical systems.

This method consists of treating subthreshold voltages of individual cells as tracking the error in the estimate of an underlying multi-dimensional state variable. When the error becomes large, a neuron in a given population representing that state variable emits a spike, updating the representation of that state. In order to ensure that neurons are efficient in their representation of such states, when a given cell spikes, it instantaneously inhibits other neurons in the population from spiking to account for that same error. Consequently, these networks can be used to represent dynamical systems defined over the represented state variables. In addition, these methods are unique in providing a clear relation to predictive coding.

These methods have been noted to have several limitations. Chief amongst these is that they apply only to linear dynamical systems. However, recent work has suggested methods for incorporating nonlinearities via dendritic computation in such networks [10]. It is not clear, however, if these methods will scale ([11], p. 156). In addition, the speed of inhibition may not be biologically plausible. As well, the work has been criticized as applying to simple point neuron integrate-and-fire models only. However, it has recently been demonstrated that similar methods can be applied to sophisticated conductance based neuronal models [9].

Given that the characterization of function is in terms of a state space, and that conductance models can be used, these methods provide multiple levels of abstraction. However, the methods do not yet provide abstractions for relating neural-level circuits to sophisticated, cognitive behavior.

Integration

Integration of simpler models should be possible with this approach because of the employed state space representations. However, there are not yet clear examples of performing such an integration. Possible barriers to integration include assumptions about the nature of inhibition and excitation included in the methods, which do not apply to subcortical areas, such as basal ganglia and cerebellum. For instance, basal ganglia consists largely of inhibitory projections, while these methods assume both excitatory and inhibitory interactions are available between all neurons in the circuit.

Empirical contact

These networks exhibit several desirable traits [6**]. For instance, they naturally exhibit Poisson-like firing variability at the level of individual spike trains. As well, they are consistent with the observation that many biological circuits balance the effects of inhibition and excitation. In addition, they are highly efficient in their use of spikes and are robust to neuronal death and other perturbations. They have been shown to efficiently implement behaviorally relevant dynamical systems including integrators, oscillators, and simple arm controllers [6**]. However, there is also evidence that some networks, like those of the oculomotor integrator, do not exhibit the kind of spike train variability expected under this characterization of neural computation [12].

Learning

Very recently, Bourdoukan and Denève [13] have begun to explore methods for learning connection weights in these kinds of networks. They have shown that with different local learning rules for fast and slow recurrent connections, it is possible to feedback a global supervised error signal that can tune the network.

Blue Brain Project

Abstraction

By contrast to Denève's theoretically driven approach, recent work from Henry Markram's lab has focussed on building large-scale brain models with significant biological detail. In short, the approach taken by the Blue Brain Project [14] and now incorporated into the Human Brain Project (HBP; [15*]) is to build models from the bottom up, with the expectation that interesting behavioral features will emerge. This work has recently culminated with the release of a paper with 82 co-authors that describes a detailed model of a small section of rodent somatosensory cortex [7**]. This model includes about 31,000 cells and 37 million synapses amongst 207 different kinds of neurons.

This approach to modeling does not bear any clear relation to abstractions of neural computation. Consequently, the model is largely used to replicate low-level experimental data that does bear a clear relation to observable behavior. While addressing such data is clearly important, the approach taken does not permit developing a theory of how neural activity relates to the behavior it controls.

Integration

In the context of this project, integration is largely focussed on methods for combining individual neuron models. Connections between neurons were determined using five rules of connectivity that provide an effective means of filling in statistically matched anatomical connections between neuron types. This permits the anatomical combination of single cell models, addressing the integration of simpler neuron models into larger circuits. However, this kind of integration does not extend to functional or behaviourally relevant networks, as such networks have not been built with these methods.

Empirical contact

These methods have very extensive empirical contact with neural data since a significant amount of the modeling effort has gone into inferring parametric information based on available neural data. There is good evidence that the rules of connectivity used in the model, along with other parameter fitting methods, are successful for generating statistically similar activity in the circuit compared to *in vivo* recordings. As well, the rules of connectivity were used to generate predictions regarding the number of connections and number of synapses per connection, as well as a wide variety of other connectivity and physiological properties of synapses.

At the circuit level, the model was used to examine the spatio-temporal patterns that result from simulated thalamic input in combination with variations in calcium levels. The model demonstrated that different levels of calcium resulted in synchronous and asynchronous neural firing in the model, in agreement with experimental

results. However, it is not clear how these methods will permit relating such local activity to observable behavior.

Nevertheless, the models provided are currently unsurpassed in the amount of molecular detail provided compared to other work, providing strong empirical contact with certain areas of neuroscience.

Learning

While it should be possible to introduce neurally plausible connection strength adaptation into such a model, it is not clear how such learning rules would affect the behavior of the circuit. This aspect of neural function is unexplored with these methods at present.

The NEF and the SPA

Abstraction

The Neural Engineering Framework (NEF) identifies three quantitatively specified principles that can be used to implement nonlinear dynamical systems in a spiking neural substrate [11]. Specifically, these principles provide methods for characterizing representation, computation, and dynamics in such networks. Consequently, given a hypothesis about what dynamical system a brain area might be implementing, the NEF can be used to build a spiking neural network that approximates that dynamical system. For instance, we might characterize working memory as a dynamical system that loads in a 2-dimensional representation during stimulus presentation and is stable during a delay period [16]. The NEF can then be used to efficiently build a network that will approximate those dynamics with spiking neurons, and allow careful and direct comparison to biological neural networks thought to underlie the same behavior. These methods have been used to propose novel models of a wide variety of neural systems [17–19] as well as to better understand more general issues about neural function [20,21].

The NEF answers questions about *how* neural systems might compute, but it does not address the issue of *what*, specifically, is computed by biological brains. We have recently addressed this second question by proposing a general neural architecture that includes specific functional hypotheses. We call this proposal the Semantic Pointer Architecture (SPA; [5^{*}]). The SPA identifies a generic means of characterizing neural representations, ‘semantic pointers,’ that are used to capture central features of perceptual, motor, and cognitive representation. Briefly, semantic pointers are neural representations of information compressed via lossy operators. The specific operator hypothesized to be used in a neural system can vary: for example, an image specific statistical operator in visual systems, a dimensionality reduced path planner in motor systems, or a compressive binding operator for representing cognitive structure. These representations are called ‘pointers’ because they are

efficient, low-dimensional (because compressed) representations of what they point to (i.e., the uncompressed information from which they were generated). They are ‘semantic’ because the contents of the pointer (unlike a typical computer pointer) are determined by the contents of what they point to. The SPA uses semantic pointers to explain, among other things, perceptual categorization, motor control, working and long-term memory, as well as conceptual binding and structure representation. Furthermore, the SPA includes an account of cognitive control that relies on basal ganglia and thalamic interactions with cortex for understanding action selection.

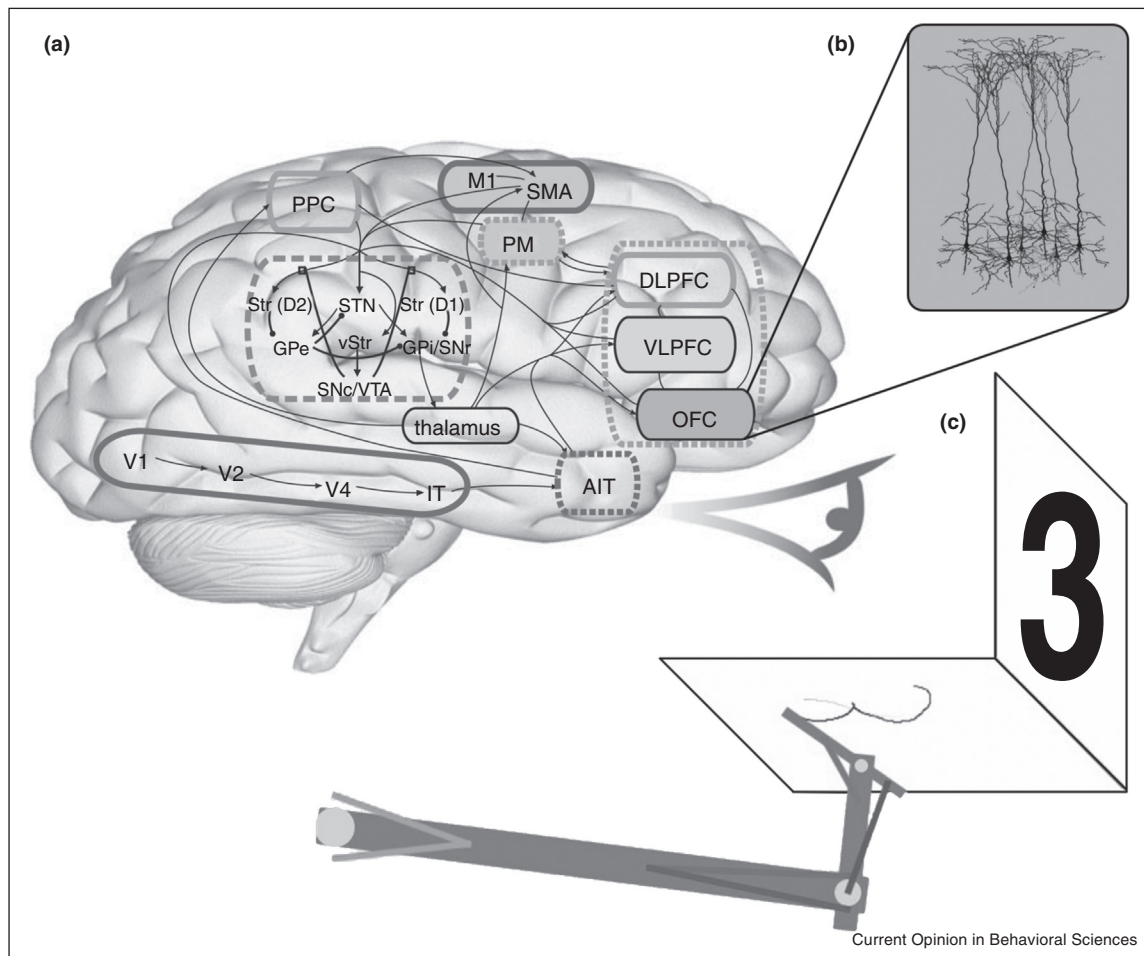
The NEF/SPA approach describes brain circuits at various levels of cognitive and neural detail. The neural simulation environment Nengo incorporates these methods, allowing the simulation of different parts of the same model at different levels of detail, including: direct solution of the dynamical system; rate neurons; or as any of many spiking point neuron models [22,23]. Tripp [24] proposed the use of surrogate population models to accelerate large-scale neural system modeling in Nengo, adding a level of simulation abstraction between neurons and the direct dynamical system solution.

Furthermore, Nengo was recently used in tandem with the NEURON simulator [25] to replace simple point neurons in the Spaun model (Figure 1; see next section) with complex conductance models [26], of the type used in the HBP models. These neurons allowed the simulation of the addition of the drug TTX to frontal cortex to examine the cognitive effects of this molecular intervention. Consequently, this one model flexibly incorporates a variety of levels of detail in a single simulation. This approach thus unifies simulation across levels of abstraction, from the molecular to the cognitive.

Integration

To demonstrate the SPA in detail, we proposed a mechanistic, functional model of the brain that uses 2.5 million spiking neurons, has about eight billion synaptic connections, and performs eight different tasks (see Figure 1; [5^{*},27^{*}]). We refer to this large-scale neural model as the Semantic Pointer Architecture Unified Network (Spaun). Spaun includes 20 anatomical areas with matched physiological properties for cells in those areas, four types of neurotransmitters (GABA, AMPA, NMDA, and Dopamine), and includes only known anatomical connections and structure. Spaun is the result of combining a variety of simpler past models that account for focussed behavioral results (e.g., working memory, motor control, action selection). Given the large number of connections in Spaun, it is a demonstration of the importance of recurrent communication throughout a large-scale network for accomplishing a variety of cognitive and non-cognitive tasks.

Figure 1



Semantic Pointer Architecture Unified Network (Spaun). **(a)** The neuroanatomical architecture of Spaun. *Abbreviations:* V1/V2/V4 (primary/secondary/extrastriate visual cortex), AIT/IT (anterior/inferotemporal cortex), DLPFC/VLPFC/OFC (dorso-lateral/ventro-lateral/orbito-frontal cortex), PPC (posterior parietal cortex), M1 (primary motor cortex), SMA (supplementary motor area), PM (premotor cortex), vStr (ventral/striatum), STN (subthalamic nucleus), GPe/i (globus pallidus externus/internus), SNc/r (substantia nigra pars compacta/reticulata), VTA (ventral tegmental area). **(b)** Replacement of leaky integrate-and-fire neurons with compartmental conductance neurons in the OFC. This replacement was used to simulate the effects of TTX on cognitive behavior. **(c)** Input to the model is through images shown to a simulated eye. Behavioral output of the model is in the form of moving a simulated arm with muscles. Modified from Eliasmith [5*].

What makes Spaun unique among large-scale brain models is its functional abilities. Spaun receives input from the environment through a single eye and manipulates the environment by moving a physically modeled arm. Spaun uses these interfaces to perform activities, ranging from perceptual-motor tasks (recreating the appearance of a perceived digit) to reinforcement learning (RL) (in a gambling task) to language-like inductive reasoning (completing abstract patterns in observed sequences of digits). It can perform the tasks in any order, they are all executed by the same model, and there are no changes to the model between tasks. To watch the model perform these tasks, see <http://nengo.ca/build-a-brain/spaunvideos>.

Empirical contact

NEF models have made a variety of successful predictions. For instance, an NEF model of path integration [28] has been credited with predicting both grid cells [29] and that path integration is independent of head direction, as confirmed by Maurer *et al.* [30].

Similarly SPA models have received broad empirical support. For instance, we have compared the performance of Spaun to human and animal data at several levels of detail [25]. Along many metrics the two align; for example, the model and the brain share; first, dynamics of firing rate changes in striatum during a gambling task; second, error rates as a function of position when

reporting digits in a memorized list; third, coefficient of variation of inter-spike intervals; fourth, reaction time mean and variance as a function of sequence length in a counting task; fifth, accuracy rates of recognizing unfamiliar handwritten digits; and sixth, success rates when solving induction tasks similar to those found on the Raven's Progressive Matrices, among other measures.

More recently, we have developed a variety of models using the SPA that demonstrate cognitive abilities not captured by Spaun. For instance, we have demonstrated the encoding and decoding of the 114,000 concepts and their relations in the WordNet database [31]. We have shown an SPA model that matches human performance on the Raven's Progressive Matrices (RPM) intelligence test, and demonstrated its ability to capture aging effects through biological manipulation [32,33]. More abstractly, we have described how the SPA naturally unifies the three most prevalent theories of conceptual representation [34*]. We have also shown simple models of language parsing [35,36], instruction following [37], and the N-back task [38]. In addition, models have been built using the SPA to explain a variety of other psychological-level phenomena, such as priming [39], emotion [40,41], and consciousness [42].

Learning

Because the NEF often employs a non-biological optimization method for network synthesis, it has been criticized as not accounting for learning [43]. However, much of our recent work has been focussed on neural adaptation in various guises. This includes simple RL in Spaun, RL models in more sophisticated semi-Markov¹ environments [44], and the first biologically plausible spiking model of hierarchical RL [45] (Rasmussen and Eliasmith, submitted for publication). We have also described a learning rule that is able to account for STDP effects as well as learn high-dimensional, nonlinear functions, including those posited by the SPA [22,23,46]. We have demonstrated a means of learning auto-associative and hetero-associative memories in a spiking network [47], and provided an account of motor control that reproduces a wide variety of adaptive motor behavior [48]. In short, the NEF/SPA has been used to produce examples of self-organized, error-driven, and RL using biologically plausible rules in spiking networks. We find the most effective method of constructing large-scale models is to combine the efficient non-biological methods with less efficient biologically plausible learning in order to address whichever kinds of adaptation are of interest in a given model.

Conclusion

Building large-scale models of biological cognition is in its infancy. We have reviewed some of methods worth

¹ Unlike a standard Markov process, the probability of there being a change in the state for a semi-Markov process depends on how much time has elapsed since the last state change.

considering in addressing this challenge. We would argue that the NEF and SPA currently have a broader demonstrated scope than current alternatives, both in terms of scaling from biological detail to cognitive behavior, and in terms of capturing crucial neural mechanisms, such as those behind working memory, action selection, motor control and learning, among others.

One common challenge for any large-scale modeling effort is accessing significant computational resources. The HBP modelers have put significant effort into developing a supercomputer infrastructure to simulate their model. Similarly, we have developed a Nengo component to compile NEF/SPA models to that same kind of infrastructure. In some sense, Nengo solves the basic technical problems involved with flexibly running many different kinds of models concurrently. This ability in Nengo is generic. However, as we have designed this software tool, we have also incorporated our theories regarding neural computation (NEF) and neural organization to generate cognitive function (SPA) that leverage the ability of the software to flexibly integrate multiple models.

More interestingly, researchers have been developing specialized, efficient 'neuromorphic' hardware to run brain-like algorithms in real-time [49–51]. Nengo is able to target this hardware as well [52,53*] opening up the possibility of testing large-scale brain models in real-world environments — perhaps the most challenging test of our ability to replicate the robust flexibility of biological brains.

Conflict of interest

Nothing declared.

Acknowledgements

This work was supported by CFI and OIT infrastructure funding and grants from Canada Research Chairs, NSERC Discovery Grant 261453, ONR Grant N000141310419, and AFOSR Grant FA8655-13-1-3084.

References and recommended reading

Papers of particular interest, published within the period of review, have been highlighted as:

- of special interest
 - of outstanding interest
1. Just MA, Keller TA, Malave VL, Kana RK, Varma S: **Autism as a neural systems disorder: a theory of frontal-posterior underconnectivity**. *Neurosci Biobehav Rev* 2012, **36**:1292-1313.
 2. Beeler JA, Frank MJ, McDaid J, Alexander E, Turkson S, Hernandez MS, Zhuang X: **A role for dopamine-mediated learning in the pathophysiology and treatment of Parkinson's disease**. *Cell Rep* 2012, **2**:1747-1761.
 3. Tyrcha J, Roudi Y, Marsili M, Hertz J: **The effect of nonstationarity on models inferred from neural data**. *J Stat Mech Theory Exp* 2013, **2013**:P03005.
 4. Eliasmith C, Trujillo O: **The use and abuse of large-scale brain models**. *Curr Opin Neurobiol* 2014, **25**:1-6.
 5. Eliasmith C: **How to Build a Brain: A Neural Architecture for Biological Cognition**. Oxford University Press; 2013.

This book describes an approach to building large-scale unified models of cognition using spiking neurons. The approach describes how the Neural Engineering Framework (NEF), a method for characterizing neural computation, underwrites the Semantic Pointer Architecture (SPA), which is a proposal for how specific cognitive functions and representations can give rise to integrated behavior.

6. Boerlin M, Machens CK, Denève S: **Predictive coding of dynamical variables in balanced spiking networks**. 2013.
This paper describes a method for building spiking networks that rely on the precise timing of spikes to represent variables in a network and compute linear functions of those variables. The methods can be used to simulate any linear dynamical system efficiently, while producing Poisson-like spiking and balancing excitation and inhibition.
7. Markram H, Muller E, Ramaswamy S, Reimann MW, Abdellah M, Sanchez CA, Muñoz-Céspedes A: **Reconstruction and simulation of neocortical microcircuitry**. *Cell* 2015, **163**:456-492.
This paper describes a highly detailed model of the rodent somatosensory cortex at the cellular level. The model includes 31,000 neurons and 37 million synapses across 207 kinds of neurons. The focus is on reproducing synaptic and cellular physiological and anatomical data. Different oscillatory properties of the network as a function of calcium concentrations are also demonstrated.
8. Zeldenrust F, Gutkin BS, Denève S: **Functional interpretation of biophysical properties of spiking neurons**. *BMC Neurosci* 2013, **14**(Suppl 1):P104.
9. Schwemmer MA, Fairhall AL, Denève S, Shea-Brown ET: **Constructing precisely computing networks with biophysical spiking neurons**. *J Neurosci* 2015, **35**:10112-10134 <http://dx.doi.org/10.1523/JNEUROSCI.4951-14.2015>.
10. Thalmeier D et al.: **Learning universal computations with spikes**. 2015arXiv:1505.07866.
11. Eliasmith C, Anderson CH: *Neural Engineering: Computation, Representation, and Dynamics in Neurobiological Systems*. Cambridge, MA: MIT Press; 2003.
12. Aksay E et al.: **Correlated discharge among cell pairs within the oculomotor horizontal velocity-to-position integrator**. *J Neurosci* 2003, **23**:10852-10858.
13. Bourdoukan R, Denève S: **Enforcing balance allows local supervised learning in spiking recurrent networks**. In *Advances in Neural Information Processing Systems*, vol 28. Edited by Cortes C, Lawrence ND, Lee DD, Sugiyama M, Garnett R. Curran Associates, Inc.; 2015:982-990.
14. Markram H: **The blue brain project**. *Nat Rev Neurosci* 2006, **7**:153-160.
15. Markram H: **The human brain project**. *Sci Am* 2012, **306**:50-55.
An accessible description of the goals and methods of the Human Brain Project. A central purpose is to build a human scale neural model with a high degree of biological plausibility.
16. Ray S, Eliasmith C: **Higher-dimensional neurons explain the tuning and dynamics of working memory cells**. *J Neurosci* 2006, **26**:3667-3678.
17. Bekolay T, Laubach M, Eliasmith C: **A spiking neural integrator model of the adaptive control of action by the medial prefrontal cortex**. *J Neurosci* 2014, **34**:1892-1902.
18. Stewart TC, Bekolay T, Eliasmith C: **Learning to select actions with spiking neurons in the basal ganglia**. *Front Neurosci* 2012:6.
19. Trujillo, Oliver, Eliasmith C: **A spiking-neuron model of memory encoding and replay in hippocampus**. *BMC Neurosci* 2014, **15**(Suppl 1):P166.
20. Tripp B, Eliasmith C: **Neural populations can induce reliable postsynaptic currents without observable spike rate changes or precise spike timing**. *Cereb Cortex* 2007, **17**:1830-1840.
21. Parisien C, Anderson CH, Eliasmith C: **Solving the problem of negative synaptic weights in cortical models**. *Neural Comput* 2008, **20**:1473-1494.
22. Bekolay T, Bergstra J, Hunsberger E, DeWolf T, Stewart TC, Rasmussen D, Eliasmith C: **Nengo: a Python tool for building large-scale functional brain models**. *Front Neuroinform* 2013:7.
23. Bekolay T, Kolbeck C, Eliasmith C: **Simultaneous unsupervised and supervised learning of cognitive functions in biologically plausible spiking neural networks**. In *Proceedings of the 35th Annual Conference of the Cognitive Science Society*. 2013:169-174.
24. Tripp BP: **Surrogate population models for large-scale neural simulations**. *Neural Comput* 2015, **27**:1186-1222 http://dx.doi.org/10.1162/NECO_a_00734.
25. Hines ML, Carnevale NT: **The NEURON simulation environment**. *Neural Comput* 1997, **9**:1179-1209.
26. Eliasmith C, Gosmann J, Choo X: **BioSpaun: a large-scale behaving brain model with complex neurons**. 2016 <http://arxiv.org/abs/1602.05220>.
27. Eliasmith C, Stewart TC, Choo X, Bekolay T, DeWolf T, Tang Y, Rasmussen D: **A large-scale model of the functioning brain**. *Science* 2012, **338**:1202-1205.
This paper describes what is currently the largest functional model of the human brain. The 2.5 million neurons and eight billion connections in the model are anatomically organized and physiologically constrained, while reproducing a variety of neural and behavioral data across eight tasks.
28. Conklin J, Eliasmith C: **A controlled attractor network model of path integration in the rat**. *J Comput Neurosci* 2005, **18**:183-203.
29. Zilli EA: **Models of grid cell spatial firing published 2005–2011**. *Front Neural Circuits* 2012, **6**:16.
30. Maurer AP, Lester AW, Burke SN, Ferng JJ, Barnes CA: **Back to the future: preserved hippocampal network activity during reverse ambulation**. *J Neurosci* 2014, **34**:15022-15031.
31. Crawford E, Gingerich M, Eliasmith C: **Biologically plausible, human-scale knowledge representation**. *Cogn Sci* 2015 <http://dx.doi.org/10.1111/cogs.12261>.
32. Rasmussen D, Eliasmith C: **A spiking neural model applied to the study of human performance and cognitive decline on Raven's advanced progressive matrices**. *Intelligence* 2014, **42**:53-82.
33. Rasmussen D, Eliasmith C: **A neural model of hierarchical reinforcement learning**. In *Proceedings of the 36th Annual Conference of the Cognitive Science Society*. 2014:1252-1257.
34. Blouw P, Solodkin E, Thagard P, Eliasmith C: **Concepts as semantic pointers: a framework and computational model**. *Cogn Sci* 2015.
This paper provides a demonstration that the SPA methods naturally account for the three major theories of concepts used in psychology.
35. Stewart TC, Choo X, Eliasmith C: **Sentence processing in spiking neurons: a biologically plausible left-corner parser**. *36th Annual Conference of the Cognitive Science Society; Cognitive Science Society: 2014:1533-1538*.
36. Stewart TC, Blouw P, Eliasmith C: *Explorations in Distributed Recurrent Biological Parsing*. 2015.
37. Choo X, Eliasmith C: **General instruction following in a large-scale biologically plausible brain model**. *Proc 35th Annu Conf Cogn Sci Soc*. 2013.
38. Gosmann J, Eliasmith C: **A spiking neural model of the n-back task**. *37th Annual Meeting of the Cognitive Science Society*. 2015:812-817.
39. Schröder T, Thagard P: **Priming: constraint satisfaction and interactive competition**. *Understanding Priming Effects in Social Psychology*. 2014:157.
40. Thagard P, Schröder T: *Emotions as Semantic Pointers: Constructive Neural Mechanisms. The Psychological Construction of Emotions*. New York: Guilford; 2013.
41. Schröder T, Stewart TC, Thagard P: **Intention, emotion, and action: a neural theory based on semantic pointers**. *Cogn Sci* 2014, **38**:851-880.
42. Thagard P, Stewart TC: **Two theories of consciousness: semantic pointer competition vs. information integration**. *Conscious Cogn* 2014, **30**:73-90.

43. O'Reilly RC, Hazy TE, Herd Seth A: **The leabra cognitive architecture: how to play 20 principles with nature and win!** *The Oxford Handbook of Cognitive Science*. 2012.
 44. Rasmussen D, Eliasmith C: **A neural reinforcement learning model for tasks with unknown time delays**. *35th Annual Conference of the Cognitive Science Society*. 2013:3257-3262.
 45. Rasmussen D: *Hierarchical reinforcement learning in a biologically plausible neural architecture*. PhD thesis, Computer Science. University of Waterloo; 2014.
 46. MacNeil D, Eliasmith C: **Fine-tuning and the stability of recurrent neural networks**. *PLoS One* 2011, **6**:e22885.
 47. Voelker AR, Crawford E, Eliasmith C: **Learning large-scale heteroassociative memories in spiking neurons**. *13th International Conference; UCNC: 2014*.
 48. DeWolf T: *A neural model of the motor control system*. 2015.
 49. Furber SB, Galluppi F, Temple S, Plana L: **The spinnaker project**. *Proc IEEE* 2014, **102**:652-665.
 50. Merolla PA, Arthur JV, Alvarez-Icaza R, Cassidy AS, Sawada J, Akopyan F, Modha DS: **A million spiking-neuron integrated circuit with a scalable communication network and interface**. *Science* 2014, **345**:668-673.
 51. Benjamin BV, Gao P, McQuinn E, Choudhary S, Chandrasekaran AR, Bussat JM, Boahen K: **Neurogrid: a mixed-analog-digital multichip system for large-scale neural simulations**. *Proc IEEE* 2014, **102**:699-716.
 52. Mundy A, Knight J, Stewart TC, Furber S: **An efficient SpiNNaker implementation of the neural engineering framework**. *International Joint Conference on Neural Networks (IJCNN); IEEE: 2015:1-8*.
 53. Choudhary S, Sloan S, Fok S, Neckar A, Trautmann E, Gao P, Stewart TC, Eliasmith C, Boahen K: **Silicon neurons that compute**. *International Conference on Artificial Neural Networks: 7552*. 2012:121-128.
- The earliest demonstration of using the NEF methods to program neuromorphic hardware. Each of the principles of the NEF are demonstrated, laying groundwork for energy efficient, real-time simulations of behaving networks.