TOPICS
TOPICS IN COGNITIVE SCIENCE

Topics in Cognitive Science (2016) 1–15

ISSN:1756-8757 print / 1756-8765 online
DOI: 10.1111/tops.12242

This article is part of the topic "Best of Papers From the Cognitive Science Society Annual Conference," Wayne D. Gray (Topic Editor). For a full listing of topic papers, see http://onlinelibrary.wiley.com/journal/10.1111/(ISSN)1756-8765/earlyview

Improving With Practice: A Neural Model of Mathematical Development

Sean Aubin, Aaron R. Voelker, Chris Eliasmith

Centre for Theoretical Neuroscience, University of Waterloo

Received 7 October 2016; accepted 19 October 2016

Abstract

The ability to improve in speed and accuracy as a result of repeating some task is an important hallmark of intelligent biological systems. Although gradual behavioral improvements from practice have been modeled in spiking neural networks, few such models have attempted to explain cognitive development of a task as complex as addition. In this work, we model the progression from a counting-based strategy for addition to a recall-based strategy. The model consists of two networks working in parallel: a slower basal ganglia loop and a faster cortical network. The slow network methodically computes the count from one digit given another, corresponding to the addition of two digits, whereas the fast network gradually "memorizes" the output from the slow network. The faster network eventually learns how to add the same digits that initially drove the behavior of the slower network. Performance of this model is demonstrated by simulating a fully spiking neural network that includes basal ganglia, thalamus, and various cortical areas. Consequently, the model incorporates various neuroanatomical data, in terms of brain areas used for calculation and makes psychologically testable predictions related to frequency of rehearsal. Furthermore, the model replicates developmental progression through addition strategies in terms of reaction times and accuracy, and naturally explains observed symptoms of dyscalculia.

Keywords: Neural engineering framework; Semantic pointer architecture; Nengo; Cognitive modeling; Mathematical ability; Dyscalculia; Skill consolidation

Correspondence should be sent to Sean Aubin, Centre for Theoretical Neuroscience, University of Waterloo Waterloo, ON, Canada N2L 3G1. E-mail: saubin@uwaterloo.ca

1. Introduction

Adaptability and scalability are two central challenges in cognitive modeling. Building a model that performs some behavior is a daunting enough task, even without considering how it might improve over time. This is made significantly more challenging when building a large-scale biologically plausible model that uses spiking neurons to represent and communicate information over time.

One example of such a biologically plausible model is the Semantic Pointer Architecture Unified Network (Spaun), which is currently the largest behaving model of the human brain (Eliasmith et al., 2012). This model flexibly performs eight different cognitive tasks by receiving images of handwritten digits through its simulated retina and outputting responses by drawing using its simulated arm. This model makes some progress in addressing the challenge of scalability (Eliasmith, 2013), but it lacks the ability to learn from previously experienced cognitive tasks to permanently improve its performance (Spaun only changes its long-term connection weights during a simple reinforcement learning task). In this work, we show that one task from Spaun's repertoire, *addition by counting*, can be extended to exhibit this cognitive ability.

For addition by counting, the model is presented with two digits and asked to draw the digit that corresponds to the sum of these two digits. The "counting" strategy methodically computes the result by adding one to a digit, for the number of times indicated by the other digit. In Spaun, this is accomplished by a number of structures that function together. Visual cortical areas compress the representation of a given image into a semantic representation, referred to as a *Semantic Pointer* (SP). Prefrontal cortical areas transform these pointers, while maintaining partial results in working memory, and a basal ganglia and thalamus control loop select actions to coordinate and drive the behavior of the system. The cortical areas of Spaun should, but currently do not, recognize and learn from previous instances of this problem. Including such adaptation should translate to performance improvements in both speed and accuracy.

DeWolf and Eliasmith (2013) have presented a neural model in which a simple motor skill is consolidated into cortex via repeated practice. However, the skill being learned in that case cannot be extended to solve the addition task because it remains unclear how a simple motor action could be generalized to count from one digit to another. Here, we consider how consolidation can occur for more sophisticated representations and demonstrate that such a perspective can help explain more complex cognitive phenomena. This is accomplished by proposing a spiking neural model that displays gradual performance improvement on the addition-by-counting task.

A previous model of this task has been implemented in ACT-R (Lebiere, 1999), however, this model has no neuroanatomical mapping and does not address the problem of transitioning between strategies.

2. Mathematical development

The development of numeracy in children is not a simple progression of skills. For instance, children learn to count before they understand how sets relate to numbers. As well, it is thought that after learning to count, they progressively learn the relationship between set sizes and numbers until the age of 5, after which they have a complete understanding of the number scale (Sarnecka & Carey, 2008). This same sort of complex learning development can be observed in children learning addition. Typically, children progress through various strategies before finally memorizing the results of addition, as shown in Table 1 (Siegler, 1987).

The *Counting* strategy involves choosing the larger number and incrementing it a number of times equal to the smaller number. The *Recall* strategy is where the two numbers form an association with a previously memorized answer, which is then recalled from long-term memory. Other strategies have been identified, although we focus here on recall and counting because the greatest developmental change is seen between these two strategies. In addition, for the sake of simplicity, we focus on the progression from the counting strategy to recall, using sums less than 10. As we demonstrate, the model initially relies entirely on the counting strategy, but it gradually learns the recall strategy, which improves reaction time and accuracy, consistent with the data from Tables 2 and 3, respectively.

To this point we have characterized the typical developmental path for children. However, there are individuals who suffer from dyscalculia. Dyscalculia is a learning disability characterized by various problems with numeracy, one of which can be understood as difficulty making the transition from counting to retrieval. Our model demonstrates how, without a parallel learning mechanism, symptoms of dyscalculia can arise. In particular, our model explains increased activation of the prefrontal cortex compared to individuals with normal numeracy (Kucian & von Aster, 2015) and a lack of progression to recall-based strategies.

3. Neural representation of digit semantics

To build our model, we make use of the Neural Engineering Framework (NEF; Eliasmith & Anderson, 2003) and the Semantic Pointer Architecture (SPA; Eliasmith, 2013) that were both used to build Spaun.

Table 1
Percentage of addition strategy use by grade level (summarized from Siegler, 1987)

Grade Level	Counting (%)	Recall (%)	Guess or No Response (%)	Other (%)
Kindergarten	30	16	30	24
Grade 1	38	44	8	10
Grade 2	40	45	5	11

Table 2
Median solution times (seconds) per addition strategy use by grade level (summarized from Siegler, 1987)

Grade Level	Counting	Recall
Kindergarten	6.0 s	3.9 s
Grade 1	6.9 s	2.1 s
Grade 2	3.9 s	1.8 s

Table 3
Percentage of errors per addition strategy use by grade level (summarized from Siegler, 1987)

Grade level	Counting (%)	Recall (%)
Kindergarten	19	29
Grade 1	4	17
Grade 2	3	7

The NEF may be understood as a "neural compiler" for mathematical functions using vector spaces. The presented model relies on two key principles of the NEF as a method for constructing networks of spiking neurons and connection weights from a mathematical description of the system. The first principle describes how a vector can be mapped onto a distributed representation over neurons. The second principle characterizes how connections between neurons can transform these vectors.

A vector $\mathbf{x}(t)$ is represented by encoding it into the spiking activity of a population of neurons. Each neuron i has an encoding vector $\mathbf{e}_i$ (which can be understood as a preferred direction in the vector space), a gain α_i, and a background current J_i^{bias}. These parameters determine how the input vector is translated into the input current $J_i(t)$ to a neural nonlinearity $G_i[\cdot]$. For our work, this neural nonlinearity is the leaky integrate-and-fire model, which converts the input current into a neural spike train $a_i(t)$.

$$a_i(t) = G_i[J_i(t)], \quad J_i(t) = \alpha_i \mathbf{e}_i \, . \, \mathbf{x}(t) + J_i^{bias} \tag{1}$$

To decode an approximation of the vector back from these spike trains, they are first convolved with a low-pass filter $h(t)$ (a decaying exponential modeled after the postsynaptic current) and then multiplied by a decoding vector $\mathbf{d}_i$:

$$\hat{\mathbf{x}}(t) = \sum_i \mathbf{d}_i (a_i * h)(t) \tag{2}$$

The decoders $\mathbf{d}_i$ are found using regularized least squares optimization to minimize the error over the range of inputs $\mathbf{x}$:

$$\int \|\mathbf{x} - \hat{\mathbf{x}}\|^2 d\mathbf{x} \tag{3}$$

To describe how two neural ensembles are connected, we define a weight matrix as the outer product of the encoders and decoders $\omega_{ij} = \mathbf{e}_i \times \mathbf{d}_j$. The second principle then shows that the input current to neuron i may be rewritten as a weighted summation of its postsynaptic potentials, allowing for the connection of multiple populations of neurons:

$$J_i(t) = \sum_j \alpha_i \omega_{ij}(a_j * h)(t) + J_i^{\text{bias}} \tag{4}$$

To apply arbitrary transformations to the vector using these connections, we can minimize the decoding error for the desired function:

$$\int \left\| f(\mathbf{x}) - \hat{f}(\mathbf{x}) \right\|^2 d\mathbf{x}. \tag{5}$$

Together, these two principles allow us to construct spiking neural models of arbitrary functions of vector spaces.

To apply the NEF to cognitive models, the SPA suggests specific architectural components and organization as well as a general kind of representation called a SP. SPs are compressed neural representations that can be efficiently manipulated, transformed, and dereferenced to retrieve deep semantic information. It has been suggested that this approach addresses the symbol grounding problem by defining the semantics of a neural representation as resulting from the compression of sensory information as well as conceptual relations (Eliasmith, 2013). Consequently, the SPA lends itself especially well to representing concepts which are grounded in multiple modalities but unified in a single representation. In this model, the SPs are used to represent digits, which have been shown to be grounded in auditory and visual modalities (Nieder, 2012), but also include conceptual relationships (e.g., TWO comes after ONE but before THREE). Conceptual relationships are captured in the SPA by adopting a compression operator from a specific type of Vector Symbolic Architecture (Gayler, 2004) called a Holographic Reduced Representation (Plate, 1995).

4. Modeling the counting strategy

Although Spaun serves as the starting point for our model, it is not practical to simulate all 2.5 million of its neurons. Instead, a counting circuit based on Spaun's design was implemented. As shown in Fig. 1, after a question input is received, the procedure consists of three main steps:

1. The digit contained in the working memory neural population is routed to a "transformation system."
2. The digit is transformed by a heteroassociative "incrementing memory" to produce the SP corresponding to the incremented digit.
3. The incremented number is returned to working memory.

This process is continued until the "Counts finished" are equal to the "Total counts to take." At which point the value in "Count result" is routed as the final answer to the output.

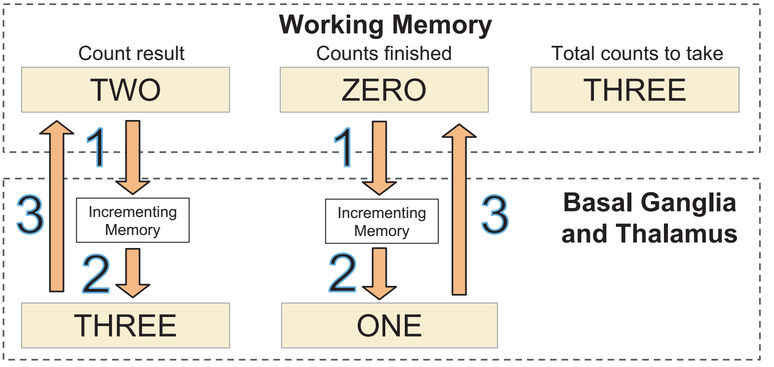

Fig. 1. Overview of the addition-by-counting procedure. See text for details.

These steps are controlled by an action selection system implemented in the basal ganglia and thalamus (Stewart, Choo, & Eliasmith, 2010) in a manner similar to Spaun.

To increment a digit, a predefined heteroassociative memory is used to associate each SP with its incremented pointer. This is implemented using a single population of neurons, with encoders tuned to each vector of the numerical vocabulary and decoders chosen to apply a transform to the output to compute the incremented SP. Finally, the output layer uses mutual inhibition to form a winner-take-all mechanism, which ensures the transformation system returns only a single pointer. This design is different from Spaun's, which relied on the repeated convolution of a single vector to indicate a count. While both approaches are viable, an associative memory is more effective for simpler, low-dimensional models.

The SP corresponding to each digit is taken from a random 10-dimensional orthonormal basis. This ensures that no prior information is given to bias the relationship between the 10 digits. Initially, the model only knows how to increment a digit, as is the situation for learning addition during childhood using solely the counting strategy.

Learning will occur in a heteroassociative memory that is identical to the aforementioned predefined memory, except its associations are learned with experience. Consequently, the action selection mechanism will not need to iterate through its associations to compute the desired result.

5. Memorization via reinforcement learning

The counting circuit that employs the predefined memory is slow as information is routed back and forth at the rate of subvocal rehearsal. We thus refer to the counting portion of the model as the "Slow-Net." In the model, other cortical areas consolidate the function of the Slow-Net by memorizing its eventual responses. The portion of the network responsible for the storage and subsequent retrieval of learned responses will be referred to as the "Fast-Net."

The purpose of the Fast-Net is to learn to associate the two digits provided as input to the Slow-Net with the Slow-Net's eventual response. This heteroassociative memory computes a discontinuous high-dimensional function with respect to the vectors that represent each digit. Specifically, it computes the function mapping from the input of two addend vectors concatenated together to the output vector of the answer. Knight, Voelker, Mundy, Eliasmith, and Furber (2016) have shown that by combining the supervised prescribed error sensitivity (PES; MacNeil & Eliasmith, 2011) learning rule with the unsupervised Vector Oja (Voja; Voelker, Crawford, & Eliasmith, 2014) learning rule, we can scalably and efficiently learn such complex functions in the NEF. The Fast-Net leverages the unsupervised learning rule to efficiently represent the incoming addends and uses the supervised learning rule to learn the correct sum from the Slow-Net.

Specifically, the PES learning rule minimizes the difference between the output of a neural population and its desired value, by adjusting its decoders ($\mathbf{d}_i$ from Eq. (2)) in response to an error signal r and given a learning rate κ:

$$\Delta \mathbf{d}_i = \kappa r a_i \tag{6}$$

However, for discontinuous high-dimensional functions such as the desired heteroassociative memory, this supervised learning rule is insufficient, as shown in Fig. 2. This is because a neuron will often fire in response to multiple inputs, in which case its decoder will be adjusted to completely overwrite its previous association (the standard "catastrophic forgetting" problem). To avoid this, a neuron should only fire for a single input, which is achievable by selecting the encoders ($\mathbf{e}_i$ from Eq. (1)) to be equal to the SPs for each of the digits and defining the neuron's threshold ahead of time, as is done in the predefined heteroassociative memory (Stewart, Tang, & Eliasmith, 2011). However, this would assume that the area of cortex designated to learn the addition task is already aware of the possible SPs within the Slow-Net.

To keep our approach general, instead of manually specifying these encoders, we use the Voja learning rule to form a sparse encoding of the possible inputs as they are presented to the Fast-Net. This is achieved by adjusting the encoders of any active neurons to become selective to only the current input. This prevents the catastrophic forgetting demonstrated in Fig. 2, in turn allowing PES to learn the correct output without overwriting past associations. This has been demonstrated by Knight et al. (2016) to scalably recall over 2,000 associations using 50 neurons per association in simulation, and over 190,000 associations in theory. Given a learning rate η and an input $\mathbf{x}$, the encoder $\mathbf{e}$ of neuron j are adjusted according to Voja as follows:

$$\Delta \mathbf{e}_j = \eta a_j (\mathbf{x} - \mathbf{e}_j). \tag{7}$$

To ensure the learned heteroassociative memory is learned correctly, the sparsity of the population is determined by setting the thresholds of the neural tuning curves. Here, we set the thresholds according to the maximum dot product between distinct inputs. This procedure is described more generally by Knight et al. (2016).

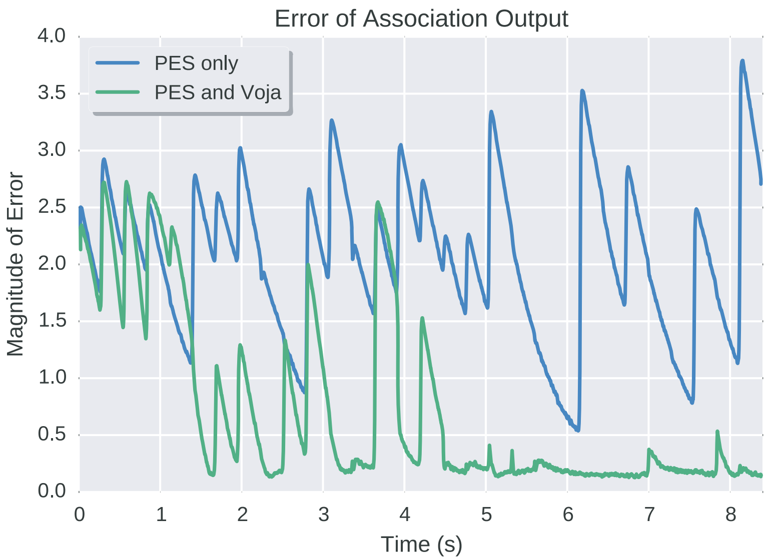

Fig. 2. Comparing two methods of learning heteroassociative memories: supervised learning only (PES only) versus both supervised and unsupervised learning (PES and Voja). An association between a 20-dimensional input and a 10-dimensional output is learned every 300 ms. We plot the error between the desired output and the learned association. The "PES only" approach suffers from catastrophic forgetting, which demonstrates the necessity of the Voja unsupervised learning rule. Simulations were run using Nengo 2.1 (Bekolay et al., 2014). See Fig. S1 for a decoder level of view.

In our model, the population of neurons in the Fast-Net memory continually learns to represent the input digits by adjusting its encoders via Eq. (7), while its decoders adjust to associate its input with the output of the Slow-Net via Eq. (6). To enable this simultaneous learning and operation, the Fast-Net is placed in parallel to the Slow-Net, as shown in Fig. 3. Both networks receive the same input, but the answer from the Slow-Net is projected to modulate the output of the Fast-Net. This modulatory error signal is a dopaminergic error signal that is sent whenever the Slow-Net responds with an answer, which then provides the feedback for learning via PES. Importantly, the network itself controls this error signal, using the same mechanisms as are used for controlling the steps of the counting process in the Slow-Net. This internal control of the dopaminergic error signal can be thought of as a type of metalearning (Doya, 2002) or controlling how to learn. In addition, such feedback could also come from the environment (e.g., in the form of a teacher correcting the student who is drilling addition facts), but this extension is outside the scope of this study.

6. Results

The learning rate of the heteroassociative memory can be adjusted to model developmentally plausible learning. At high learning rates the model learns mappings after being shown only a single example and at lower learning rates it gains confidence gradually,

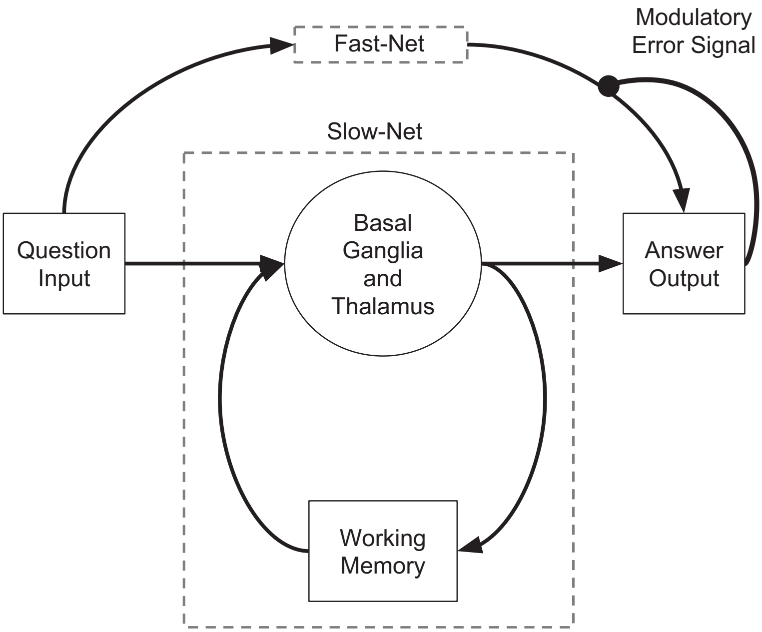

Fig. 3. High-level model architecture, featuring the parallel Slow-Net and Fast-Net. The input is provided to both networks simultaneously. The Fast-Net learns its responses via a modulatory error signal projected from the output of the Slow-Net.

covering the spectrum of human variability and demonstrating the versatility of the heteroassociative memory. In the following simulations we chose a lower learning rate to imitate the gradual accumulation of experience typical in development.

The model was built and simulated using Nengo 2.1 (Bekolay et al., 2014), whereas the results were plotted using Seaborn 0.7.1 (Waskom et al., 2016). The results of the Slow-Net, which implement the counting strategy, are shown in Fig. 4.

As expected, the magnitude of error between the answer from Fast-Net and the correct answer decreases with rehearsal, as shown in Fig. 5. Each epoch of training consists of 20 randomized example additions with no repetition. Given the fast learning rate of the model, after each epoch there is a significant drop in error (during each subsequent epoch, the model is seeing problems it has encountered before).

Which network drives the overall response is determined by a separate basal ganglia control loop contained in the Answer Output module of Fig. 3. After the magnitude of error between the Fast-Net output and any numerical output drops below the arbitrarily set certainty threshold of 0.5, the Fast-Net response will drive the model's response instead of Slow-Net, as it is confident in its answer. Any decrease in error magnitude past the certainty threshold reflects an increase in the certainty of the answer.

In addition, once the Fast-Net becomes confident enough in its responses to override the Slow-Net, the speed of responses becomes faster and more uniform, as shown in

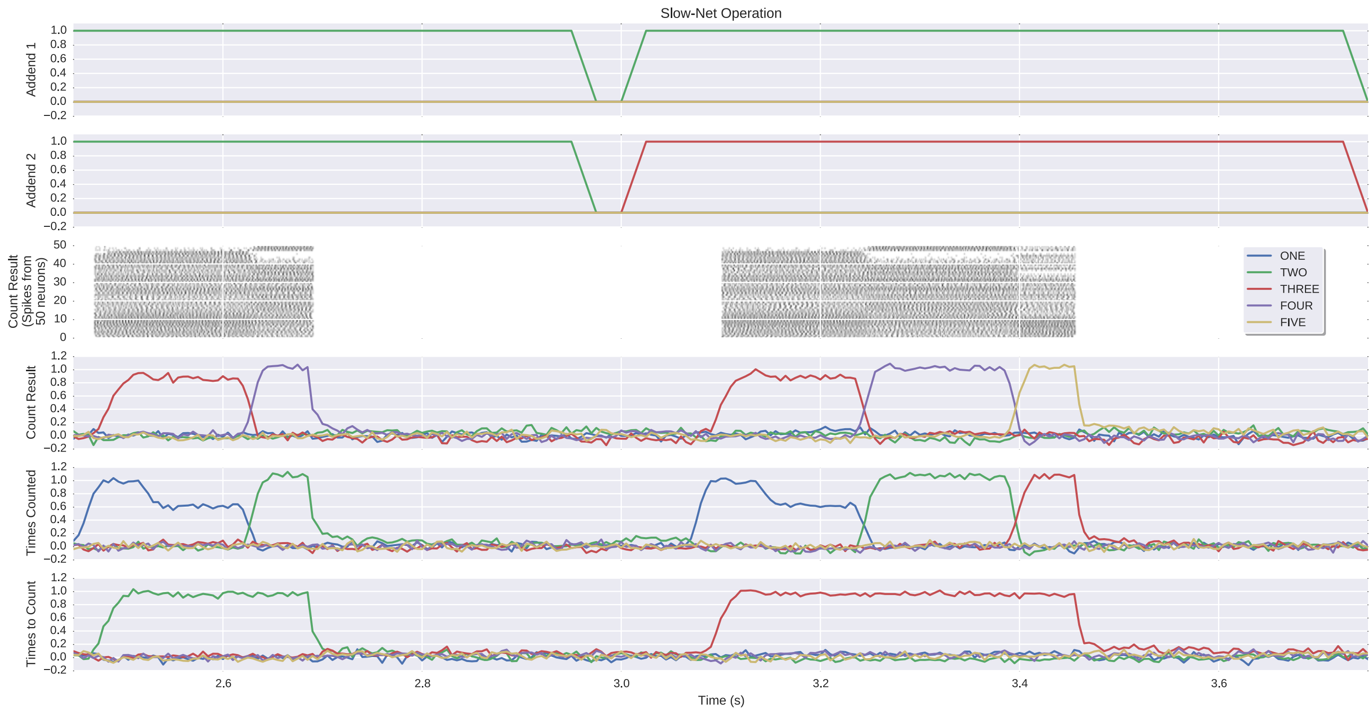

Fig. 4. The Slow-Net (counting network) answering the questions 2 + 2 and 1 + 3. The plots show similarity of neural activity to Semantic Pointers over time (colors indicate which pointer). As shown in “Count Result” the model is progressing through each intermediate digit before reaching the answer. Lines are the similarity between the spiking pattern in the area and the ideal spiking pattern for each number in the numerical vocabulary.

Fig. 5. Error magnitude in the Fast-Net decreasing with training received from the feedback of each trial.

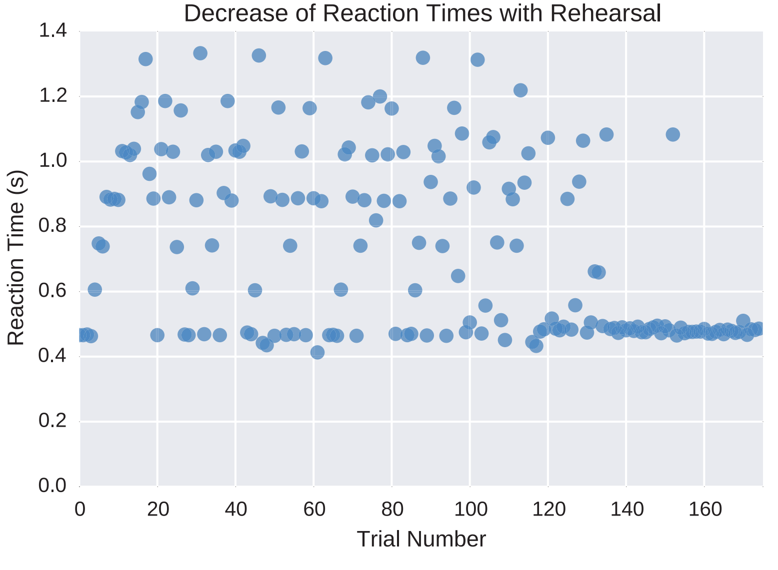

Fig. 6. Reaction times decreasing with rehearsal as the Fast-Net takes over for the Slow-Net for increasingly more additions. Note that these reaction times do not take into account motor planning for communicating the result and are thus much faster.

Fig. 6. The fact that reaction times regularize with transition to a recall-based strategy matches experiments investigating addition strategies (Siegler, 1987). In over 100 trials, the Slow-Net only failed three times to produce a correct answer and instead overcounted.

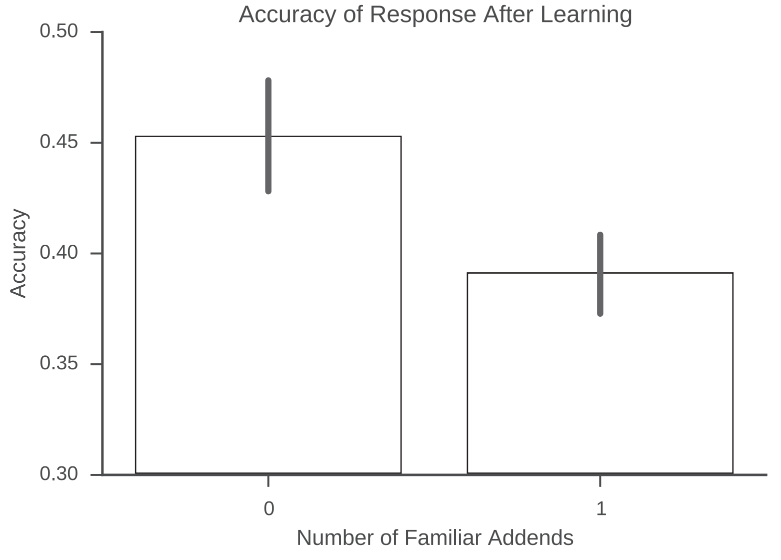

Fig. 7. Heteroassociative memory accuracy (similarity between memory output and desired output) given various addends and after training. Ticks on bars represent bootstrapped 95% confidence intervals.

These three failures are omitted from Fig. 6, as they are considered as outliers which could either be corrected by tweaking the model further or by implementing an introspective error monitoring mechanism.

To gauge the effects of previous experience on the learning process, we also ran an experiment where the Fast-Net was simulated separately and trained on a set of addition facts that only used the set 1, 2, 5. These facts were repeated until convergence. After this initial training, we presented new addition problems where either none or one of the addends belonged to the set 1, 2, 5.

Our model predicts proactive interference from previously learned addition facts with familiar addends. As shown in Fig. 7, the new problems with unfamiliar addends were learned more accurately. This is a direct consequence of the SP representation chosen for the addends. In particular, as we chose to concatenate the two digits to form the inputs to the heteroassociative memory, the SPs with common addends have a higher dot-product similarity than those without any common addends. This results in fewer encoders being available for these SPs after Voja has converged. Binding the two addends together, as in a Holographic Reduced Representation, would result in statistically different SP similarities. Psychological experiments involving patterns of interference during development would help infer which representation is more plausible.

7. Neuroanatomical mappings and dyscalculia

As discussed in Spaun's mapping of counting (Eliasmith et al., 2012), parietal areas are more active for stable, learned transformations, while prefrontal areas are more active for transient, working memory representations. Given that the Fast-Net responses eventually replace those from the Slow-Net, we would expect that with practice brain activity will move from the prefrontal cortex to frontal-parietal areas. Although this is quite difficult to measure during the natural development of a child, this hypothesis is supported by a few observations.

First, when doing mental calculation, age is correlated with the use of parietal brain areas. However, it is inversely correlated with the use of prefrontal and hippocampal brain areas, as well as the use of the dorsal basal ganglia area (Rivera, Reiss, Eckert, & Menon, 2005). Within the model, this transition can be framed as older children using fewer iterative processes (no use of working memory in prefrontal, no iterative process control requiring the basal ganglia, and no loading of instructions from the hippocampus) and more memorization.

Second, those with dyscalculia show greater activation of the prefrontal cortex compared to individuals with normal numeracy (Kucian & von Aster, 2015). Although this model make no claims about why dyscalculia occurs, given that it is a complicated disability usually accompanied by various comorbidities and no direct cause (Rubinsten & Henik, 2009), it does provide part of a possible explanation as to why such compensation occurs. Specifically, those with dyscalculia are unable to consolidate the functional role of the prefrontal cortex during the counting task within the frontal-parietal region and must instead rely on their working memory. Given an excessively noisy input, inaccurate feedback, or inappropriate modulation of the error signal, the Fast-Net could fail to learn the mapping between addends and sum. Consequently, there would be limited progression from counting to recall and a continued dependence on working memory.

8. Discussion and conclusions

We presented a biologically plausible model that progresses from a methodical counting procedure to recalling the response for an addition task. As specified in Table 3 and shown in Fig. 5, the accuracy of the Fast-Net memory progresses from noisy to accurate. Once a sufficient accuracy threshold is reached, the memory takes over the process of addition from the Slow-Net, increasing the reaction times, as specified in Table 2 and shown in Fig. 6.

The model as presented is clearly limited in several respects. For instance, numerical representation in the brain consists of more structure than the randomly selected orthonormal vectors used here. This assumption is reasonable for small magnitudes, but it is untenable for numerical representation in general. For instance, there is evidence that neurons tuned to numerical size comparisons are proportional to a log scale and can be highly sensitive to task saliency (Nieder & Dehaene, 2009). Capturing such properties will require different numerical representations.

Although the model proposes one possible mechanism for parallel learning, additional mechanisms are required in more complex situations, such as chunking sequential motor commands (Ramkumar et al., 2016). Many frameworks have been proposed for accomplishing such tasks (Verwey, Shea, & Wright, 2015); however, we are unaware of any that use spiking neurons. One potential mechanism is a neural stack-like memory structure that produces compressed representations of action plans (Arsalidou & Taylor, 2011). The resulting SPs could be used within our model to learn associations between action plans.

Finally, this model only describes a handful of the brain areas associated with numerical calculation. One of the most surprising areas involved in addition is the cerebellum. It has been suggested that cerebellar activity might be a developmental artifact persisting from when addition is first learned as a physical combination of objects (Blouw, Eliasmith, & Tripp, 2016). To model cerebellar involvement, counting and grouping objects would need to be rehearsed via explicit motor plans. This could be captured with a more in-depth model that includes a motor control hierarchy and a visual system similar to those found in Spaun.

Acknowledgments

This work was supported by CFI and OIT infrastructure funding as well as the Canada Research Chairs program, NSERC Discovery grant 261453, AFOSR grant FA8655-13-1-3084, and NSERC graduate funding. The authors are especially grateful to Terry Stewart for his mentorship throughout this work.

References

Arsalidou, M., & Taylor, M. J. (2011). Is 2+ 2= 4? Meta-analyses of brain areas needed for numbers and calculations. *Neuroimage*, *54*(3), 2382–2393.

Bekolay, T., Bergstra, J., Hunsberger, E., DeWolf, T., Stewart, T. C., Rasmussen, D., X., Choo, A. R., Voelker & Eliasmith, C. (2014). Nengo: A Python tool for building large-scale functional brain models. *Frontiers in Neuroinformatics*, *7*(48). doi: 10.3389/fninf.2013.00048

Blouw, P., Eliasmith, C., & Tripp, B. (2016). A scaleable spiking neural model of action planning. In A. P. Dan Grodner Dan Mirman & J. Trueswell (Eds.), *Proceedings of the 38th Annual Conference of the Cognitive Science Society* (pp. 1583–1588). Philadelphia, PA: Cognitive Science Society.

DeWolf, T., & Eliasmith, C. (2013). A neural model of the development of expertise. In R. L. West & T. C. Stewart (Eds.), *12th international conference on cognitive modelling* (pp. 119–124). Ottawa, ON: Carleton University.

Doya, K. (2002). Metalearning and neuromodulation. *Neural Networks*, *15*(4), 495–506.

Eliasmith, C. (2013). *How to build a brain: A neural architecture for biological cognition*. New York: Oxford University Press.

Eliasmith, C., & Anderson, C. H. (2003). *Neural engineering: Computation, representation, and dynamics in neurobiological systems*. Cambridge, MA: MIT Press.

Eliasmith, C., Stewart, T. C., Choo, X., Bekolay, T., DeWolf, T., Tang, Y., & Rasmussen, D. (2012). A large-scale model of the functioning brain. *Science*, *338*, 1202–1205.

Gayler, R. W. (2004). Vector symbolic architectures answer Jackendoffs challenges for cognitive neuroscience. *arXiv preprint cs/0412059*.

Knight, J., Voelker, A. R., Mundy, A., Eliasmith, C., & Furber, S. (2016). Efficient SpiNNaker simulation of a heteroassociative memory using the Neural Engineering Framework. In *The 2016 international joint conference on neural networks (IJCNN)*. British Columbia: Vancouver.

Kucian, K., & von Aster, M. (2015). Developmental dyscalculia. *European Journal of Pediatrics*, *174*(1), 1–13.

Lebiere, C. (1999). The dynamics of cognition: An ACT-R model of cognitive arithmetic. *Kognitionswissenschaft*, *8*(1), 5–19.

MacNeil, D., & Eliasmith, C. (2011). Fine-tuning and the stability of recurrent neural networks. *PloS ONE*, *6*(9), e22885.

Nieder, A. (2012). Supramodal numerosity selectivity of neurons in primate prefrontal and posterior parietal cortices. *Proceedings of the National Academy of Sciences*, *109*(29), 11860–11865.

Nieder, A., & Dehaene, S. (2009). Representation of number in the brain. *Annual Review of Neuroscience*, *32*, 185–208.

Plate, T. A. (1995). Holographic reduced representations. *IEEE Transactions on Neural Networks*, *6*(3), 623–641.

Ramkumar, P., Acuna, D. E., Berniker, M., Grafton, S. T., Turner, R. S., & Kording, K. P. (2016). Chunking as the result of an efficiency computation trade-off. *Nature Communications*, *7*.

Rivera, S. M., Reiss, A., Eckert, M. A., & Menon, V. (2005). Developmental changes in mental arithmetic: Evidence for increased functional specialization in the left inferior parietal cortex. *Cerebral Cortex*, *15* (11), 1779–1790.

Rubinsten, O., & Henik, A. (2009). Developmental dyscalculia: Heterogeneity might not mean different mechanisms. *Trends in Cognitive Sciences*, *13*(2), 92–99.

Sarnecka, B. W., & Carey, S. (2008). How counting represents number: What children must learn and when they learn it. *Cognition*, *108*(3), 662–674.

Siegler, R. S. (1987). The perils of averaging data over strategies: An example from childrens addition. *Journal of Experimental Psychology: General*, *116*(3), 250.

Stewart, T. C., Choo, X., & Eliasmith, C. (2010). Symbolic reasoning in spiking neurons: A model of the cortex/basal ganglia/thalamus loop. In S. Ohlsson & R. Catrambone (Eds.), *32nd Annual Conference of the Cognitive Science Society*. Portland, OR: Cognitive Science Society.

Stewart, T. C., Tang, Y., & Eliasmith, C. (2011). A biologically realistic cleanup memory: Autoassociation in spiking neurons. *Cognitive Systems Research*, *12*, 84–92. doi:10.1016/j.cogsys.2010.06.006

Verwey, W. B., Shea, C. H., & Wright, D. L. (2015). A cognitive framework for explaining serial processing and sequence execution strategies. *Psychonomic Bulletin & Review*, *22*(1), 54–77.

Voelker, A. R., Crawford, E., & Eliasmith, C. (2014). Learning large-scale heteroassociative memories in spiking neurons. In O. H. Ibarra, L. Kari & S. Kopecki (Eds.), *13th international conference on unconventional computation and natural computation*. London, ON: Springer.

Waskom, M., Botvinnik, O., Drewokane, Hobson, P., Halchenko, Y., Lukauskas, S., ... Rocher, L. (2016). *seaborn: v0.7.0 (january 2016)*. doi:10.5281/zenodo.45133

Supporting Information

Additional Supporting Information may be found online in the supporting information tab for this article:

Appendix S1. Source code and decoder analysis.