Learning large-scale heteroassociative memories in spiking neurons

University of Waterloo

CNRG computational neuroscience research group
compneuro.uwaterloo.ca

Aaron Russell Voelker, Eric Crawford, Chris Eliasmith {arvoelke, e2crawfo, celiasmith}@uwaterloo.ca
Centre for Theoretical Neuroscience, University of Waterloo <http://ctn.uwaterloo.ca>

Introduction

A heteroassociative memory is a collection of associations between keys (input concepts) and values (associated concepts).

The hippocampus likely performs the function of a heteroassociative memory (Treves and Rolls, 1994). Most cognitive architectures require the ability to store and recall associations.

Crawford et al. (2013) constructed a scalable heteroassociative memory by specifying the connection weights in a feedforward network of spiking neurons.

We demonstrate how to learn these connection weights online.

Approach

Concepts are taken as vectors, x, in a high-dimensional space.

In the Neural Engineering Framework (NEF), each neuron has an encoding vector that maps x into the current driving the neural nonlinearity $G[\cdot]$.

$$a_i = G[\alpha_i \mathbf{e}_i \cdot \mathbf{x} + J_i^{bias}]$$

For each neuron, we find a decoding vector that maps the spiking back to a function of x.

$$\hat{f}(\mathbf{x}(t)) = \sum \mathbf{d}_i^f a_i(t)$$

Connection weights are the outer product between the encoding and decoding vectors.

$$\mathbf{W} = \mathbf{e}\mathbf{d}^T$$

If we choose a neuron's encoding vector to be a particular key, it will fire selectively for that key. The decoding vector for this neuron is optimized to estimate the value associated with its key.

Our novel learning rule modifies the encoding vectors of active neurons to be selective to an input vector. This tunes the connection weights so that a small number of distinct neurons respond to each key.

$$\Delta \mathbf{e} = \eta a(t) \left(\begin{pmatrix} \mathbf{x}^T \\ \vdots \\ \mathbf{x}^T \end{pmatrix} - \mathbf{e} \right)$$

The PES learning rule (Bekolay et al., 2013) updates the decoding vectors of these neurons to estimate the current value.

Scalability

The number of neurons grows linearly with the number of associations.

This approach is scalable to over 100,000 concepts, the size of an adult human's vocabulary, while using only approximately 14.7 mm^2 of cortex (Crawford et al., 2013).

Model Architecture

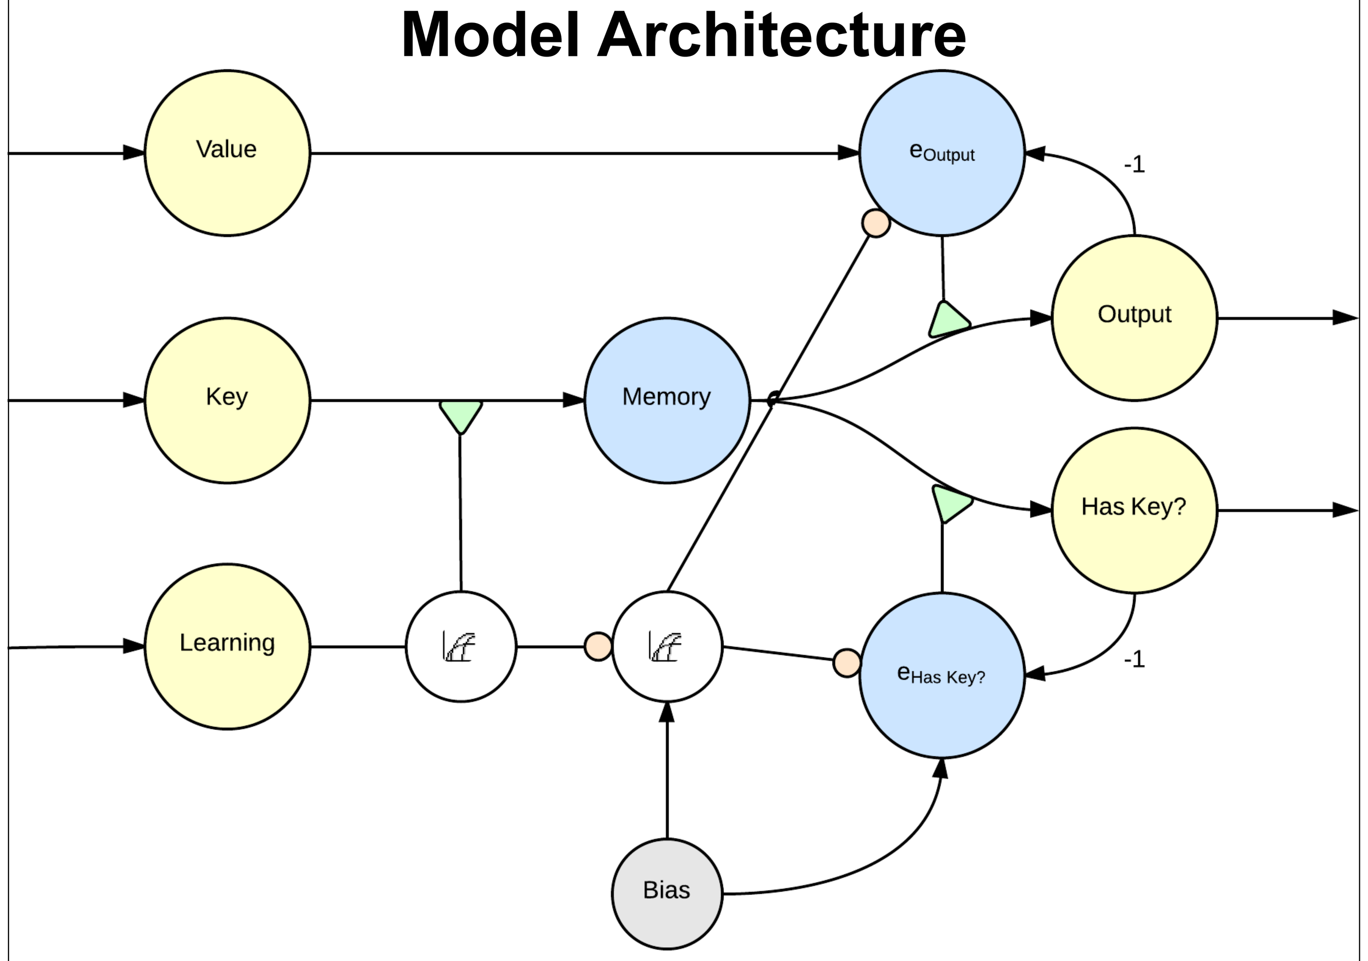

Yellow circles denote inputs and outputs, blue circles are NEF ensembles, white circles are scalar [0, 1] ensembles with positive intercepts and encoders, orange connections are inhibitory, and green connections are modulatory.

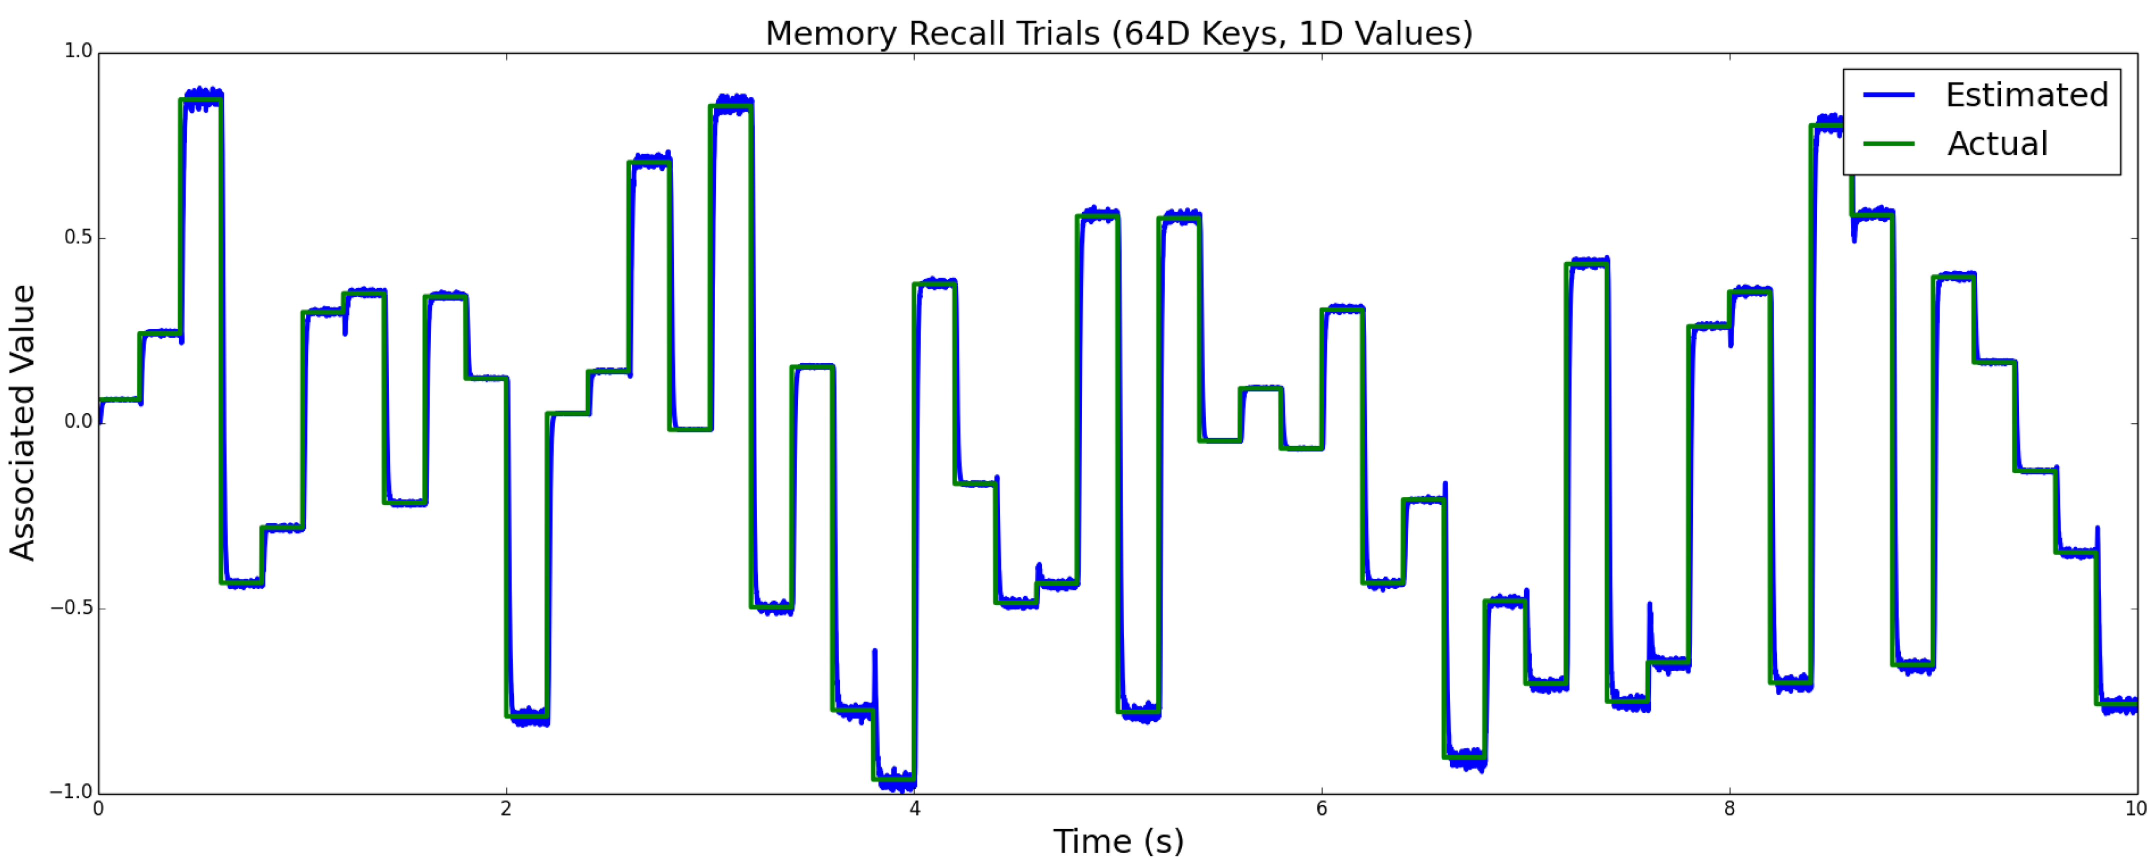

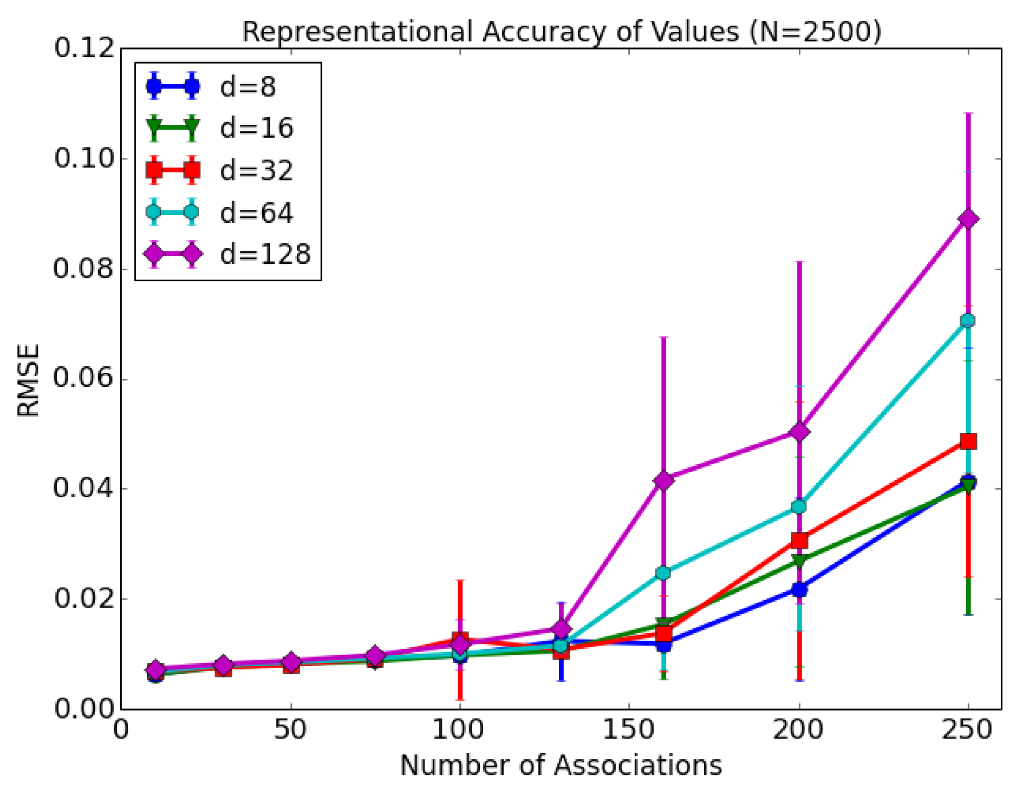

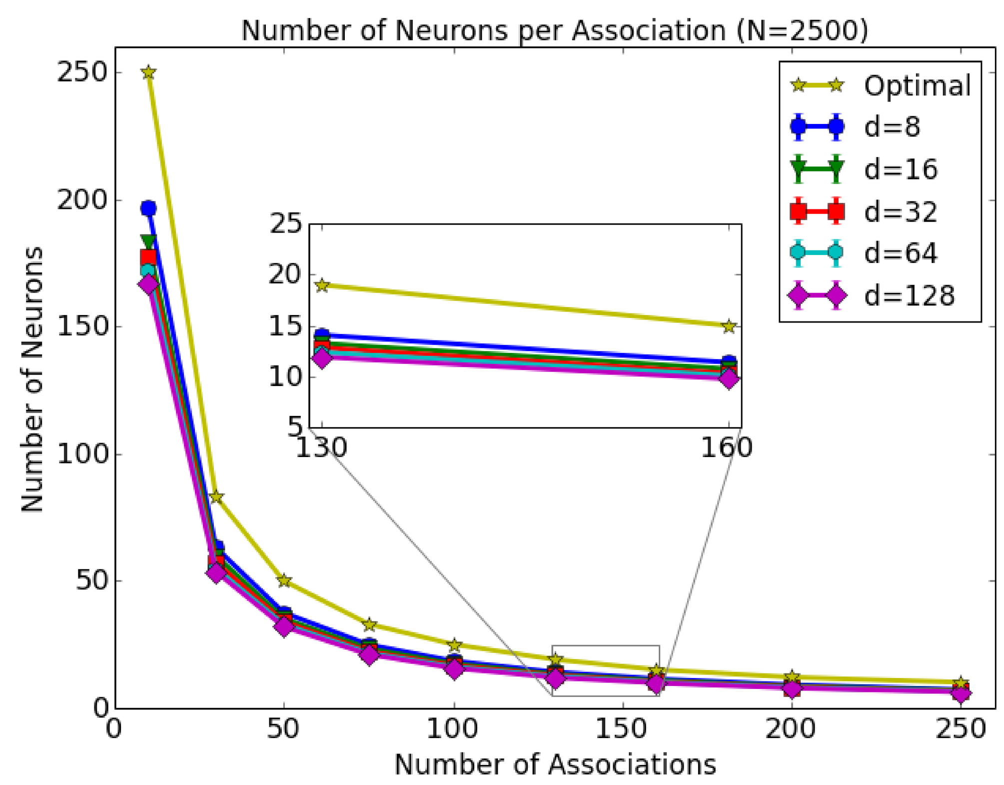

Encoder Self-Organization

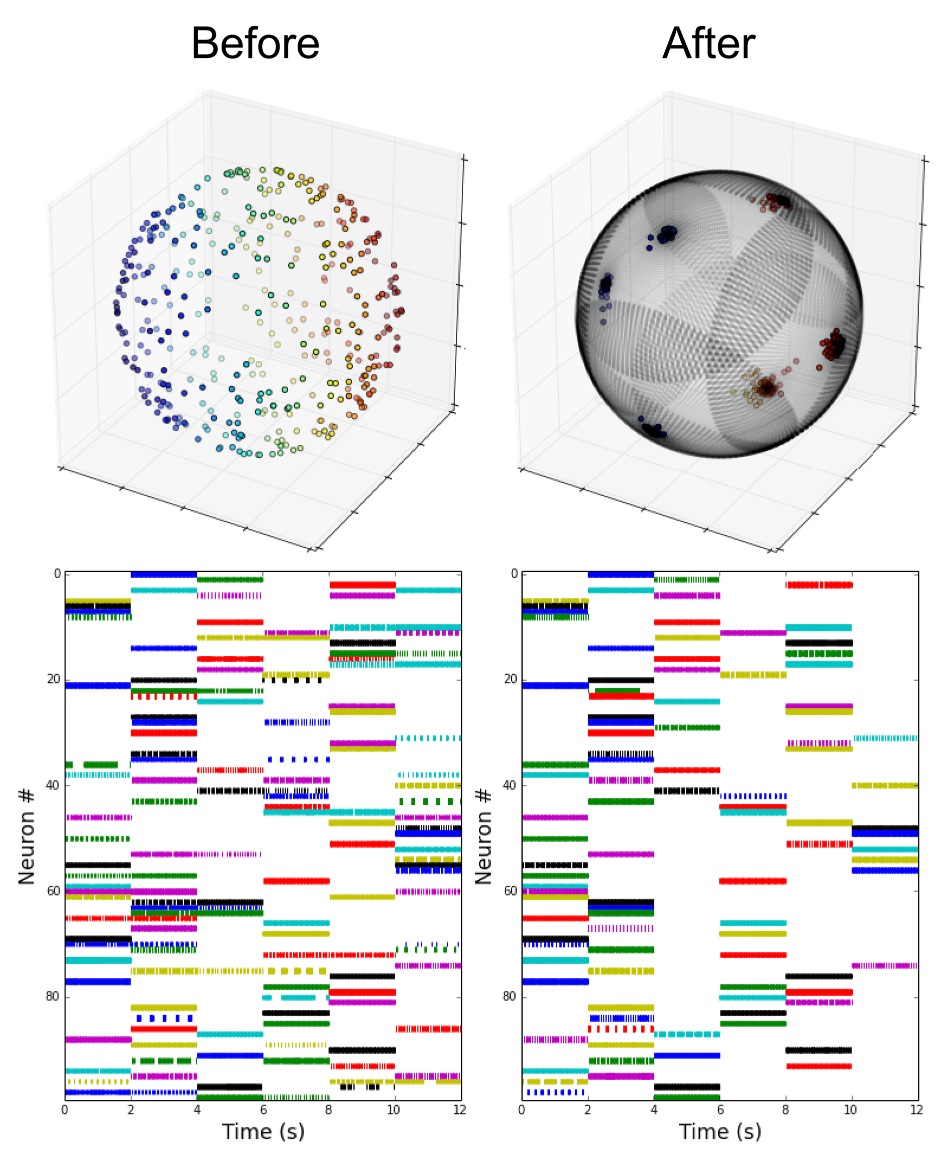

Results

The memory is tested by presenting each key to the network for a small interval of time, and decoding estimates of the associated values.

With sufficiently dissimilar keys, the RMSE is around 0.01 when allocated at least twenty neurons per item.

This error is mostly due to noise.

Conclusion

Using the Neural Engineering Framework, we can encode associations as pairs of high-dimensional vectors represented by neural activations in spiking neurons.

These associations can be stored in the connection weights between populations of neurons, which are learned online using biologically plausible learning rules.

This feedforward network can recall an association with high accuracy in just a few milliseconds, and scales efficiently to the size of vocabularies observed in human adults.