

Vector Symbolic Spiking Neural Network Model of Hippocampal Subarea CA1 Novelty Detection Functionality

Claus Agerskov

clausagerskov@hotmail.com

Department of Physics, Technical University of Denmark, Kongens Lyngby, 2800, Denmark

A neural network model is presented of novelty detection in the CA1 subdomain of the hippocampal formation from the perspective of information flow. This computational model is restricted on several levels by both anatomical information about hippocampal circuitry and behavioral data from studies done in rats. Several studies report that the CA1 area broadcasts a generalized novelty signal in response to changes in the environment. Using the neural engineering framework developed by Eliasmith et al., a spiking neural network architecture is created that is able to compare high-dimensional vectors, symbolizing semantic information, according to the semantic pointer hypothesis. This model then computes the similarity between the vectors, as both direct inputs and a recalled memory from a long-term memory network by performing the dot-product operation in a novelty neural network architecture. The developed CA1 model agrees with available neuroanatomical data, as well as the presented behavioral data, and so it is a biologically realistic model of novelty detection in the hippocampus, which can provide a feasible explanation for experimentally observed dynamics.

1 Introduction ---

Memory loss is a prevalent problem in modern society, and this condition is expected to increase in severity because of a number of societal changes. Among these are increases in life expectancies, as well as the fact that the percentage of old people in the general population is on the rise (Alzheimer's Association, 2014).

In experimental neuroscience, the perspective from disease effects has been the major source of information about brain function. Afflictions such as Alzheimer's disease and traumatic brain injury, as well as other neurodegenerative diseases, can have crippling effects on an individual. A cure for these ailments is still unavailable. One major reason for this is our ignorance of how the human brain works on a very basic level. On one hand, we have to understand the complex interactions of different brain areas from a systems scale to the scale of synapse function. This is especially

true for Alzheimer's disease, as some of the first symptoms of this debilitating disease are the impairment of memory on the short-term scale, and its conversion into long-term memory, and then slowly eroding other brain functions, eventually leading to death. On the other hand, the only hope of improvement for people with structural damage to brain areas such as the medial temporal lobe (from war or accident) seems to be brain prosthetics that replicate the function of the damaged areas. Such prosthetics, however, require that we know about the circuits in the brain and their intricate functioning that give rise to the complex behavior associated with normal human life. To completely rehabilitate people who are unable to form new memories, we need to understand exactly how this memory system works in a healthy person.

With recent developments in computer science and neural network (NN) algorithms, it is now feasible to do larger-scale simulations of neural networks to emulate brain function. This is the focus of the newly emerging field of computational cognitive neuroscience, which seeks to draw links of causality between the scales of synaptic transmission and cognition. Furthermore, by having models of brain regions available, we can begin to make some of the many theories and concepts of cognitive psychology more concrete and provide testable predictions that can then be verified or falsified using the standard methods in the field of experimental neuroscience.

For the reasons stated above, there is considerable interest in studying the neural basis of memory. In this article, the focus is on exactly that: investigating the construction of functional neural network models of the hippocampal formation (HCF). Whereas the transferral of memory to long-term storage, hypothesized to be located in the neocortex, is very slow, the hippocampus seems to enable rapid storage of new information in a highly sparse and distributed form (O'Reilly & McClelland, 1994; O'Reilly, Bhattacharyya, Howard, & Ketz, 2014). An issue that could arise from this functionality is the redundant storage of the new information in the same short-term memory circuits. This very sparse representation of the hippocampus (HC) makes sure that two temporally separated but semantically similar sensory inputs do not overlap, which would result in catastrophic interference (see French, 1999, for a contemporary overview of the concept). What is needed is some signal proportional to the novelty of the sensory input. This signal could also serve as a gateway to long-term storage, a functionality attributed to the HC by Lisman and Grace (2005). The HC subdomain CA1, in particular, has been linked to such a novelty signal in several behavioral studies (Larkin, Lykken, Tye, Wickelgren, & Frank, 2014; Nitz & McNaughton, 2004). Specifically, Larkin et al. (2014) found that the CA1 area generates a generalized novelty signal irrespective of the nature of the novelty. Nitz and McNaughton (2004) additionally found that novelty is associated with a decrease in CA1 inhibitory interneuron activity. It is the focus of this article to demonstrate how such a novelty signal can be generated, using an artificial neural network model, restricted by anatomical and behavioral data, to ensure biological realism.

2 Methods

2.1 Semantic Representations and the NEF. The simulated neural networks all consist of point-like, leaky integrate-and-fire (LIF) neurons, and the model development is supported by the neural engineering Framework (NEF) (Eliasmith & Anderson, 2003), carried out in the Nengo NN simulation software (Bekolay et al., 2014). Furthermore, following the foundation of the NEF, the individual neurons are assigned preferred direction vectors, φ , so that the input current I_{in} to any neuron in a network is proportional to the dot product between this vector and an input vector x , representing sensory information:

$$I_{in} = \alpha(\varphi \cdot x) + I_{bias}, \quad (2.1)$$

where α is a conversion factor correcting the units and I_{bias} is the background current resulting from the activity of other nearby neurons. This is a phenomenological model of how neurons represent information in a hierarchical, highly distributed manner by virtue of their position in a complex network. This model was originally inspired by Georgopoulos, Schwartz, and Kettner (1986), who found that certain neurons in the macaque monkey brain encode two-dimensional movement directions.

Generalizing to high dimensions, this preferred direction vector can be viewed as representing abstract semantic information, which forms the basis of vector symbolic architectures (VSA). The input current to each neuron in a network results in a spiking rate $a(x)$. The neuron output presented in this study is the decoded, represented value computed internally in the simulation software by finding the associated least-squares linearly optimal decoding vectors, d_i , calculated from the following formula:

$$d = \Gamma^{-1}\gamma$$

where $\Gamma_{ij} = \int a_i(x)a_j(x)dx$

and $\gamma_i = \int a_j(x)xdx,$ (2.2)

where $a_i(x)$ and $a_j(x)$ is the activity of two neurons in the population being decoded. This decoded value is an estimate of the original value x (see Eliasmith & Anderson, 2003, for details). In the implemented VSA, the dot product between two high-dimensional vectors is seen as a measure of similarity: vectors with a smaller angle θ between them are more similar than distant vectors. From the definition of the dot-product, we have the following for two vectors A and B :

$$A \cdot B = \|A\| \|B\| \cos(\theta). \quad (2.3)$$

Having unit vectors, the dot-product gives a measure between 0 and 1, 0, meaning orthogonal vectors (i.e., no similarity), and 1, meaning identical information. The dot-product can also be negative, meaning inverse proportionality.

Operations on the sensory input, as well as on represented input (e.g., a linear transformation Wx), can be implemented in the connection weights W between neurons, explicitly derived in the NEF with the same decoding vectors, d_i using the equation

$$\omega_{ij} = \alpha_j \phi_j W d_i, \quad (2.4)$$

where ω_{ij} is the synaptic weight between neuron i and j and Wx represents a linear transformation on x . This principle can be generalized to nonlinear functions as well, with a change in γ , replacing x with $f(x)$, where $f(x)$ is the value of the function desired to be applied on x .

2.2 Long-Term Memory Network. To be able to determine what is novel, one must first establish what is known. In other words, a storage network of known semantic information (vectors) must be implemented. This can then be an input to the network calculating the novelty signal, together with the original sensory input, to allow comparison. The challenge of storing and retrieving information in long-term memory (LTM) in a biologically realistic autoassociative way is a longstanding unsolved problem in the science of neural networks. The basic problem is as follows: store m patterns of activity in a group of n mutually connected neurons in such a way that if the pattern m_i is presented to the network, it converges to that same stored pattern and returns it as an output.

The classic solution takes the form of the Hopfield network (Hopfield, 1982). This model of LTM is powerful, but it is limited to neurons with binary activation. Also, the network always converges to a stored pattern, no matter the input. As such, it is not suitable for our applications. Instead the long-term memory network architecture developed by Stewart, Tang, and Eliasmith (2011) and elaborated by Voelker, Crawford, and Eliasmith (2014) and Crawford, Gingerich, and Eliasmith (2013) is used to represent the known information. This model is also a VSA developed with the NEF. Specifically, it works by setting the preferred direction vectors of every neuron in a group to be a stored vector. In this study, 100 neurons are used per stored semantic vector, allowing for a memory network that is robust to cell death. Crucially for our purposes, the memory network responds with noise when presented with an unknown vector (i.e., one that is not stored in any of the groups). This is illustrated by comparing Figure 1 (left), showing the input to the storage, to Figure 1 (right), showing the resulting output of

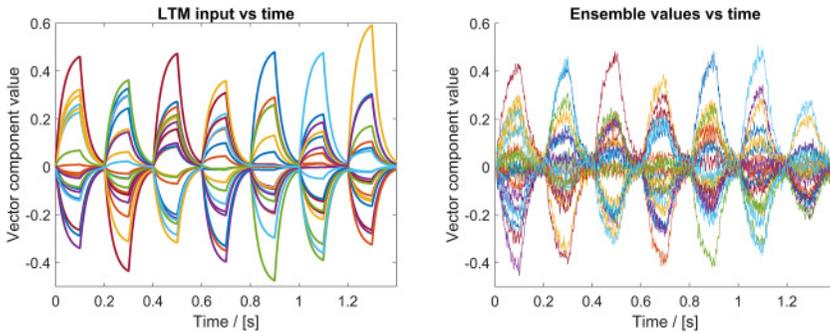


Figure 1: (Left) Input vector components of the 20-dimensional vectors representing the semantic input to the long-term memory (LTM) network. Each vector is presented for 0.1 second, followed by 0.1 second of zero input. The input is gradually rising and falling to each value according to the neuron time constants. (Right) Response of the LTM network to the presented inputs, revealing autoassociative dynamics for known inputs and noisy response for unknown inputs.

the memory network. The input node is at any time directly representing either a 20-dimensional vector or zero. Six vectors are stored separately in six different storage neuron groups. By having the input cycle through representing the six stored vectors as shown in Figure 1 (left), the memory responds with the same vectors as output. Note that in the graph, seven different represented vectors are shown, each presented to the network for 0.1 second, separated by 0.1 second of zero input. The seventh input is an unknown vector that is not stored in the memory and is used to probe the system's recognition capabilities. We see exactly the desired dynamics of noisy response to unknown input.

3 Results

3.1 Higher-Dimensional Novelty Detection. To detect the degree of novelty (as opposed to the noisy output of the LTM network) in systems carrying semantic information in a VSA, we need to compare some stored information to a high-dimensional input, external to the simulated brain region, and produce a signal that is proportional to the similarity between those two pieces of information. Eliasmith (2013) has noted that one can use the intrinsic, nonlinear properties of hippocampal pyramidal neurons to create what can be called a cleanup memory. A cleanup memory that can quickly and accurately pick out a noise-reduced output, out of a storage, with maximum similarity to some input, was developed by Stewart et al. (2011) and demonstrated in use in the ordinal serial encoding (OSE) model

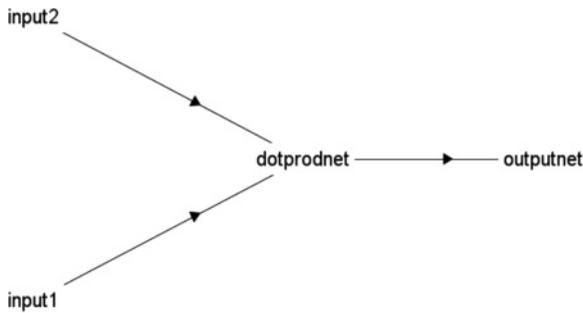


Figure 2: Simple multidimensional novelty detection system, without a storage network. Two inputs project to a neuron group. The weights have been adjusted according to the NEF to compute the dot-product between the two inputs.

of serial memory by Choo (2010). The issue with this cleanup memory for the current purposes is that it has programmatically explicit access to the full vocabulary of possible vectors and directly compares similarity to the input. Instead, we are interested in two external outputs, simulating a projection from, for example, a high-level sensory cortex and a memory network. Essentially we need to make a network compute the dot-product between our LTM output and the simulated input. Figure 2 shows such an implemented network that is attempting to perform dot-product operation. Two input nodes connect to an ensemble with a dimensionality of the number of inputs times the dimensionality of the inputs (i.e., for 5D inputs, $2 \times 5 = 10$). This ensemble then projects to another ensemble of neurons that represents the scalar recognition signal. It is by adjusting the weight matrix between the dotprodnet ensemble and the outputnet ensemble that one is able to compute the dot-product function on the two inputs.

To avoid time-consuming training of the network, the weights are derived using the NEF. This has been done in Nengo, using its implementation of the NEF. The inputs and outputs of the network are shown in Figures 3 and 4, respectively.

As shown in the two graphs of Figure 3, the two inputs each have two different 5D vectors of unity length represented. These inputs are alternating between similar (dot-product around 0.9) and dissimilar (dot-product below 0.15) pairs of vectors. The graphs in Figure 4 show the recognition signal computed by connections between the dotprodnet and the outputnet ensembles for 300 (blue graph) and 1000 (green graph) neuron ensembles with radius 2, in the dotprodnet of Figure 2. The red graph is the directly computed dot-product between the two inputs.

As desired, the recognition signal is higher when the two inputs are similar compared to when the inputs are dissimilar. If we compare the graph for 300 neurons to the direct calculation, however, we see a deviation

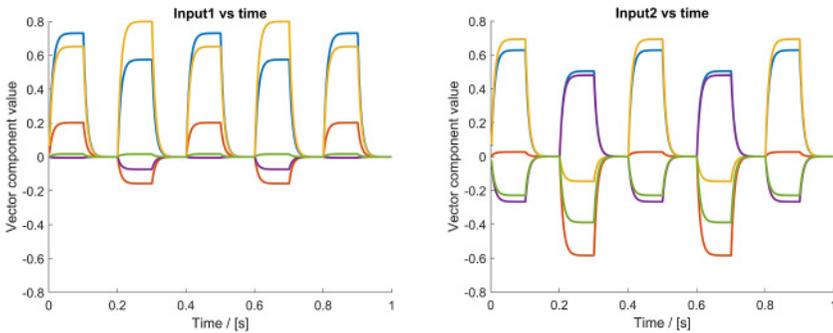


Figure 3: Inputs to the simple novelty detection network showing alternating dissimilar and similar vectors.

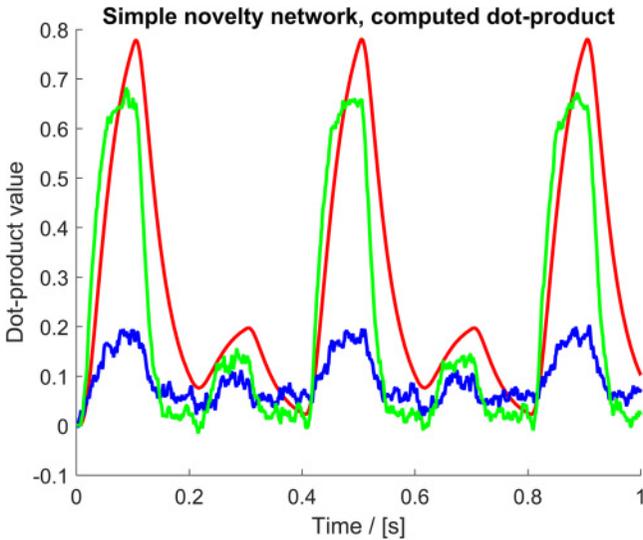


Figure 4: Recognition signal, or calculated dot-product, using 300 (blue) and 1000 (green) neurons to compute the operation. Direct calculation is shown in red. There is still some disagreement between the neuronal and direct computation, even when tripling the number of neurons, so it is concluded that the dot-product function is not performed accurately.

between the two graphs, indicating that the network weights have not fully implemented the dot-product operation. Allocating more than three times as many neurons to the task does not improve the performance to completely agree with the direct calculation. It seems that the one-ensemble

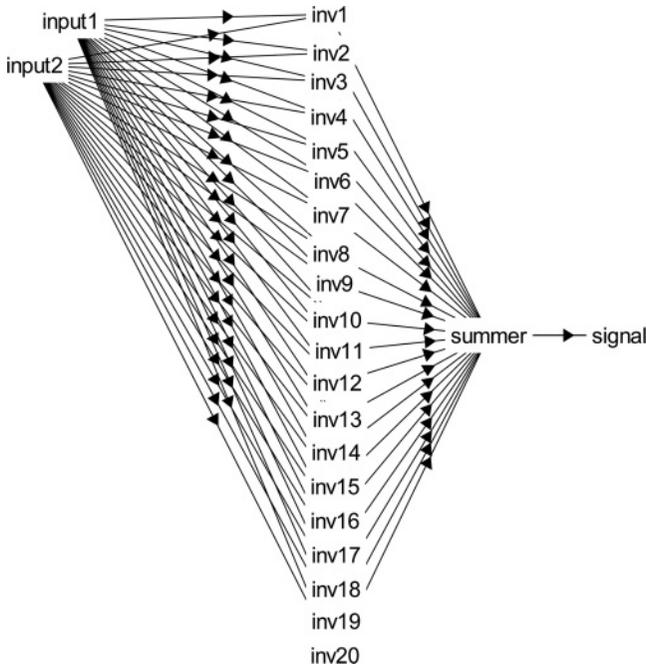


Figure 5: Pairwise product and summation (PPnS) implementation of the novelty detection network. Each input now projects its vector components pairwise to the corresponding group of neurons (inv1–inv20) (each 150 neurons) that multiplies them together using weight adjustments according to the NEF. The next layer (the summer ensemble of 500 neurons) then adds all the components together, using neural integration, and the computed novelty signal is represented by the signal group (500 neurons). In total, the implementation uses 4000 neurons.

implementation of the dot-product operation is not an economical solution to the problem at hand.

The most obvious way of reducing the complexity of the operation would be to divide it into two stages: the multiplication stage and the summation stage. By having individual smaller groups multiply each component of the input vectors pairwise and then summing all of these results together, we can perform the same operation, but in a way that should be much easier to implement in a network. Such a pairwise product and summation (PPnS) network implementation is shown in Figure 5 for 20D inputs. As in the simple novelty implementation shown in Figure 2, the PPnS network takes two inputs, now each 20D, and computes the dot-product between them. In this network, however, we first split the two inputs into individual vector components. These then project pairwise to the invi ensembles,

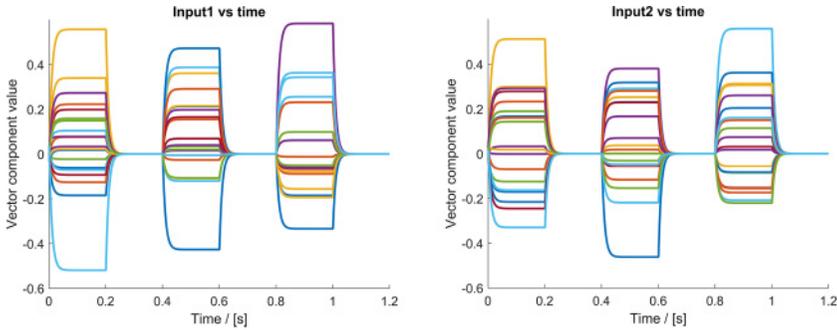


Figure 6: 20D vector inputs for the PPnS network.

each 150 neurons, for the i th input vector components of input 1 and 2. The multiplication part of the dot-product is then implemented in the weights between the *invi* and summer ensemble, with 500 neurons. Simultaneously the products are integrated into a 1D value (the summation part of the dot-product operation) by virtue of the intrinsic neuron dynamics of integrating dendritic inputs to the summer ensemble. The final signal group then represents the computed signal, using 500 neurons. In total the PPnS 20D network uses 4000 spiking neurons.

A test of this network's capabilities is done by having the two inputs represent alternately similar and dissimilar inputs, as done before with the one group novelty network. The vectors represented by the two inputs are shown in Figure 6. The resulting recognition signal (dot-product) is shown in Figure 7.

The two curves in the signal graph represent the recognition signal computed using spiking neurons (black curve) and direct calculation of the dot-product (red curve). By comparing the curves, we can see that there is excellent agreement of the signal magnitude between the two, meaning an improvement of performance compared to the one-neuron group architecture, even with a 15D increase in input.

To further investigate the difference in performance of the two dot-product implementations, a series of accuracy tests is run for different dimensionalities (and consequently number of neurons) and optimization radii. For both the one-ensemble (OE) and PPnS implementations, 5D and 10D dot-products are tested, using 1000 and 2000 neurons, respectively. For the OE, a 15D test is done, and for the PPnS, a 20D. For both the OE 15D and PPnS 20D tests, the number of neurons is set to 4000. Each network is constructed using a range of optimization radii and its performance evaluated on 100 runs with different, randomly generated input vectors, of length 1, and components between -1 and 1 . In each run, the performance of the network is compared to a direct calculation of the dot-product between two

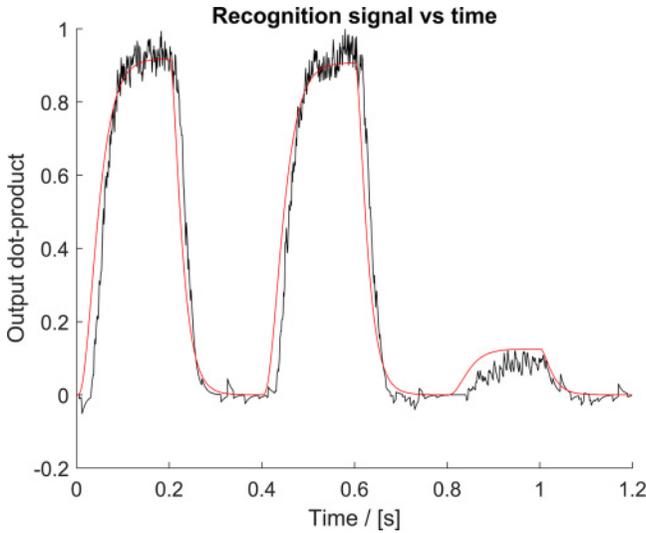


Figure 7: Computed recognition signal (dot-product) by the PPnS network using spiking neurons (black curve) and direct calculation (red curve) and showing excellent agreement in signal magnitude between the two, signifying that the network has implemented the dot-product operation well.

inputs, and a root-mean-squared error (RMSE) is calculated. A mean RMSE across the 100 runs is reported as a measure of the implementation's performance, and the standard deviation of this RMSE is chosen as a measure of error on the reported performance. Figure 8 shows the dependence of the RMSE on the optimization radius. Example runs for the OE 5D, 10D, and 15D tests are shown in red, together with the RMSE for the PPnS for 5D, 10D, and 20D tests in blue. For the OE dot-product implementation, there is a clear degradation of performance when going from 5D to 10D inputs, as well as for 10D to 15D, albeit the degradation is smaller. For the PPnS implementation, using the same number of neurons for the different dimensionalities, there is a more stable performance illustrated by the RMSE values staying at approximately 0.15, even when using 20D inputs. All in all, there is superior performance in the PPnS network and better scaling, so a smaller increase in the number of neurons to conserve the level of performance is needed. There also seems to be a smaller variation between trials when testing with constant radius for the PPnS implementation, but the implications of this result are not clear.

3.2 Connecting the Novelty System. To create the full novelty system, an input is connected in parallel with the LTM network to the developed PPnS novelty network implementation. The input also connects to the LTM

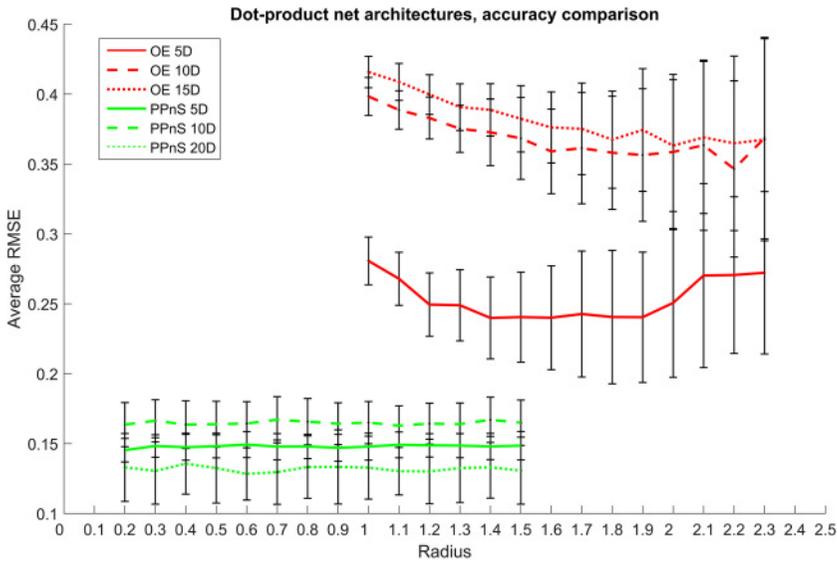


Figure 8: RMSE versus optimization radius for OE and PPnS architectures computing the dot-product for different dimensionalities. The 5D tests use 1000 neurons and 10D 2000 neurons; the 15D and 20D tests both use 4000 neurons. It is seen how the PPnS architecture achieves stable RMSE values across the different dimensionalities. For the OE implementation, there is a degradation of performance for rising dimensionality.

ensembles to evoke a memory response. The PPnS network compares the input to the recalled memory and computes the similarity (i.e., dot-product). The signal passes on to a group of neurons with a threshold of 0.4, so that only significant recalls are captured. This group then broadcasts the generalized recognition signal. This threshold is not necessarily biological, but it sets a lower limit, where there could be little difference between the calculated dot-product and background noise. The full implementation of the novelty detection system is shown in Figure 9. Notice that the PPnS network has been organized into a separate network object, called CA1, the name of the hippocampal subdomain whose function it emulates. The reason for this change is to simplify the network diagram, and the implementation of the CA1 subnetwork is exactly as shown in Figure 5. Again the inputs used to probe the system are seven vectors: the first six are stored in the memory, and the seventh is an unknown vector that probes the network's response to novelty. The individual storage networks have a random 20D vector assigned as their collective, preferred direction vectors, serving as a memory store. This is an implementation of the LTM described in Stewart et al. (2011), Crawford et al. (2013), and Voelker et al. (2014). The storagesum

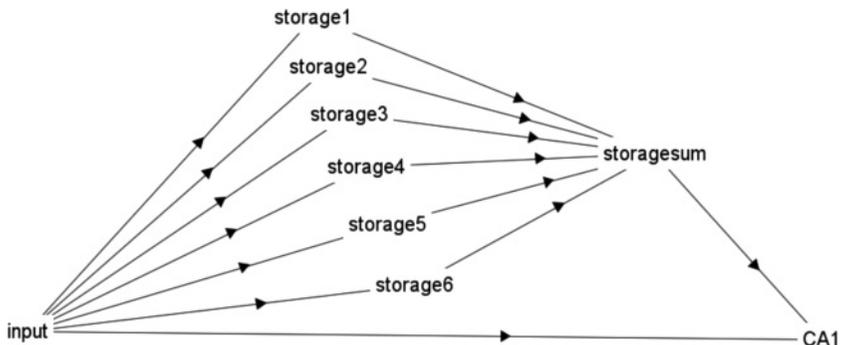


Figure 9: Fully connected novelty detection network with processed sensory input, long-term memories (100 neurons per ensemble), and simulated CA1 hippocampal subdomain with novelty detection capabilities (the PPnS network), consisting of 4000 spiking neurons.

ensemble simply adds up the six storage ensembles' activities, creating a 20D output of the memory stores.

Figure 10 shows the generated recognition signal as a black curve. The blue lines correspond to the times each probe vector is presented to the network. It is seen how the network has a large signal every time something known is presented, but as soon as a novel vector is shown, the signal drops to zero. These dynamics are similar to the functionality of the action selection system of the basal ganglia described by Stewart, Choo, and Eliasmith (2010). In this system, the action or inference selected by the network is the one that is inhibited the least.

The functionality of the developed novelty detection system, when it is situated in the hippocampus, could then be that the CA1 domain sends a signal that excites inhibitory interneurons, which prevents learning in, for example, the Schaffer collaterals between HC domains CA3 and CA1. This would effectively turn off learning according to the Ketz-O'Reilly computational model of HC functionality (Ketz, Morkonda, & O'Reilly, 2013). When a novel piece of information then appears, the recognition signal drops to zero, which consequently makes the inhibitory interneuron activity drop. The Schaffer collaterals can then be used for learning. This change in interneuron activity in the presence of novelty is similar to the dynamics reported by Nitz and McNaughton (2004). Furthermore, by comparing the network architecture of the novelty detection circuit to the described anatomy and circuitry of the hippocampal formation (Andersen, Morris, Amaral, Bliss, & O'Keefe, 2007), similar topographic projections are seen as from the sensory input projections (entorhinal cortex cellular layer III) to the CA1 and subiculum (which projects back to the deeper EC layer V).

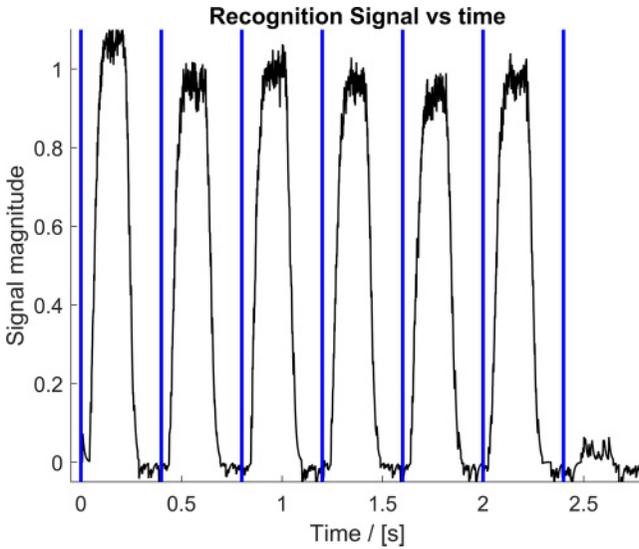


Figure 10: Recognition signal generated by the fully connected system. The blue lines correspond to the onset of input vector, showing a big difference in network response between known (first six inputs) and novel (last input) information.

Following the previous considerations, we conclude that the presented functional model of the CA1 subdomain agrees with both functional (Nitz & McNaughton, 2004; Larkin et al., 2014) and anatomical data (Andersen et al., 2007). It also provides a possible explanation for how switching learning on and off in the hippocampal short-term memory system could function.

4 Conclusion

Supported in part by anatomical data about the HCF, its circuits and connections with other brain areas (Andersen et al., 2007), and in part by behavioral studies done on rats (Larkin et al., 2014; Nitz & McNaughton, 2004), a novel model of the functioning of the hippocampal subdomain CA1 has been developed. From behavioral studies, the CA1 is proposed to be a novelty detection system, capable of generating a generalized novelty signal. Using the neural engineering framework (NEF), developed by Eliasmith and Anderson (2003), it was possible to develop a novelty detection system that could compare two inputs, in the form of 20-dimensional vectors (though generalizable to arbitrary dimensionality), representing semantic information according to the semantic pointer hypothesis (Eliasmith, 2013), and report the similarity between them by implementing the dot-product operation in a neural network.

This was achieved in practice by constructing a three-layer network, the pairwise product and summation (PPnS), the first layer taking parallel pairwise vector component projections and computing the products of each pair. These were then integrated by simply projecting all the components to the next layer, integrating the inputs by virtue of simulated intrinsic neuronal dynamics. Finally, the last layer represents this difference in input, effectively computing a recognition signal. One of the major results of this modeling work is that it is generally less economical and scalable to implement the dot-product operation in one neuronal ensemble compared to a more component-wise implementation. This was demonstrated for various optimization radii and dimensionalities, showing an RMSE value between the OE network signal and the direct computation of $\text{RMSE}_{\text{OE15D}} = 0.3632 \pm 0.0471$ for 15D inputs using 4000 neurons with an optimization radius of 2.0. In comparison, the PPnS implementation achieved an RMSE of $\text{RMSE}_{\text{PPnS20D}} = 0.1283 \pm 0.0187$ for 20D using 4000 neurons with an optimization radius of 0.6. Splitting the dot-product computation into components yields a more efficient assignment of the operation to the neural network, which can be scaled to higher dimensionalities, in a modular fashion by adding more components.

How this hippocampal subdomain model would fit into a declarative memory system was then demonstrated by first implementing an autoassociate long-term memory (LTM) to emulate neocortex storage, inspired by Stewart et al. (2011). The LTM was then connected with the novelty detection network and the input connected in parallel to the LTM and the novelty unit, allowing comparison between the retrieved memory and the original input. Adding a thresholded neural layer, with a threshold of 40%, resulted in novelty detection dynamics with a large signal difference between novel and known inputs. As the LTM responds with noise to unknown inputs, this results in a large recognition signal when no novelty is present and no signal when novelty is presented to the system. This is in accordance with available electrophysiological evidence of novel detection in rats, assuming interneuron activity similar to that in the basal ganglia action selection system, where the least inhibited action is selected.

Parallels can be drawn from this to the developed novelty network. The default state, when no novelty is present, has high tonic inhibitory interneuron activity, driven by the excitatory neurons representing the dot-product signal. This could block learning if the interneurons projected to the Schaffer collaterals between hippocampal layers CA3 and CA1. When novelty is then presented to the simulated network, the activity of the excitatory neurons representing the dot-product would drop to low levels and, consequently, the inhibitory interneuron activity as well. With decreased interneuron activity, the Schaffer collaterals are unblocked, and learning can occur.

The developed model is in general agreement with anatomical data, as the architecture is structurally similar to the topographical map between

the inputs (EC layers) and the CA1 layer (the novelty detection network) as reported in anatomical studies and summarized in Andersen et al. (2007). The model of the CA1 also agrees with the experiments of Nitz, & McNaughton (2004), showing decreased interneuron activity in the presence of novelty. In conclusion, our study presents a biologically realistic model of the functioning of the CA1 subdomain of the hippocampal memory system, explaining how memory storage can be switched on and off according to a recognition signal.

References

- Alzheimer's Association. (2014). 2014 Alzheimer's disease facts and figures. *Alzheimer's and Dementia*, 10(2), 47–92.
- Andersen, P., Morris, R., Amaral, D., Bliss, T., & O'Keefe, J. (2007). *The hippocampus book*. New York: Oxford University Press.
- Bekolay, T., Bergstra, J., Hunsberger, E., DeWolf, T., Stewart, T. C., Rasmussen, D., ... 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
- Choo, X. (2010). *The ordinal serial encoding model: Serial memory in spiking neurons*. Master's thesis, University of Waterloo.
- Crawford, E., Gingerich, M., & Eliasmith, C. (2013). Biologically plausible, human scale knowledge representation. In *Proceedings of the 35th Annual Conference of the Cognitive Science Society* (pp. 412–417). Austin, TX: Cognitive Science Society.
- Eliasmith, C. (2013). *How to build a brain*. New York: Oxford University Press.
- Eliasmith, C., & Anderson, C. H. (2003). *Neural engineering: Computation, representation, and dynamics in neurobiological systems*. Cambridge, MA: MIT Press.
- French, R. M. (1999). Catastrophic forgetting in connectionist networks. *Trends in Cognitive Sciences*, 3(4), 128–135.
- Georgopoulos, A. P., Schwartz, A., & Kettner, R. E. (1986). Neuronal population coding of movement direction. *Science*, 233, 1416–1419.
- Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proc. Natl. Acad. Sci.*, 79, 2554–2558.
- Ketz, N., Morkonda, S. G., & O'Reilly, R. C. (2013). Theta coordinated error-driven learning in the hippocampus. *PLOS Computational Biology*, 9(6). doi: 10.1371/journal.pcbi.1003067
- Larkin, M. C., Lykken, C., Tye, L. D., Wickelgren, J. G., & Frank, L. M. (2014). Hippocampal output area CA1 broadcasts a generalized novelty signal during an object-place recognition task. *Hippocampus*, 24(7), 773–783.
- Lisman, J. E., & Grace, A. A. (2005). The hippocampal-VTA loop: Controlling the entry of information into long-term memory. *Neuron*, 46(5), 703–713.
- Nitz, D., & McNaughton, B. (2004). Differential modulation of CA1 and Dentate gyrus interneurons during exploration of novel environments. *Journal of Neurophysiology*, 91(2), 863–872.
- O'Reilly, R. C., Bhattacharyya, R., Howard, M. D., & Ketz, N. (2014). Complementary learning systems. *Cognitive Science*, 38(6), 1229–1248.

- O'Reilly, R. C., & McClelland, J. L. (1994). Hippocampal conjunctive encoding, storage, and recall: Avoiding a trade-off. *Hippocampus*, 4(6), 661–682.
- Stewart, T. C., Choo, X., & Eliasmith, C. (2010). Symbolic reasoning in spiking neurons: A model of the cortex/basal ganglia/thalamus loop. In *Proceedings of the 32nd Annual Meeting of the Cognitive Science Society* (pp. 1100–1105). Austin, TX: 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.
- Voelker, A. R., Crawford, E., & Eliasmith, C. (2014). *Learning large-scale heteroassociative memories in spiking neurons*. London, Ontario: Unconventional Computation and Natural Computation.

Received June 2, 2015; accepted December 14, 2015.