

*To appear in The Rewiring Brain: A Computational Approach to Structural Plasticity in the Adult Brain
edited by Arjen van Ooyen and Markus Butz, Elsevier.*

Restricted Boltzmann Machine models of hippocampal coding and neurogenesis

Rory Finnegan
McMaster University
McMaster Integrative Neuroscience Discovery & Study
204-583-4012
finnegrj@mcmaster.ca

Mark Shaw
University of Toronto
Department of Psychology
647-702-7984
mark.shaw@mail.utoronto.ca

Suzanna Becker
McMaster University
Department of Psychology, Neuroscience and Behaviour
905-525-9140 ext23020
beckers@mcmaster.ca

Restricted Boltzmann Machine models of hippocampal coding and neurogenesis

Rory Finnegan, Mark Shaw and Suzanna Becker

Abstract: The hippocampus has been traditionally viewed as a memorization device, creating orthogonalized representations (pattern separation) in the dentate gyrus, and performing associative retrieval (pattern completion) in the CA3. Moreover, neurogenesis in the dentate gyrus is widely assumed to increase pattern separation. Evidence that neurogenesis is important for behavioural discrimination has been erroneously taken as supporting the pattern separation assumption. Instead, we propose that the hippocampus forms a probabilistic, generative model of its input, using forward and feedback connections for encoding versus reconstruction. Using the Restricted Boltzmann Machine, we model the developmental trajectory of adult-generated neurons from hyperactive, hyperplastic, sparsely connected young neurons to less plastic, more densely connected mature neurons under tight inhibitory control. Models with neurogenesis are more robust against interference, while paradoxically generating more overlapping representations (less pattern separation). When applied to more realistic grid cell and boundary vector cell inputs, the model learns place cell representations. Finally, we simulate a full multi-layer hippocampal model with neurogenesis, and discuss how it can learn representations of sequential, complex events.

Keywords: hippocampus, neurogenesis, dentate granule cells, recall, neural networks, retroactive interference, proactive interference, Restricted Boltzmann Machine, grid cells, boundary vector cells

Introduction

The hippocampus has intrigued researchers for many decades because of its vital role in the formation and retrieval of complex associative memories. Inspired by Marr's simple memory theory of neural coding in the archicortex,¹ a computational framework that included an initial coding stage followed by an associative retrieval stage, subsequent computational modellers have further fleshed out the contributions of different hippocampal sub-regions to these memory functions.²⁻⁷ Under this widely accepted, traditional view of the hippocampus, the initial coding stage takes place in the dentate gyrus (DG), followed by an associative retrieval stage in the CA3 and CA1 regions. In these models, the DG serves to generate less overlapping, more orthogonalized representations (pattern separation) through sparse coding, while the CA3, with its dense recurrent collateral connections, performs cued retrieval (pattern completion). The associative pathway from CA3 to CA1 in some models further contributes to pattern completion. The basic circuit of this traditional model is illustrated in Figure 1.

INSERT FIGURE 1 ABOUT HERE

The traditional view of the hippocampus as a memorization and cued retrieval device, as described above, has garnered considerable empirical support. Direct evidence from unit recordings of ongoing neuronal activity and immediate early gene (IEG) markers of recent neuronal activity support a role for the DG in sparse coding and pattern separation.⁸⁻¹² Similarly, there is direct evidence of a role for the CA3 in cued recall and pattern completion.¹³⁻¹⁵ However, more recent data suggest that this traditional view is too simplistic. Firstly, a large body of evidence suggests that the hippocampal "memory system" is as much involved in

imagining and predicting the future as it is in remembering the past. For example, patients with medial temporal lobe (MTL) damage that includes the hippocampus not only show episodic and contextual memory deficits; they also have great difficulty imagining future scenarios.¹⁶⁻¹⁹ Events that they attempt to imagine, like the past events that they attempt to remember, are notably lacking in episodic detail. Secondly, recent data call into question whether the MTL memory system is merely retrieving and remembering information, or may also be involved in the ongoing classification and perception of stimuli. A prime example is the coding of space. In both rodents and humans, hippocampal “place cells” reflect the animal’s current location in space and are critical for ongoing navigation.^{20,21} The importance of the MTL for representing perceptual information also extends to non-spatial perceptual judgements. A growing body of evidence points to a central role for the hippocampus in representing higher order perceptual features and making complex perceptual decisions (for a review, see Lee et al²²), for example, discriminating faces and objects from novel stimuli, even when those faces and objects were not subsequently well remembered.²³

The fact that the hippocampal system seems to be involved in perception, imagery and predicting the future is inconsistent with the notion that its chief function is exact memorization for subsequent cued retrieval. Instead, the traditional view of the hippocampus (Figure 1) must be re-thought. An updated view is that the hippocampus operates integrally with the neocortex to encode, represent, perceive and predict information. As a first step toward modelling the role of the hippocampus within this broader view, Kali and Dayan²⁴ proposed a hierarchical, probabilistic model of the cortex, with the hippocampus situated at the top of the hierarchy. The hippocampus and neocortex jointly form a multilayer model of the sensory input. Kali and Dayan’s model provides a compelling alternative account of observations such as “sequence

replay”, the re-activation of recently experienced sequences of place representations during sleep.^{25,26} The standard view is that hippocampal sequence replay serves to consolidate memories in the neocortex, making them independent of the hippocampus. In contrast, Kali and Dayan propose that event replay serves to maintain the correspondence between neocortical and hippocampal representations, as both may be evolving over time. This view fits within the broader framework of the brain as a probabilistic, generative, hierarchical model that includes the hippocampus.

Another key property of the hippocampus that is missing from the traditional model is the occurrence of ongoing neurogenesis throughout the lifespan. Adult neurogenesis in the hippocampus has profound implications for neural coding. Mature dentate granule cells (DGCs) fire very sparsely, as they are under tight control by inhibitory interneurons, and might reasonably be assumed to contribute to pattern separation in the DG. On the other hand, the newly generated neurons have very different properties. By about 3-4 weeks of age they have matured to the point that they are able to fire action potentials and contribute to neural coding, and yet at this stage they are much more plastic^{27,28} and more highly excitable^{29,30} than mature DGCs. Over the next several weeks, the young neurons mature and become progressively more like adult DGCs. How would this varying population of DGCs affect learning and memory? In this chapter, we consider how the idea of a probabilistic, generative model of encoding and prediction could be implemented within the hippocampal circuit using the Restricted Boltzmann Machine (RBM)³¹ to simulate the learning process. Further, we consider the contribution of neurogenesis to coding in this model. We then consider how the model could be extended to include more realistic inputs in the form of grid cells and boundary vector cells in the entorhinal cortex and subiculum / pre-subiculum. Simulated models with these more naturalistic inputs lead

to the emergence of place cells when applying the RBM learning equations. Finally, we present new simulations of a more complete hippocampal model that includes multiple stacked RBMs, and we outline how this full hippocampal model should account for a wide range of neurobiological data including sequence coding and replay.

A RBM Model of Learning and Neurogenesis in the DG

The RBM is a type of artificial neural network which learns a set of weights so as to form a probabilistic, generative model of a dataset.³¹ The network consists of a set of reciprocally connected stochastic units, partitioned into visible and hidden layers of units, with no within layer connections. The architecture of a one-layer RBM can be seen in the middle panel of Figure 2. These bidirectional connections between the visible and hidden units form a bipartite graph. Weights in the network are updated based on the difference between the data-driven representations (and the probability distribution of the dataset) and the model's expectation-driven representations. While exact calculation of these expectations is intractable, they can be approximated through brief Gibbs sampling.³¹ We used 1-step Gibbs sampling, and the contrastive divergence (CD) learning rule³¹ in our RBM based model:

$$\Delta W_{ij} = \epsilon \left((v_i h_j)_{data} - (v_i h_j)_{recon} \right) \quad (1)$$

where ΔW_{ij} is the update to the weight W_{ij} between the i^{th} visible unit and j^{th} hidden unit, calculated by taking the difference between the data-driven and reconstructed Hebbian terms (product of i th visible state and j th hidden state), multiplied by a learning rate ϵ . The input vector v_{data} is the data-driven visible state vector and h_{data} is the data-driven hidden state generated by clamping the visible units' states to v_{data} and sampling the hidden units' states

according to equation 5. Thus, $(v_i h_j)_{data}$ represents the product of the i^{th} visible and j^{th} hidden activations in the data-driven phase. v_{recon} is a reconstruction of the input vector generated by clamping the states of the hidden units to the data-driven pattern h_{data} and sampling the states of the visible units according to equation 4. h_{recon} is then created in the same way as h_{data} , but by clamping the visible units' states to v_{recon} .

$$\Delta a_i = \epsilon(v_{i_{data}} - v_{i_{recon}}) \quad (2)$$

$$\Delta b_j = \epsilon(h_{j_{data}} - h_{j_{recon}}) \quad (3)$$

Units' binary states are updated probabilistically according to equations 4 and 5 below:

$$p(v_i = 1|h) = \sigma\left(a_i + \sum_j h_j w_{ij}\right) \quad (4)$$

$$p(h_j = 1|v) = \sigma\left(b_j + \sum_i v_i w_{ij}\right) \quad (5)$$

where a_i and b_j represent biases which provide a mechanism for shifting the output of the sigmoid activation function $\sigma(x) = 1/(1 + e^{-x})$. A unit's bias is similar to the negative of a unit's threshold in other neural network models. The weight updates presented in equations 2 and 3 show that these bias weights are updated using the same positive and negative Hebbian terms used in updating W as shown in equation 1.

We can see in equation 1, that the positive Hebbian term $(v_i h_j)_{data}$ associates data-driven input and hidden state vectors, while the negative Hebbian term $(v_i h_j)_{recon}$ tries to “unlearn” the association between the corresponding reconstructed visible and hidden state vectors. Theoretically, the learning procedure should converge when the probability distribution of its internal reconstructions of the training patterns exactly match that of the corresponding data-driven states. In general, an RBM model’s reconstructions of the training patterns are obtained by brief Gibbs sampling, alternatingly sampling hidden and visible unit states that are nearby data-driven states using the model’s bottom-up and top-down weights respectively. Similar to a Hopfield network, the RBM utilizes a local and unsupervised learning rule, which also minimizes the free energy within the network. However, the presence of hidden units, along with the ability to stack RBMs to form deep networks, provides greater memory capacity. Furthermore, the ability to run the RBM in unclamped, top-down or generative mode, may provide a way of simulating dreaming along with memory reconstruction and cued recall.

The Role of Young DGCs in Memory Encoding

We recently proposed a novel RBM based model of the DG which incorporates the developmental trajectory of adult-born DGCs.³² In this model, a single RBM represents the entorhinal cortex (EC) input and DGC layer with its visible and hidden units respectively. As the model DGCs undergo development, they become progressively less plastic, more sparse in their firing, and more densely connected to their entorhinal inputs. We demonstrate how these properties can explain the importance of adult-generated DGCs for memory across both short and long time scales.

In the model, the maturational trajectory of adult born DGCs is loosely based on mouse

data, for DGCs from the third week of maturation onward. It is at about 3-4 weeks of age that adult born DGCs establish synaptic afferent and efferent connections and are able to fire action potentials.³³ As compared to more mature neurons, young DGCs have a higher input resistance, lower capacitance, lower activation threshold and a slower membrane time constant. As a result, 3-4 week old DGCs can be described as being more excitable, while having a smaller and slower action potentials.^{27,28} Moreover, the young neurons are more sparsely connected to their perforant pathway inputs from the EC relative to mature DGCs.³⁴ From weeks five through eight the young neurons undergo a gradual decline in synaptic plasticity and are increasingly regulated by feedback inhibition.³⁵ By the eighth week, the physiological properties of adult-generated DGCs are largely indistinguishable from their mature counterparts.^{34,35}

While several replacement and additive models of neurogenesis have looked at how new neurons affect learning,^{36,37} few models have considered the full range of unique properties of adult hippocampal neurogenesis including the developmental trajectory of adult-generated neurons: changes in plasticity, connectivity, excitability and survival versus apoptosis. We model the neural trajectory of young DGCs by incorporating additional constraints into the learning equation, including a dynamic learning rate and separate penalty terms for sparse activation (or excitability) versus sparse connectivity. These dynamic constraints are calculated using a Gompertz function $g(t)$ (a roughly S-shaped curve)³⁸ to model each of the neural growth and decay parameters. At each simulated time step, the age of a hidden unit is increased, and its constraint parameters are updated as follows. A sparse connectivity constraint describes the level of interconnectedness between the visible and hidden layers. Young DGCs are randomly connected to only 30% percent of EC input units, while mature DGCs are fully connected to the input layer. The connectivity proportion varies smoothly from 30% to 100% in direct proportion

to $g(t)$, the Gompertz function. Thus, young neurons are initially sparsely connected, and transition to becoming fully connected to the input layer. Similarly, the learning rate, which can be thought of as a neuron's plasticity level, is defined as $1 - g(t)$ normalized to lie between 0.0025 and 0.1. Thus, the plasticity level decreases over time as the neuron matures. Finally, our sparse activation cost increases over time, in direct proportion to $g(t)$, scaled to lie between 0.0 (young DGCs) and 0.9 (mature DGCs). Thus, young neurons are unaffected by the sparseness constraint, while mature neurons are heavily constrained to have low firing probabilities, and therefore adopt sparse codes, simulating the increasingly tight regulation by feedback inhibition over DGCs as they mature. Given these variable properties, the learning rule can be redefined as

$$\Delta W_{ij} = \epsilon_j \left((v_i h_j)_{data} - (v_i h_j)_{recon} \right) - (\lambda_j W_{ij}) - cost(q_j - p) \quad (6)$$

where the learning rate ϵ , weight decay λ and sparsity cost term q are now each weighted by dynamically changing vectors of values rather than static hyper-parameters, while p represents the current probability of hidden unit j activating.

Returning to our primary objective, what impact does the developmental trajectory of young DGCs have on hippocampal learning and memory? To investigate this, we designed a set of experiments to monitor proactive and retroactive memory interference over short and long time scales. Rodents with reduced neurogenesis show greater interference on challenging memory tasks, including *proactive interference* between previously learned information and newly learned information,³⁹ and *retroactive interference* of newly acquired memories on the retrieval of previously learned information,⁴⁰ when interfering items are learned in different sessions spaced over a relatively long time period of days to weeks. It is relatively easy to see

why neurogenesis would reduce interference between similar items that are encountered several days apart, as there would be a substantial change in the population of young highly plastic DGCs when the second item is encountered. Interestingly, there is also evidence for a role for neurogenesis in resolving interference between similar items that are interleaved within a single learning/testing session.⁴¹ It is more challenging to explain why neurogenesis would help in this situation, as the population of young DGCs would not be expected to change substantially over a time scale of minutes. We hypothesized that the very sparse, random connectivity of young neurons might be crucial for generating distinct representations of similar items encountered within the same session. This could allow different young DGCs to respond to different features of the input, even though they are hyper-excitabile and highly plastic.

Given the findings in the literature described above, we sought to test whether the model would exhibit proactive and/or retroactive interference, either within or across learning sessions, as a function of neurogenesis levels. Interference was created by training our models iteratively on highly overlapping sets of patterns with the expectation that new patterns would be more difficult to learn when similar patterns were learned previously (proactive interference), and distally learned patterns would be more difficult to retrieve when new similar patterns were subsequently learned (retroactive interference). We also examined whether this interference would manifest within a single learning session versus across a longer time scale of several days of simulated learning sessions.

In the same-session test conditions, the presence of a few young plastic neurons was expected to help with rapid encoding without creating interference between similar patterns. Likewise, for multi-session testing, given that a new pool of young neurons would be available at each learning session, it was expected that the neurogenesis models would perform even better

than same-session tests. While this initial set of simulations^{32,42} focused on comparing the encoding performance of a single layer model with just the hippocampal DG layer, with and without neural turnover, later in this chapter we present new simulation results with a full multilayer hippocampal model.

The results of these experiments showed that the incorporation of neurogenesis led to improved encoding performance on both short and long time scales, in both the same and multi-session tests, relative to a control model consisting of an RBM with a sparsely coded DG layer (but lacking neurogenesis). In the case of same session tests, the presence of a few young, more plastic, neurons likely provided the necessary advantage without introducing interference. Interestingly, the neurogenesis models had more overlap among hidden unit activations than the control sparse RBM network. Thus, although the RBM models were free to learn either a sparse code or a highly distributed code to represent the training patterns, our simulation results indicate that the neurogenesis models learned that it was advantageous to have less sparse activations. This may seem paradoxical, as sparse coding and pattern separation are widely believed to be the main mechanisms by which the DG avoids interference between overlapping items. Instead, our simulations suggest an additional mechanism for mitigating interference. The hyperactive young neurons generate highly distributed, overlapping codes for similar items, which allows them to learn distinct features despite being very sparsely connected to their inputs. At the same time, the neurogenesis models showed reduced proactive interference. The increase in accuracy on subsequently learned sets of overlapping patterns suggests that the neurogenesis models may be better at distinguishing novel and common elements to each group of patterns.

In the multi-session tests, the incorporation of neural maturation and turnover provided less benefit to overall performance than expected. While the non-sparsely connected neurogenesis

model did see about a 1% increase in performance over the same session tests, the addition of sparse connectivity imparted no further improvement; this model performed about the same as its non-sparsely connected counterpart. Thus, sparse connectivity was most important in protecting against interference between patterns learned within a single session, while neural turnover itself was sufficient to mitigate interference at longer time scales. These results suggest that sparse connectivity and neural turnover work in equilibrium with each other, and either one may be relatively more important, depending on the learning task demands.

Simulating the Emergence of Place Cells in RBMs with More Naturalist Inputs

Our simulations described up to this point used highly simplified synthetic data sets of random patterns with varying degrees of overlap. An important next step in developing more realistic hippocampal models is to use more biologically realistic inputs. Since the discovery of place cells within the hippocampus by O'Keefe and Dovstrovsky,⁴³ their role in spatial memory and learning have been a focus of hippocampal research for decades. Located in the DG, CA3 and CA1 layers,⁴⁴ place cells fire in response to very specific spatial positions or place fields.⁴³ It is believed that these place cells help form a spatial map in hippocampal memory, but how are these place fields formed?

There are two major types of spatial inputs to place cells that have been relatively well characterized: grid cells and boundary vector cells. Either or both of these inputs could drive the formation of place cells, as explored in the simulations reported in this section. Grid cells, first discovered by Hafting et al. in medial entorhinal cortex,⁴⁵ have subsequently also been found in pre- and para-subiculum,⁴⁶ all of which are prominent inputs to the hippocampus. Each grid cell has a tuning curve that exhibits a multimodal pattern of firing within an enclosed region of space,

such that the cell fires in multiple regularly spaced locations. Remarkably, the locations of peak firing of a given grid cell correspond to the vertices of a hexagonal grid. An example of a simulated input grid cell's tuning curve can be seen in the left panel of Figure 2. Thus, unlike place cells, grid cells respond to multiple locations.^{45,47} It is believed that the diversity of the grid cell receptive fields, with varying spatial scales and phase offsets, can encode the spatial locations of an entire environment.⁴⁷ A second type of spatial feature that provides input to hippocampal place cells comes from boundary vector cells (BVCs), which are tuned to the distances and directions to environmental boundaries. For example, a particular boundary vector cell might respond best to a boundary along the west edge of an environment, whenever the rat is a short distance from that wall, but respond equally well to all locations along the wall. BVCs were first predicted by a place cell model developed by Burgess and colleagues,^{48,49} and empirical evidence for boundary vector cells was later discovered in the subiculum,⁵⁰ an important input to the hippocampus.

We have shown that an RBM that receives input from a visible layer consisting of simulated grid cells can generate spatially localized place-tuned neurons in the hidden layer of the RBM.⁵¹ The grid cell input was generated by applying the oscillatory interference model of grid cell firing developed by Burgess and colleagues⁵² and adapted by Hasselmo et al.⁵³ The data taken as input to the model is real data recorded by Hafting et al.⁴⁵ as a rat traversed a square environment. Hasselmo applied Burgess' model to this spatial trajectory data, using three head direction cells which respond in relation to the speed of the rat and its orientation. At the same time, their model incorporates an inherent oscillation function, which together cause the formation of approximately realistic grid cells. These grid cell patterns are calculated along the same pathway of the rat recorded by Hafting et al. An example of a grid cell pattern imposed

over the rat's pathway (grey lines) is shown in the left panel of Figure 2. The activities of a population of 100 grid cells with different spatial scale and phase tuning formed the input to the RBM at each point in time. The dentate layer was also modelled by a population of 100 neurons. The expectation was that the RBM model would successfully be able to re-encode the simulated grid cell fields into place cell fields, which correspond to highly localized unimodal place fields within the source spatial trajectory data.

The results of this study were somewhat mixed. About 40% of DG cells formed small unimodal, localized place fields, such as the one shown in the right panel of Figure 2. These are typical of place tuning of DG cells. Another 25% formed bimodal place fields, with two different locations causing the cell to fire, and the remainder showed no specific tuning. For simulations with only 10 DG cells, in contrast, the model developed multi-modal firing fields.

These simulations demonstrate that the RBM model was able to re-encode the grid cell fields into a place cell representation. However, in order to identify the expected unimodal place cell firing, the model needed to be trained for 4,000 iterations with at least as many place cells as grid cell inputs.⁵¹ Nonetheless, this model demonstrates how an RBM can be used to simulate a real-world situation based on biologically realistic input patterns, bridging the gap between theoretical and experimental research on hippocampal coding and memory.

INSERT FIGURE 2 ABOUT HERE

While many have assumed that grid cells are the primary source of input driving the spatial selectivity of place-tuned hippocampal neurons, there is evidence that other factors influence place cell tuning, including the locations of environmental boundaries and context.^{54,55}

Moreover, it has been argued that BVCs, rather than grid cells, might be the major driving force for the development of place fields.⁵⁶ According to this alternative view, rather than grid cells and place cells forming a hierarchy of representations, they constitute complementary, interacting representations, with the grid cells maintaining one type of spatial representation driven by self-motion, and the place cells maintaining another type of spatial representation driven by environmental features such as landmarks (BVCs), objects, and other contextual features. As originally demonstrated by Hartley et al.⁴⁸, a place-selective tuning curve can be computed from a combination of appropriately tuned BVCs. We have further shown how this representation could be learned by an RBM with an input layer consisting of differently tuned BVCs.⁵⁷ The model was trained by “traversing” an enclosed environment, generating corresponding input patterns across the population of BVCs inputs. From this input, the model learned the mapping from BVCs to place-tuned neurons.⁵⁷

Our simulations with realistic BVC or grid cell inputs lend further support to the notion that the hippocampus forms generative, probabilistic models of its input, and that the RBM can provide a biologically plausible mechanistic account of the learning process. However, all of the simulations reported up to this point have been with single layer networks. We next consider a multi-layer stacked RBM model of the full hippocampal circuit.

The Effect of Young DGCs on Learning in the Full Hippocampal Model

The one-layer RBM models considered up to this point provide a good account of place tuning and neurogenesis in the DG, but in the absence of the CA3 and CA1 layers, they lack the capacity to model cued recall, sequence learning and replay. As a continuation of our study on the role of young DGCs in hippocampal learning,⁴² we now present a full multilayer

hippocampal model to explore the role of young DGCs in cued recall performance. In order to model cued recall within the hippocampus we must extend our model to include the associative CA3 layer, which by the “traditional account” discussed in the introduction, performs associative retrieval and pattern completion. Another view of the role of the CA3, with its rich recurrent collateral interconnections, as well as its associative connections to area CA1, is that CA3 (possibly jointly with CA1) plays a crucial role in the formation of temporal associations, anticipatory coding and sequence learning.⁵⁸⁻⁶⁰ This is the view adopted here.

Interestingly, by adding another layer to the DG model we can address an important flaw in the original DG model. Our existing DG model requires reciprocal connectivity between the input and output layers, whereas the known anatomy of the DG does not support this architecture: DGCs do not project back to the EC; however, this can be addressed by incorporating reciprocal connections between the CA3 and the DG,^{61,59} and between the CA3 and the EC via a stacked RBM architecture. Specifically, by converting the DG layer into a Conditional Restricted Boltzmann Machine (CRBM), this layer maintains a set of bidirectional connections to the CA3, and only requires a set of conditional unidirectional weights from the EC. This full hippocampal circuit model is used to explore the functional impact of young vs mature DGCs on hippocampal learning, when investigating the performance changes on memory recall (pattern completion) and sequence replay tasks.

In order to incorporate the CA3 layer into the model, and model associative and temporal learning, a way of modelling the recurrent collaterals is required. While recurrent neural network learning procedures such as back-propagation through time,⁶²⁻⁶⁴ long short-term memory networks⁶⁵ and liquid state machines⁶⁶ have proven effective for modelling sequential data,^{67,68} they are not considered to be biologically plausible. We would like to retain the appealing

features of the RBM, while extending the operation of the network to modelling sequential data. One such extension is the CRBM,⁶⁹ which offers a tractable approximation to a fully recurrent network. The CRBM extends the RBM by adding visible-to-visible and visible-to-hidden autoregressive weights from other (or conditional) visible inputs,⁶⁹ but treats these extra inputs from the previous time step as fixed, rather than simulating a fully recurrent network. However, the CRBM is not limited to conditioning on these historical observations. By making the DG layer a CRBM with direct input from the EC, and including bidirectional connectivity between the CA3 and DG, we can eliminate the need for the feedback projection from the DC to the EC. To simplify, we combine the CA3 and CA1 regions into a single recurrently connected layer, implemented as a CRBM with input from the EC and hidden-to-hidden layer temporal connections. This multilayer architecture is shown in Figure 3. Our bidirectional weights between the CA and DG layers represent the mossy fibre inputs to the CA3 and back-projections from the CA3 to the DG. Further, we can use the autoregressive visible-to-hidden weights to represent the EC to DG connections. By doing so, the DG will learn patterns of activation in the CA3 by conditioning on the EC. This provides an RBM multilayer hippocampal model that correctly accounts for the directionality of the known connectivity within the hippocampal structure.

One advantage of multi-layer models built upon the RBM is that they can be trained one layer at a time and stacked together. By greedily training a 2-layer network and incorporating the developmental trajectory of newly generated neurons (as in our previous DG model) to the CRBM, we can build a full hippocampal model that supports both the recurrent connections in the CA3 and the unidirectional connections from the EC to the DG. We begin by training the CA3 layer on the EC input. We then lock the EC-to-CA3 layer weights and transform the EC

input through these weights, producing inputs for training the DG layer. We proceed by training the DG layer on these CA3 outputs, conditioned on the initial EC patterns. While the CA3 layer supports the conditional input from previous time-steps and will be included in future work, we did not model this property in the simulation experiments reported here. Once both layers of the network are trained, the cued recall performance can be tested by presenting a test pattern as the initial EC input, passing through both layers, sampling the DG hidden layer and reconstructing the pattern back through to the EC. This architecture is based on one proposed by Becker and Hinton⁷⁰ and is shown in Figure 3.

INSERT FIGURE 3 ABOUT HERE

To evaluate the role of young DGCs on learning in the full hippocampal model, the same experimental design from the previous study was reused in the current simulations. As before, we incorporated the same mechanisms for simulating the developmental trajectory of young DGCs. To review, we designed a set of experiments to monitor proactive and retroactive memory interference over short and long time scales. The model was trained for multiple sessions on a different group of patterns in each session, and there was interference across sessions as well as within sessions. This interference was created by training our models iteratively on highly similar patterns. The expectation was that new patterns would be more difficult to learn when similar patterns were learned previously (proactive interference), and distally learned patterns would be more difficult to retrieve when new similar patterns were subsequently learned (retroactive interference). We also examined whether this interference would manifest within a single learning session and/or across a longer time scale of several days

of simulated learning sessions. The expectation was that the same improvements to encoding performance in the single layer model would be observed in the cued recall experiments with the full hippocampal model.

The results of these experiments are summarized in Figure 4. As with the one-layer DG model, the multi-layer hippocampal model showed reduced proactive and retroactive interference when neurogenesis was added, as shown in Figure 4, panels A and C respectively. As with the one-layer model, the degree of DG hidden unit overlap versus input pattern overlap was higher for models with neurogenesis (a decrease in pattern separation), as the more active young DGCs are less selective in their firing patterns (Figure 4, panel B). Across different simulation runs, the neurogenesis models achieved varying degrees of pattern separation (Figure 4, panel B) but in all cases the DG overlap was much greater than that of the sparsely coded models lacking neurogenesis. In spite of this reduced pattern separation, the models with neurogenesis usually outperformed those without (Figure 4, panel D). Interestingly, the improved performance for the neurogenesis models appears to be magnified relative to the single EC-DG layer network from the earlier study.

INSERT FIGURE 4 ABOUT HERE

Discussion

In this chapter, we have presented several simulation studies that demonstrate how the RBM can be used to model the contribution of neurogenesis in the DG to learning, and to model hippocampal learning more generally. We have demonstrated in both a single-layer RBM , and a

multi-layer model of the full hippocampal circuit, that neurogenesis contributes to interference reduction. We have also shown that the model can be extended to incorporate more biologically realistic grid cell and boundary vector cell inputs. The incorporation of these more realistic inputs led, in both cases, to the emergence of place cells.

Two novel insights emerge from our simulations of models with neurogenesis. First, neurogenesis contributes to interference reduction at both short and long time scales. At short time scales, when there is interference between similar items encountered within a single learning session, neural turnover is less important than simply having a pool of highly plastic, highly active and very sparsely connected neurons available for coding. There is very little turnover within the course of a single learning session. All else being equal, a set of highly plastic, highly active young neurons could potentially lead to even greater interference, by being equally activated for all items encountered within a session. However, in our simulations, the incorporation of sparse connectivity prevented this from happening, and in fact, sparse connectivity was critical in allowing the young neurons to generate distinct representations of overlapping items within a session. Over longer time scales of days to weeks, neurogenesis can reduce interference even without the sparseness constraint. This is because of ongoing proliferation and maturation of the DGCs. Thus, from one session to the next, the pool of available hyper-plastic and hyper-excitabile young neurons will have evolved. This continuously evolving pool of young plastic neurons ensures that slightly different subsets of young neurons will be recruited for encoding different memories.

A second novel insight that emerges from our simulations of neurogenesis is that the young, highly active neurons can contribute to interference reduction, perhaps paradoxically, while at the same time decreasing pattern separation. Here, we use pattern separation in the manner in

which it was originally intended, referring to the generation of more orthogonalized, or less overlapping, neural codes for overlapping input patterns. In the recent literature, the term pattern separation has also come into widespread use to refer to almost any behavioural task that requires discrimination between similar items. This is rather problematic, as discussed in further detail by Becker,⁷⁴ as pattern separation (orthogonalization) and behavioural discrimination do not necessarily go hand in hand. As shown in our simulations, the addition of neurogenesis caused our models to generate more overlapping, and yet distinct neural codes for similar items. Thus, orthogonalization (pattern separation, as originally conceived by computational modellers) is not the only possible mechanism for achieving interference reduction.

Future experimental studies The widespread view that neurogenesis increases pattern separation has also been called into question by some empirical findings. For example, it is well established that neurogenesis levels decline with age.^{77,78} If neurogenesis contributes to increased pattern separation, then pattern separation should also decline with age. However, direct assessment of neural activation patterns generated in similar contexts using immediate early gene markers indicate that neural overlap declines with age (indicating stronger pattern separation) and this reduced overlap correlates with a decline in behavioural discrimination of similar contexts.⁷⁹ Thus, consistent with our simulation results, an age-related reduction in neurogenesis leads to greater pattern separation, and at the same time, greater interference between similar memories. Further empirical studies are required to tease apart the roles of the young versus older DG neurons in sparse coding and pattern separation, in encoding versus retrieval, in memory and interference reduction at short (same session) versus long time scales, and in representing novel information and fine details versus more general information.

While our simulations indicate that in the model, young DGCs contribute to reduced pattern separation within the DG as a whole, it has been suggested that the young neurons could indirectly contribute to pattern separation amongst the mature DGC population, by generating increased inhibitory feedback over mature DGCs.⁷⁵ It also should be noted that some, but not all, behavioural studies have found that neurogenesis knockdown leads to improved performance on high interference memory tasks.⁷⁶ These studies appear to further contradict the theory that neurogenesis increases sparse coding, and that this sparse coding reduces memory interference. Thus, further empirical research is needed to fully characterize the functional significance of neurogenesis for mitigating memory interference.

Future modeling studies

The simulations reported in this chapter demonstrated how an RBM model could be used to model place cell fields, neurogenesis and the developmental trajectory of young DGCs with a hippocampal circuit. We also showed how our one-layer RBM model can be extended to a stacked, multi-layer RBM model of the full hippocampal circuit, but have not yet demonstrated the potential of this multi-layer model to simulate sequential behaviour. Future developments of the model will combine the more realistic inputs, i.e. grid cells and boundary vector cells, with the full multi-layer model including temporal associative learning in the CA3 region. This model should learn to encode place sequences. Moreover, in the top-down or generative mode, the model provides a natural account of both forward and reverse sequence replay. A related model has been developed by Fox, Prescott and colleagues based on particle filters.⁷¹⁻⁷³ Combining

temporal sequence learning with neurogenesis, within a single model, should enable the encoding of distinct event memories for overlapping temporal sequences.

References

1. Marr D. Simple Memory: A Theory for Archicortex. *Philos Trans R Soc B Biol Sci.* 1971;262(841):23-81. doi:10.1098/rstb.1971.0078.
2. McNaughton BL, Morris RGM. Hippocampal synaptic enhancement and information storage within a distributed memory system. *Trends Neurosci.* 1987;10(10):408-415. doi:10.1016/0166-2236(87)90011-7.
3. Treves A, Rolls ET. Computational constraints suggest the need for two distinct input systems to the hippocampal CA3 network. *Hippocampus.* 1992;2:189-200.
4. Hasselmo ME, Schnell E. Laminar selectivity of the cholinergic suppression of synaptic transmission in rat hippocampal region CA1: computational modeling and brain slice physiology. *J Neurosci.* 1994;14(6):3898-3914.
<http://www.scopus.com/inward/record.url?eid=2-s2.0-0028302804&partnerID=40&md5=018c11a327f4a4dd96ca4d671cfb034f%5Cnhttp://www.jneurosci.org/content/14/6/3898.full.pdf>.
5. Hasselmo ME, Wyble BP, Wallenstein G V. Encoding and retrieval of episodic memories: Role of cholinergic and GABAergic modulation in the hippocampus. *Hippocampus.* 1996;6(6):693-708. doi:10.1002/(SICI)1098-1063(1996)6:6<693::AID-HIPO12>3.0.CO;2-W.

6. McClelland JL, McNaughton BL, O'Reilly RC. Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychol Rev.* 1995;102(3):419-457.
7. O'Reilly RC, McClelland JL. Hippocampal conjunctive encoding, storage, and recall: Avoiding a trade-off. *Hippocampus.* 1994;4(6):661-682. doi:10.1002/hipo.450040605.
8. Jung MW, McNaughton BL. Spatial selectivity of unit activity in the hippocampal granular layer. *Hippocampus.* 1993;3(2):165-182.
9. Gothard KM, Hoffman KL, Battaglia FP, McNaughton BL. Dentate gyrus and ca1 ensemble activity during spatial reference frame shifts in the presence and absence of visual input. *J Neurosci.* 2001;21(18):7284-7292. doi:21/18/7284 [pii].
10. Chawla MK, Guzowski JF, Ramirez-Amaya V, et al. Sparse, environmentally selective expression of Arc RNA in the upper blade of the rodent fascia dentata by brief spatial experience. *Hippocampus.* 2005;15(5):579-586. doi:10.1002/hipo.20091.
11. Leutgeb JK, Leutgeb S, Moser M-B, Moser EI. Pattern separation in the dentate gyrus and CA3 of the hippocampus. *Science (80-).* 2007;315(5814):961-966. doi:10.1126/science.1135801.
12. Satvat E, Schmidt B, Argraves M, Marrone DF, Markus EJ. Changes in task demands alter the pattern of zif268 expression in the dentate gyrus. *J Neurosci.* 2011;31(19):7163-7167. doi:10.1523/JNEUROSCI.0094-11.2011.
13. Nakazawa K, Quirk MC, Chitwood RA, et al. Requirement for Hippocampal CA3 NMDA Receptors in Associative Memory Recall. *Science (80-).* 2002;297(July):211. doi:10.1126/science.1071795.
14. Neunuebel JP, Knierim JJ, Alme CB, et al. CA3 Retrieves Coherent Representations from

- Degraded Input: Direct Evidence for CA3 Pattern Completion and Dentate Gyrus Pattern Separation. *Neuron*. 2014;81(2):416-427. doi:10.1016/j.neuron.2013.11.017.
15. Wills TJ, Lever C, Cacucci F, Burgess N, O'Keefe J. Attractor dynamics in the hippocampal representation of the local environment. *Science* (80-). 2005;308:873-876.
 16. Tulving E. Memory and consciousness. *Can Psychol Can*. 1985;26(1):1-12. doi:10.1037/h0080017.
 17. Hassabis D, Kumaran D, Vann SD, Maguire EA. Patients with hippocampal amnesia cannot imagine new experiences. *Proc Natl Acad Sci*. 2007;104(5):1726-1731. doi:10.1073/pnas.0610561104.
 18. Rosenbaum RS, Gilboa A, Levine B, Winocur G, Moscovitch M. Amnesia as an impairment of detail generation and binding: Evidence from personal, fictional, and semantic narratives in K.C. *Neuropsychologia*. 2009;47(11):2181-2187. doi:10.1016/j.neuropsychologia.2008.11.028.
 19. Andelman F, Hoofien D, Goldberg I, Aizenstein O, Neufeld MY. Bilateral hippocampal lesion and a selective impairment of the ability for mental time travel. *Neurocase*. 2010;16(5):426-435. doi:10.1080/13554791003623318.
 20. O'Keefe J. Place units in the hippocampus of the freely moving rat. *Exp Neurol*. 1976;51:78-109.
 21. Ekstrom AD, Kahana MJ, Caplan JB, et al. Cellular networks underlying human spatial navigation. *Nature*. 2003;425(6954):184-188.
 22. Lee ACH, Yeung L-K, Barense MD. The hippocampus and visual perception. *Front Hum Neurosci*. 2012;6. doi:10.3389/fnhum.2012.00091.
 23. Barense MD, Henson RNA, Graham KS. Perception and Conception: Temporal Lobe

- Activity during Complex Discriminations of Familiar and Novel Faces and Objects. *J Cogn Neurosci*. 2011;23(10):3052-3067. doi:10.1162/jocn_a_00010.
24. Káli S, Dayan P. Off-line replay maintains declarative memories in a model of hippocampal-neocortical interactions. *Nat Neurosci*. 2004;7(3):286-294. doi:10.1038/nm1202.
 25. Nádasdy Z, Hirase H, Czurkó A, Csicsvari J, Buzsáki G. Replay and time compression of recurring spike sequences in the hippocampus. *J Neurosci*. 1999;19(21):9497-9507.
 26. Lee AK, Wilson MA. Memory of sequential experience in the hippocampus during slow wave sleep. *Neuron*. 2002;36(6):1183-1194. doi:10.1016/S0896-6273(02)01096-6.
 27. Schmidt-Hieber C, Jonas P, Bischofberger J. Enhanced synaptic plasticity in newly generated granule cells of the adult hippocampus. *Nature*. 2004;429(6988):184-187. doi:10.1038/nature02553.
 28. Snyder JS, Kee N, Wojtowicz JM. Effects of adult neurogenesis on synaptic plasticity in the rat dentate gyrus. *J Neurophysiol*. 2001;85(6):2423-2431.
 29. Mongiat LA, Espósito MS, Lombardi G, Schinder AF. Reliable Activation of Immature Neurons in the Adult Hippocampus. Reh TA, ed. *PLoS One*. 2009;4(4):e5320. doi:10.1371/journal.pone.0005320.
 30. Marin-Burgin A, Mongiat LA, Pardi MB, Schinder AF. Unique processing during a period of high excitation/inhibition balance in adult-born neurons. *Science (80-)*. 2012;335:1238-1242.
 31. Hinton GE. Training products of experts by minimizing contrastive divergence. *Neural Comput*. 2002;14(8):1771-1800. doi:10.1162/089976602760128018.
 32. Finnegan R, Becker S. Neurogenesis paradoxically decreases both pattern separation and

- memory interference. *Front Syst Neurosci.* 2015;136(9). doi:10.3389/fnsys.2015.00136.
33. Zhao C, Teng EM, Summers Jr RG, Ming G-L, Gage FH. Distinct morphological stages of dentate granule neuron maturation in the adult mouse hippocampus. *J Neurosci.* 2006;26(1):3-11. doi:10.1523/JNEUROSCI.3648-05.2006.
 34. Piatti VC, Ewell LA, Leutgeb JK. Neurogenesis in the dentate gyrus: carrying the message or dictating the tone. *Front Neurosci.* 2013;7. doi:10.3389/fnins.2013.00050.
 35. Temprana SG, Mongiat LA, Yang SM, et al. Delayed Coupling to Feedback Inhibition during a Critical Period for the Integration of Adult-Born Granule Cells. *Neuron.* 2015;85(1):116-130.
 36. Becker S. A computational principle for hippocampal learning and neurogenesis. *Hippocampus.* 2005;15(6):722-738. doi:10.1002/hipo.20095.
 37. Weisz VI, Argibay PF. A putative role for neurogenesis in neuro-computational terms: inferences from a hippocampal model. *Cognition.* 2009;112(2):229-240. doi:10.1016/j.cognition.2009.05.001.
 38. Gompertz B. On the Nature of the Function Expressive of the Law of Human Mortality, and on a New Method of Determining the Value of Life Contingencies. *Phil Trans Roy Soc London.* 1832;123:513-585.
 39. Luu P, Sill O, Gao L, Becker S, Wojtowicz JM, Smith D. The Role of Adult Hippocampal Neurogenesis in Reducing Interference. *Behav Neurosci.* 2012;126(3):381-391.
 40. Winocur G, Becker S, Luu P, Rosenzweig S, Wojtowicz JM. Adult hippocampal neurogenesis and memory interference. *Behav Brain Res.* 2012;227(2):464-469. doi:10.1016/j.bbr.2011.05.032.
 41. Clelland CD, Choi M, Romberg C, et al. A functional role for adult hippocampal

- neurogenesis in spatial pattern separation. *Science*. 2009;325(5937):210-213.
doi:10.1126/science.1173215.
42. Finnegan R. Masters Thesis - Computational Modelling of Adult Hippocampal Neurogenesis (MIT "Expat" License).
 43. O'Keefe J, Dostrovsky J. The hippocampus as a spatial map: preliminary evidence from unit activity in the freely-moving rat. *Brain Res*. 1971;34:171-175.
 44. Park E, Dvorak D, Fenton AA. Ensemble place codes in hippocampus: CA1, CA3, and dentate gyrus place cells have multiple place fields in large environments. *PLoS One*. 2011;6(7). doi:10.1371/journal.pone.0022349.
 45. Hafting T, Fyhn M, Molden S, Moser MB, Moser EI. 2005-02-01826A, Hafting et al. *Nature*. 2005;436(7052):1-8. doi:10.1038/nature03721.
 46. Boccara CN, Sargolini F, Thoresen VH, et al. Grid cells in pre- and parasubiculum. *Nat Neurosci*. 2010;13(8):987-994. doi:10.1038/nn.2602.
 47. Rolls ET, Stringer SM, Elliot T. Entorhinal cortex grid cells can map to hippocampal place cells by competitive learning. *Network*. 2006;17(4):447-465.
doi:10.1080/09548980601064846.
 48. Hartley T, Burgess N, Lever C, Cacucci F, O'Keefe J. Modeling place fields in terms of the cortical inputs to the hippocampus. *Hippocampus*. 2000;10(4):369-379.
doi:10.1002/1098-1063(2000)10:4<369::AID-HIPO3>3.0.CO;2-0.
 49. O'Keefe J, Burgess N. Geometric Determinants of the Place Fields of Hippocampal Neurons. *Nature*. 1996;381:425-428.
 50. Lever C, Burton S, Jeewajee A, O'Keefe J, Burgess N. Boundary Vector Cells in the Subiculum of the Hippocampal Formation. *J Neurosci*. 2009;29(31).

51. Shaw M. Modelling Place Cells Through Machine Learning by. 2012;(April).
52. Burgess N, Barry C, O'Keefe J. An oscillatory interference model of grid cell firing. *Hippocampus*. 2007;17(9):801-812. doi:10.1002/hipo.20327.
53. Hasselmo ME, Giocomo LM, Zilli EA. Grid cell firing may arise from interference of theta frequency membrane potential oscillations in single neurons. *Hippocampus*. 2007;17(12):1252-1271. doi:10.1002/hipo.20374.
54. Burgess N, Cacucci F, Lever C, O'Keefe J. Characterizing multiple independent behavioral correlates of cell firing in freely moving animals. *Hippocampus*. 2005;15(2):149-153. doi:10.1002/hipo.20058.
55. Anderson MI, Jeffery KJ. Heterogeneous modulation of place cell firing by changes in context. *J Neurosci*. 2003;23:8827-8835. doi:23/26/8827 [pii].
56. Bush D, Barry C, Burgess N. What do grid cells contribute to place cell firing? *Trends Neurosci*. 2014;37(3):136-145. doi:10.1016/j.tins.2013.12.003.
57. Byrne P, Becker S. A principle for learning egocentric-allothetic transformation. *Neural Comput*. 2008;20(3):709-737. doi:10.1162/neco.2007.10-06-361.
58. Levy WB. A Sequence Predicting {CA3} is a Flexible Associator That Learns and Uses Context to Solve Hippocampal-Like Tasks. *Hippocampus*. 1996;6:579-590.
59. Jensen O, Lisman JE. Hippocampal CA3 region predicts memory sequences: accounting for the phase precession of place cells. *Learn Mem*. 1996;3(2-3):279-287. doi:10.1101/lm.3.2-3.279.
60. Mehta MR. From synaptic plasticity to spatial maps and sequence learning. *Hippocampus*. 2015;25(6):756-762. doi:10.1002/hipo.22472.
61. Myers CE, Scharfman HE. Pattern separation in the dentate gyrus: a role for the CA3

- backprojection. *Hippocampus*. 2011;21(11):1190-1215. doi:10.1002/hipo.20828.
62. Williams R ~J., Zipser D. A Learning Algorithm for Continually Running Fully Recurrent Neural Networks. *Neural Comput*. 1989;1(2):270-280.
 63. Pineda F ~J. Generalization of Back Propagation to Recurrent and Higher Order Neural Networks. In: *Proceedings of IEEE Conference on Neural Information Processing Systems*. Denver, Colorado; 1987.
 64. Pearmutter B ~A. Learning state space trajectories in recurrent neural networks. *Neural Comput*. 1989;1(2):263-269.
 65. Hochreiter S, Schmidhuber J. Long short-term memory. *Neural Comput*. 1997;9(8):1735-1780. doi:doi:10.1162/neco.1997.9.8.1735.
 66. Maass W, Markram H. On the computational power of circuits of spiking neurons. *J Comput Syst Sci*. 2004;69(4):593-616. doi:http://dx.doi.org/10.1016/j.jcss.2004.04.001.
 67. Graves A. *Supervised Sequence Labelling with Recurrent Neural Networks*. Vol 385. Springer; 2012.
 68. Schmidhuber J, Wierstra D, Gomez FJ. Evolino: Hybrid Neuroevolution / Optimal Linear Search for Sequence Learning. In: Edinburgh: Proceedings of the 19th International Joint Conference on Artificial Intelligence (IJCAI); 2005:853-858.
 69. Taylor G, Hinton G, Roweis S. Modeling human motion using binary latent variables. *NIPS*. 2007;19:1345-1352.
 70. Becker S, Hinton GE. Caching and replay of place sequences in a Temporal Restricted Boltzmann Machine model of the hippocampus. In: Society for Neuroscience; 2007.
 71. Fox C, Prescott T. Hippocampus as unitary coherent particle filter. *Proc Int Jt Conf Neural Networks*. 2010. doi:10.1109/IJCNN.2010.5596681.

72. Fox C, Prescott T. Learning in a unitary coherent hippocampus. *Lect Notes Comput Sci (including Subser Lect Notes Artif Intell Lect Notes Bioinformatics)*. 2010;6352 LNCS(PART 1):388-394. doi:10.1007/978-3-642-15819-3_52.
73. Boorman LW, Damianou AC, Martinez-Hernandez U, Prescott TJ. Extending a Hippocampal Model for Navigation Around a Maze Generated from Real-World Data. In: ; 2015:441-452. doi:10.1007/978-3-319-22979-9_44.
74. Becker S. Neurogenesis and pattern separation: Time for a divorce. *to Appear Wiley Interdiscip Rev Cogn Sci*. 2016.
75. McAvoy K, Besnard A, Sahay A. Adult hippocampal neurogenesis and pattern separation in DG: a role for feedback inhibition in modulating sparseness to govern population-based coding. *Front Syst Neurosci*. 2015;9(120). doi:10.3389/fnsys.2015.00120.
76. Hvoslef-Eide M, Oomen CA. Adult neurogenesis and pattern separation in rodents: A critical evaluation of data, tasks and interpretation. *Front Biol (Beijing)*. 2016. doi:DOI: 10.1007/s11515-016-1406-2.
77. Seki T, Arai Y. Age-related production of new granule cells in the adult dentate gyrus. *Neuroreport*. 1995;6(18):2479-2482. doi:10.1097/00001756-199512150-00010.
78. Kuhn HG, Dickinson-Anson H, Gage FH. Neurogenesis in the dentate gyrus of the adult rat: Age-related decrease of neuronal progenitor proliferation. *J Neurosci*. 1996;16(6):2027-2033.
79. Marrone DF, Adams AA, Satvat E. Increased pattern separation in the aged fascia dentata. *Neurobiol Aging*. 2011;32(12):2317.e23-2317.e32. doi:10.1016/j.neurobiolaging.2010.03.021.
80. Brown TH, Johnston D. Voltage-clamp analysis of mossy fiber synaptic input to

hippocampal neurons. *J Neurophysiol.* 1983;50:487-507.

Figure 1: Traditional view of the hippocampus as a memorization and cued retrieval device, with pattern separation in the dentate gyrus (DG) and pattern completion in the CA3. The CA3 and CA1 regions receive both direct input from the entorhinal cortex (EC) via the perforant path, and indirect input through the trisynaptic circuit via the DG. The DG granule cells project to the CA3 field via mossy fiber synapses, which are few in number but are among the largest in the brain, such that only a few Mossy fiber synapses may be sufficient to activate a CA3 pyramidal cell⁸⁰. It has therefore been suggested that these terminals act as ‘‘detonator synapses,’’ so that during encoding, a sparse pattern of activation in the DG mandatorily causes a postsynaptic CA3 cell to fire^{2,3}. On the other hand, during retrieval, the CA3 recurrent collaterals and CA3-to-CA1 Shaffer collaterals may dominate in driving CA3 and CA1 cells to perform associative recall³⁻⁵. Another prominent source of input to the hippocampus, frequently left out of the standard model, is from the subiculum.

Figure 2: Sample activation of a grid cell input and corresponding dentate gyrus (DG) place cell activation in our Restricted Boltzmann Machine (RBM) model. Left: Example of firing pattern of a single grid cell used as input to the model. The red points show the grid cell activation, while the simulated rat’s path is shown in grey. Centre: Diagram of the RBM used to re-encode the grid cell firing patterns into place cell fields. The RBM model consisted of a set of visible and hidden units, fully inter-connected, with no within-layer connections. Right: Sample firing pattern of a single place cell generated by the model’s hidden units. Again, the red points show the place cell activation, while the simulated rat’s path is shown in grey.

Figure 3: A full hippocampal model consisting of stacked Conditional Restricted Boltzmann Machines (CRBMs). A combined CA3 & CA1 layer is trained on entorhinal cortex (EC) input, and conditioned on previous EC input. This CA3 & CA1 layer subsequently acts as the input to train the dentate gyrus (DG) layer via the bidirectional connections, representing the mossy fibres and CA3 backprojections. In this architecture, the EC is treated as conditional input to the DG layer, replacing the bidirectional weights in our previous RBM based DG models with unidirectional autoregressive weights. This architecture is based on one proposed by Becker and Hinton (2007) using Temporal Restricted Boltzmann Machines (TRBMs).

Figure 4: Performance of the multi-layer RBM hippocampal models with and without neurogenesis on across-session cued recall tests. The models were trained sequentially for 11 sessions, on 11 different groups of 90 patterns. Each group of patterns learned in a given session had high overlap with patterns learned in other sessions. The model was tested on noisy versions of these training patterns after each session to test proactive interference, and after all sessions had completed to test retroactive interference. (A) Proactive interference for cued recall accuracies during training. (B) Retroactive interference for cued recall accuracies on each group of patterns. (C) The relationship between post training recall accuracy and DG hidden unit activation overlap. (D) The distribution of post training accuracy over all sequentially learned groups of patterns.