Modeling the Effects of Neurogenesis on Interference and Pattern Separation for Proximal Similar Events

Caitlin Coyiuto

NEUR 335

Introduction

The function of the hippocampus has long been implicated in learning and memory, with different regions in the hippocampus playing distinctive roles in these complex processes. Regions of interest include the CA3, which may be involved in the formation of associative memories (Deng, Aimone, & Gage, 2010). An associative network is characterized by a process known as pattern completion, in which memory cues are capable of activating a set of neurons representative of the stored memory (Deng et al., 2010). Research has also investigated the role of the dentate gyrus (DG), which provides inputs to CA3 through mossy fibers, and may be responsible for preprocessing information before it is relayed to downstream regions (Jonas & Lisman, 2014). Furthermore, the DG appears to be capable of reducing the interference between similar memories, specifically through pattern separation, the process of primary interest in this paper (Deng et al., 2010).

Pattern separation is defined as a computational process that minimizes interference through the creation of distinct, non-overlapping neural codes (Deng et al., 2010), and the DG seems to be anatomically suitable for this process. Firstly, the DG possesses roughly five to ten times more neurons than its input, the entorhinal cortex (EC) (Amaral, Scharfman, & Lavenex, 2007), thereby presenting a large capacity for encoding multiple distinct neural codes. Additionally, dentate gyrus cells (DGCs) utilizes a sparse coding scheme, in which these neurons provide highly selective inputs to its downstream region, CA3 (Jung & McNaughton, 1993). This is due to a large degree of inhibition from neighboring interneurons, thereby facilitating the sparse DGC activity (Jung & McNaughton, 1993). Effectively, the properties of the DG are highly appropriate for pattern separation, and thereby facilitate the encoding of similar memories with minimal interference. This is especially vital to an associative network such as the CA3,

which the DG projects to. In associative networks, there is a risk that a single memory cue may activate multiple sets of neurons due to the nature of pattern completion, and may thereby result in a converging of multiple memories into a single, inappropriate representation (Deng et al., 2010). The DG helps minimize the storage of these inappropriate representations through pattern separation, which enables distinct, separate encoding patterns for different memories.

Research has endeavored to identify additional processes that may influence pattern separation in the DG – one such process includes neurogenesis, which may facilitate pattern separation. An empirical study on spatial discrimination with mice reveals that for items spaced close together, mice with ablated neurogenesis were impaired in a discrimination task (Clelland et al., 2009). The inverse also appears to be true for mice with elevated levels of neurogenesis, in which those with higher levels performed better on the spatial discrimination task (Creer, Romberg, Saksida, van Praag, & Bussey, 2010). Asides from assisting in pattern separation for proximal similar items, neurogenesis also may assist in pattern separation by minimizing retroactive and proactive interference for items placed distally, or across a longer period of time. Experiments on rats revealed how impaired neurogenesis was associated with proactive interference, in which learning a list of odor stimuli was impaired by exposure to a previous, highly similar list a few days ago (Luu et al., 2012). Similarly, for retroactive interference, rats with suppressed neurogenesis experienced impairments in retrieving previously learned information on a visual discrimination task (Winocur, Becker, Luu, Rosenzweig, & Wojtowicz, 2012).

Although multiple computational models have been produced as means to understand the role of neurogenesis in pattern separation in the hippocampus, only a single model has investigated neurogenesis with respect to similar items learned distally or proximal to each other

(Finnegan & Becker, 2015). This paper attempts to produce a simplification of the model by Finnegan and Becker (2015), in which the effects of neurogenesis on pattern separation are investigated for similar items learned within a short time span. The current model utilizes a version of the Restricted Boltzmann Machine (RBM) to represent the network between EC inputs to DGCs. As means to assess the effects of neurogenesis, the sparsity and learning rate constants were manipulated to reflect young or mature DGCs. Effects of proactive interference and degree of pattern separation were measured respectively by: a) accuracy, the model's ability to reconstruct new input memory patterns after learning a set of similar patterns, and b) overlap, or specifically, the degree of pattern overlap stored within the DGC units.

Methods

As Restricted Boltzmann Machines (RBMs) have been commonly used to model neural network learning in other applications (Hinton, Osindero, & Teh, 2006), an RBM was used as means to represent the connections between the EC and DG. The RBM consists of two layers, the first being a series of visible units and the second a series of hidden units, with each unit having reciprocal connections with the other layer. In order to preserve the ratio between EC units and DGCs, the current model used 10 visible units and 50 hidden units. Learning is modeled by updating the weights, or connections between each EC and DGC unit after each presentation of a pattern. Testing is represented by presenting a new set of similar patterns, but no longer updating the weights after each pattern to indicate that learning has already occurred. After testing, accuracy is computed in terms of the difference between the initial pattern and the reconstruction by the visible units, along with calculating the degree of overlap between the patterns stored in the hidden units. In order to assess the effects of neurogenesis, two RBMs

were used, each with manipulated sparsity constraints as means to represent young (neurogenesis model) or mature (non-neurogenesis) DGCs. To evaluate the effects of proactive interference for items learned within a short time frame, all neurons in the model remained young or old to simulate no long passage of time.

Data Initalization

The initial data for the visible unit layer was represented through 10 similar binary patterns, and thereby effectively acted as the patterns that must be learned by the model. Additionally, the use of binary values simply translated to a given unit being activated (1) or inactivated (0). In order to compute a series of similar patterns, an initial pattern of 10 values in length was randomly generated to have 30% of its values set to ones. This pattern then had 20% of its values reset to zero randomly 10 times, and thereby generating the series of 10 similar binary patterns. These 10 patterns are analogous to the "protoypes" defined in the model by (Finnegan & Becker, 2015).

To initialize the values for the weights, the matrix was populated by taking random values from a normal distribution with a sigma of 0.0001 and a mean of 0.0 (Finnegan & Becker, 2015).

RBM Training

To model how the artificial network learned these series of patterns, the weights between the visible and hidden units were updated after each presentation of a pattern. The update equation utilizes a constrastive divergence learning procedure (Finnegan & Becker, 2015; Hinton, 2002) as provided below:

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

To compute the values of the hidden and visible units after each pattern presentation, the following equations was used using a brief Gibbs sampling rule (Finnegan & Becker, 2015) :

$$p(h_j = 1 \mid v) = \sigma(\sum_i v_i w_{ij}) \quad (2)$$

$$p(v_i = 1 \mid h) = \sigma(\sum_i h_i w_{ij}) \quad (3)$$

Equation (2) produces the values for the hidden units by calculating the probability that a given hidden unit is activated – this is dependent on the values of the visible units and weights for that specific unit. To generate the reconstruction of the patterns, or the visible unit reconstruction, the hidden unit data generated in Equation (2) is used in Equation (3) to similarly compute probabilities that a given visible reconstruction unit is activated. Finally, the visible reconstruction data are used back into Equation (2) as means to obtain the hidden unit reconstructions. Once the hidden data, visible and hidden reconstructions were computed for a given pattern, the weights of each of these units were updated using Equation (1). This procedure was repeated 10 times, for each of the 10 input patterns.

Sparsity Constants for Differentiating Neurogenesis and Non-Neurogenesis RBMs

In order to assess the effects of neurogenesis and non-neurogenesis on pattern separation and interference, sparsity constants were manipulated for their respective RBMs. To represent the neurogenesis RBM, the model was simplified from that of Finnegan and Becker's (2015), in that all neurons in the RBM were young DGCs, while that in the non-neurogenesis RBM were mature DGCs. Two variables, learning rate (E) and sparsity cost (cost) were manipulated between these RBMs, to reflect the characteristics of these young and old neurons, as shown below in an updated version of Equation (1):

$$\Delta W_{ij} = E((v_i h_j)_{data} - (v_i h_j)_{recon})) - (\lambda \times W_{ij}) - cost \times (q\text{-}p) \quad (4)$$

Young DGCs tend to have low input specificity, and high synaptic plasticity, while mature DGCs are highly selective but with low plasticity (McAvoy, Besnard, & Sahay, 2015). Firstly, the value of learning rate (E), effectively reflective of plasticity, was set at 0.1 for the non-neurogenesis and 0.3 for the neurogenesis RBM. For sparsity cost, a value of 0.9 was set for the non-neurogenesis RBM to enforce the selective, and low degree of activation found in mature DGCs. For the neurogenesis RBM, cost was set at 0, to reflect the young neurons' low selectivity, and greater excitability. The cost term was additionally multiplied by the difference between q, the mean activation of the hidden units, and p, which was set at 0.05. These parameters were taken from Finnegan and Becker's (2015) model.

RBM Testing

To assess the performance of both models with respect to items learned within a short time frame, and to identify whether learning new items would impact reconstruction for older items, computations for hidden data, visible and hidden reconstructions were performed using Equations (2) and (3) again for a series of 10 additional similar patterns. These testing patterns originated from the same core pattern from the first 10 training patterns, so were therefore each 20% different from the original 10 patterns. Additionally, the weights were no longer being updated after each presentation of the new patterns, to represent how the model had now finished learning and was going to be tested on its ability to reconstruct older patterns.

Evaluating Model Performance

After completion of the testing phase, degree of overlap and the accuracy of the reconstructions for the testing patterns were computed. The following equations from Finnegan and Becker's (2015) model were used to compute percent accuracy between the initial patterns and the reconstructions during the testing phase:

$$D(V_{data}, V_{recon}) = \sum_{i=1}^{n} | (V_{data,i} - V_{recon,i}) | \quad (5)$$

$$M(V_{data}, V_{recon}) = 1 - D(V_{data}, V_{recon})/l \quad (6)$$

A percentage was calculated for each of the test patterns presented, yielding a data set of 10 points.

Hidden unit overlap of the testing phase was calculated in a similar fashion to that as accuracy. For each pattern: 1) A sum was taken of the number of active units (1's) and subtracted from every other pattern's sum, 2) This absolute difference was divided by the total number of hidden units (n = 100), and finally, 3) These differences for each pattern with every other pattern were summed, and divided by the number of patterns (n = 10). Each pattern therefore had a single value representative of its degree of overlap with all other patterns in the set.

Results

Computations for accuracy reveal that the neurogenesis RBM is able to successfully reproduce all test patterns after learning a set of similar patterns within a short time frame (Figure 1A). This is shown by the 100% accuracy of the visible units' reconstruction of the inputs across all 10 patterns. In contrast, the non-neurogenesis model performed poorer than that of the neurogenesis RBM. This suggests how proactive interference, or the learning of the first set of 10 patterns, hindered reconstruction accuracy of a new set of patterns for the non-neurogenesis model, but not for the neurogenesis trials.

When assessing degree of hidden unit overlap (Figure 1B), the non-neurogenesis RBM yields lower values of overlap, indicative of greater pattern separation. In contrast, the neurogenesis RBM produces greater overlap values, suggestive of reduced pattern separation. These findings paradoxically demonstrate how despite lower levels of pattern separation, the neurogenesis RBM benefits from greater reconstruction accuracy, while the non-neurogenesis model yields greater levels of pattern separation, but lower reconstruction accuracies.

Discussion

Paradoxical Decrease in Interference and Pattern Separation

Our results, and that of Finnegan and Becker (2015), demonstrate that when presented with similar items over a short period of time, neurogenesis appears to paradoxically reduce proactive interference, yet yields a lower degree of pattern separation. Studies investigating feedback inhibition suggest how young DGCs may exert such a great degree of inhibition to their mature counterparts, that the greater degree of excitability elicited by the presence of these young neurons may be counteracted by this inhibition (McAvoy et al., 2015). Young DGCs facilitate this inhibition through mossy fiber cells and hilar interneurons, which would, by suppressing mature DGC activity, increase the degree to which these older neurons are sparsely activated. Experiments done in mice reveal how greater levels of young DGCs result in a reduced spread of voltage sensitive dye (VSD) and activation strength in the DG after electrical stimulation (Ikrar et al., 2013). In contrast, mice with reduced levels of young DGCs produced greater levels of excitability in the DG (Ikrar et al., 2013). Further analysis also revealed how there were a greater number of synapses between DGCs and hilar interneurons for mice with greater levels of young DGCs, suggestive of how these interneurons may be the underlying

mechanism to the increased sparse mature DGC activity (Ikrar et al., 2013). In conclusion, this implies that although new, young DGCs do not help facilitate separation, they compensate for this by greatly inhibiting mature DGCs, and thereby greatly reducing interference when memories are encoded.

Although inhibition from young DGCs may explain heightened reconstruction accuracy, the current model cannot concretely support this claim. Due to simplifications made to the original model by Finnegan and Becker (2015), the neurogenesis RBM consisted of solely young DGCs, while the non-neurogenesis RBM consisted of mature DGCs. As there were no mature neurons in the neurogenesis model to begin with, the paradoxical relationship between reduced interference and pattern separation cannot be explained by inhibitory feedback from young DGCs onto old DGCs in the current model. The implications of these changes done on the current model will be discussed in the next section.

Simplifications of the Current Model

The current model made a significant amount of simplifications from that of Finnegan and Becker's (2015), particularly in the modeling of age. Firstly, although like the original model, which simulated the learning of proximal similar items by keeping all plasticity and sparsity values constant, the current model made the assumption that the neurogenesis RBM would contain solely new DGC units, while the non-neurogenesis model exclusively mature DGCs. This was not the case with Finnegan and Becker's model, to which the DGC population in the neurogenesis model contained both young and old neurons. This may have been the reason to why a ceiling effect was achieved in the neurogenesis model for reconstruction accuracy, and not for its non-neurogenesis counterpart.

Furthermore, neuron growth was not simulated in the current model. Finnegan and Becker's (2015) model utilized the Gompertz function (Gompertz, 1825) as means to model the temporal growth of the neurons. This function was thereby implemented into the computations for plasticity and sparsity cost, in which plasticity values would be greater and cost values would be lower for younger neurons. In their trials for the learning of proximal similar items, sparsity cost and plasticity values were kept constant across presentations of different pattern stimuli, yet were not constant across all neurons in the population. Additionally, since age was kept constant, the current model also omitted the use of a sparse connectivity parameter, which serves to manipulate a greater degree of connectivity for older DGCs, and a sparse degree of connectivity for younger DGCs. In Finnegan and Becker's (2015) model, this connectivity constraint was dependent on the age of neurons, and would thereby affect each weight for the neuron population. Given that the current model does not take into account the degree of connectivity between old and young neurons, and assumes that all neurons are of the same age in either the neurogenesis or non-neurogenesis RBM, the current model is not an accurate representation of a neural population.

An additional weakness to the current model was that neuronal turnover was not implemented. The original model utilized an additional parameter to identify older units that had poor differentiation between hidden units – specifically, in that hidden units with similar weights to all its visible units would be less effective in differentiating patterns. Finnegan and Becker's (2015) model penalized 5% of the neurons depending on their age and differentiation capacities by simulating apoptosis, to which these neurons would be reset back to new, young DGCs. Although this may have not been a great issue when testing a model's ability to retrieve

memories across a short period of time, it is nonetheless another integral process excluded in the current model.

Finally, the current model only tests for the effects of proactive interference, and for a single group of similar patterns at a very small pattern and sample size. The number of visible units used in the current model's RBMs was 10, and the number of hidden units was set at 50. Finnegan and Becker's (2015) model utilized 200 visible units, 1000 hidden units, and 11 groups consisting of 90 patterns each. This may have also been the reason why a ceiling effect was observed in the current model's neurogenesis RBM, as a size of 10 for a given pattern may not have provided enough variation when 20% of the values were randomly reset – this meant that 2 values in a given pattern were reset to 0, and since the length of a pattern was only 10, it is highly probable that the 2 reset values were 0's in the first place. Additionally, due to Finnegan and Becker's (2015) vastly greater sample and test sizes, testing of proactive interference was performed, in which after learning each group of patterns, visible unit reconstructions could be immediately assessed. Fundamentally, the original model was capable of testing the effects of neurogenesis on pattern separation and interference on a much larger, and comprehensive scale than the current model.

Limitations and Future Directions

Considerations for future computational models of neurogenesis in the hippocampus include: a) the form of the network model, in this case, an RBM, b) the use of an addition rather than replacement model, and c) integrating other pathways in the DG, such as the hilar interneurons. First, the RBM utilizes reciprocal connections between its set of visible and hidden units; however, in the DG circuitry, these connections are not reciprocal in that DGCs do not

project back to the EC. Although utilizing an RBM provides benefits to modeling a neural network, in that it is capable of utilizing the data it trained on to generate new sets of data (Finnegan and Becker, 2015), it does not accurately portray the circuitry that it is modeling.

Secondly, the model proposed by Finnegan and Becker (2015) utilizes a replacement model, in which old DGCs are removed and replaced with new neurons. This method was suggested to actually enhance pattern separation, as the replacement helps facilitate that a set of DGCs is responsible for a certain memory (Deng et al., 2010). In contrast, the additive model assumes that the cells continue to age even as new neurons continue to be added (Deng et al., 2010). An additive approach may be able to model temporal properties of neurogenesis on pattern separation more realistically, as the effects of young DGCs on pattern separation are noticeable after days or weeks (Deng et al., 2010).

Finally, considering how Finnegan and Becker's (2015) model produced the paradoxical result of reduced pattern separation and interference in their neurogenesis RBM, it would be beneficial to model the other pathways involved in the DG, specifically that of the hilar interneurons, which send inhibitory projections back to the DG. If greater levels of young DGCs influence the degree of inhibitory inputs from the hilar interneurons to the DG (Ikrar et al., 2013), factoring the role of these interneurons into a computational model may help elucidate whether these inhibitory connections are responsible for the paradoxical findings between pattern separation and interference.

Figures

VARIABLES	DEFINITIONS/PARAMETERS
Hidden Data (h_{data})	Binary matrix representing DGCs.
Hidden Reconstruction (h_{recon})	Reconstruction of hidden data.
l	Length of Vdata or Vrecon.
Learning Rate (E)	Normalized to lie between 0.1-0.3. *For non-neurogenesis RBM: 0.1* *For neurogenesis RBM: 0.3*
q	Mean of hidden data activation.
Sigmoid Activation Function (σ)	Used to shift probabilities of activation. $\sigma = 1/(1+e^{-x})$.
Sparsity Cost (cost)	Normalized to lie between 0-0.9. *For non-neurogenesis RBM: 0.9* *For neurogenesis RBM: 0*
Target activation probability (p)	Probability of activation for a unit hidden unit. *Set to 0.05*.
Visible Data (v_{data})	Binary matrix representing EC input.
Visible Reconstruction (v_{recon})	Reconstruction of visible data.
Weight (w)	Matrix representing synaptic weights between each EC and DGC unit.
Weight Decay (λ)	Weight decay per hidden unit. *For non-neurogenesis RBM: 0* *For neurogenesis RBM: 0.005*
EQUATIONS	
(1) $\Delta W_{ij} = E((v_i h_j)_{data} - (v_i h_j)_{recon}))$	
(2) $p(h_j = 1 \mid v) = \sigma(\sum_i v_i w_{ij})$	
(3) $p(v_i = 1 \mid h) = \sigma(\sum_i h_i w_{ij})$	
(4) $\Delta W_{ij} = E((v_i h_j)_{data} - (v_i h_j)_{recon})) - (\lambda \times W_{ij}) - cost \times (q-p)$	
(5) $D(V_{data}, V_{recon}) = \sum_{i=1}^{n} \mid (V_{data,i} - V_{recon,i}) \mid$	
(6) $M(V_{data}, V_{recon}) = 1 - D(V_{data}, V_{recon})/l$	

Table 1: Variables and equations used within the model.

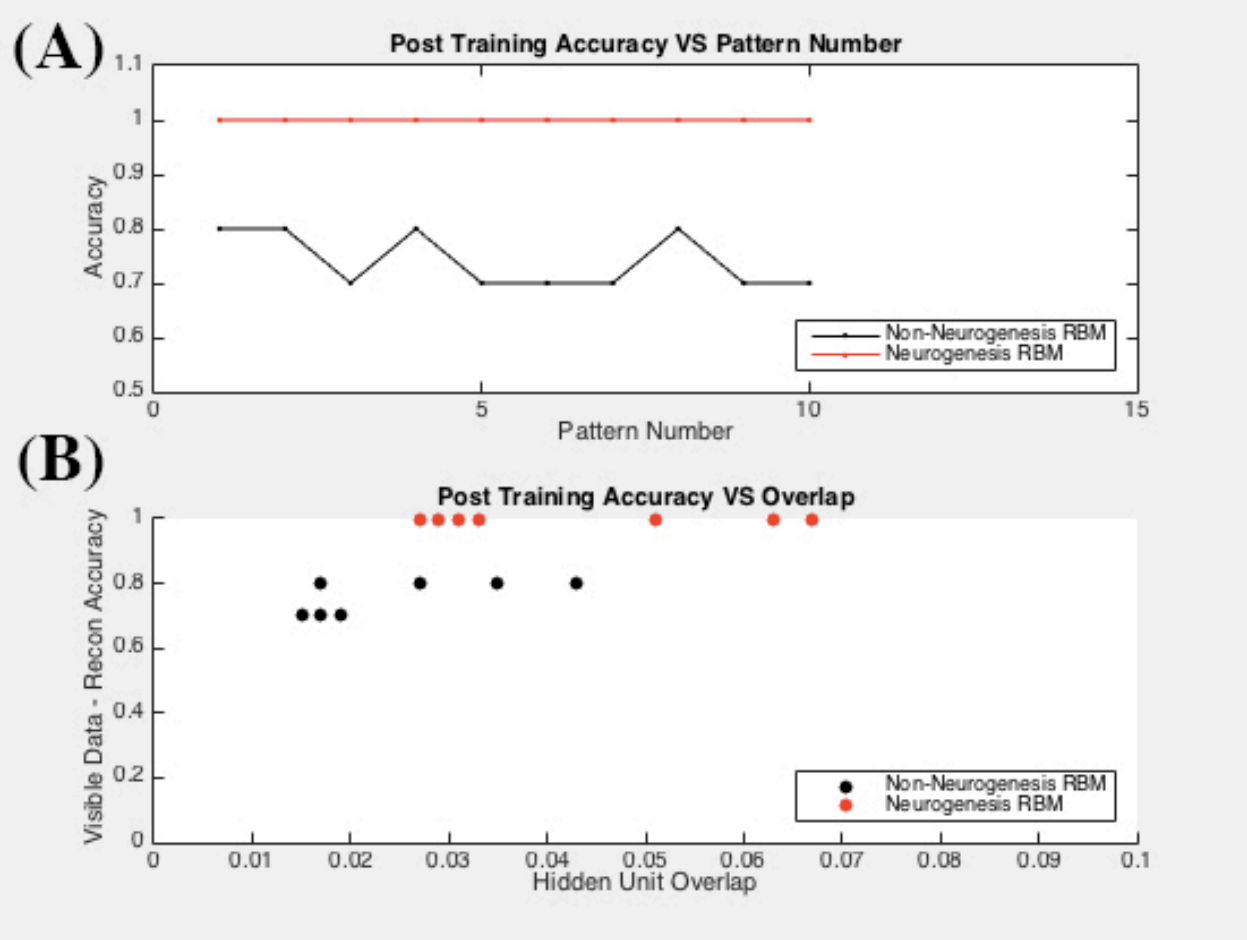

Figure 1: Performance of RBM models with or without neurogenesis for proximal similar patterns. Both models were first trained on 10 patterns that varied 20% from each other, and then were tested on an additional 10 patterns that also differed by 20% from each other. **(A)** Demonstrates accuracy between the test patterns and the visible units' reconstructions for each test pattern. **(B)** Demonstrates accuracy against degree of hidden unit overlap for each pattern. Each coordinate represents each pattern presented to the models. Note that some coordinates overlap each other.

Bibliography

Amaral, D. G., Scharfman, H. E., & Lavenex, P. (2007). The dentate gyrus: fundamental neuroanatomical organization (dentate gyrus for dummies). *Progress in Brain Research, 163*, 3-22. doi: 10.1016/S0079-6123(07)63001-5

Clelland, C. D., Choi, M., Romberg, C., Clemenson, G. D., Jr., Fragniere, A., Tyers, P., . . . Bussey, T. J. (2009). A functional role for adult hippocampal neurogenesis in spatial pattern separation. *Science, 325*(5937), 210-213. doi: 10.1126/science.1173215

Creer, D. J., Romberg, C., Saksida, L. M., van Praag, H., & Bussey, T. J. (2010). Running enhances spatial pattern separation in mice. *Proc Natl Acad Sci U S A, 107*(5), 2367-2372. doi: 10.1073/pnas.0911725107

Deng, W., Aimone, J. B., & Gage, F. H. (2010). New neurons and new memories: how does adult hippocampal neurogenesis affect learning and memory? *Nature Reviews: Neuroscience, 11*(5), 339-350. doi: 10.1038/nrn2822

Finnegan, R., & Becker, S. (2015). Neurogenesis paradoxically decreases both pattern separation and memory interference. *Front Syst Neurosci, 9*, 136. doi: 10.3389/fnsys.2015.00136

Gompertz, B. (1825). On the Nature of the Function Expressive of the Law of Human Mortality, and on a New Mode of Determining the Value of Life Contingencies. *Philosophical Transactions of the Royal Society of London, 115*, 513-583.

Hinton, G. E. (2002). Training products of experts by minimizing contrastive divergence. *Neural Computation, 14*(8), 1771-1800. doi: 10.1162/089976602760128018

Hinton, G. E., Osindero, S., & Teh, Y. W. (2006). A fast learning algorithm for deep belief nets. *Neural Computation, 18*(7), 1527-1554. doi: 10.1162/neco.2006.18.7.1527

Ikrar, T., Guo, N., He, K., Besnard, A., Levinson, S., Hill, A., . . . Sahay, A. (2013). Adult neurogenesis modifies excitability of the dentate gyrus. *Front Neural Circuits, 7*, 204. doi: 10.3389/fncir.2013.00204

Jonas, P., & Lisman, J. (2014). Structure, function, and plasticity of hippocampal dentate gyrus microcircuits. *Front Neural Circuits, 8*, 107. doi: 10.3389/fncir.2014.00107

Jung, M. W., & McNaughton, B. L. (1993). Spatial selectivity of unit activity in the hippocampal granular layer. *Hippocampus, 3*(2), 165-182. doi: 10.1002/hipo.450030209

Luu, P., Sill, O. C., Gao, L., Becker, S., Wojtowicz, J. M., & Smith, D. M. (2012). The role of adult hippocampal neurogenesis in reducing interference. *Behavioral Neuroscience, 126*(3), 381-391. doi: 10.1037/a0028252

McAvoy, K., Besnard, A., & Sahay, A. (2015). Adult hippocampal neurogenesis and pattern separation in DG: a role for feedback inhibition in modulating sparseness to govern population-based coding. *Front Syst Neurosci, 9*, 120. doi: 10.3389/fnsys.2015.00120

Winocur, G., Becker, S., Luu, P., Rosenzweig, S., & Wojtowicz, J. M. (2012). Adult hippocampal neurogenesis and memory interference. *Behavioural Brain Research, 227*(2), 464-469. doi: 10.1016/j.bbr.2011.05.032