HIPPOCAMPUS 15:722–738 (2005)

A Computational Principle for Hippocampal Learning and Neurogenesis

Suzanna Becker

ABSTRACT: In the three decades since Marr put forward his computational theory of hippocampal coding, many computational models have been built on the same key principles proposed by Marr: sparse representations, rapid Hebbian storage, associative recall and consolidation. Most of these models have focused on either the CA3 or CA1 fields, using "off-the-shelf" learning algorithms such as competitive learning or Hebbian pattern association. Here, we propose a novel coding principle that is common to all hippocampal regions, and from this one principal, we derive learning rules for each of the major pathways within the hippocampus. The learning rules turn out to have much in common with several models of CA3 and CA1 in the literature, and provide a unifying framework in which to view these models. Simulations of the complete circuit confirm that both recognition memory and recall are superior relative to a hippocampally lesioned model, consistent with human data. Further, we propose a functional role for neurogenesis in the dentate gyrus (DG), namely, to create distinct memory traces for highly similar items. Our simulation results support our prediction that memory capacity increases with the number of dentate granule cells, while neuronal turnover with a fixed dentate layer size improves recall, by minimizing interference between highly similar items. ©2005 Wiley-Liss, Inc.

KEY WORDS: computational model; paired associate learning; episodic memory; neurogenesis

INTRODUCTION

The hippocampus has been called the "darling of neuroscientists" (Churchland and Sejnowski, 1994) because of its key role in memory, coupled with its unique anatomical and electrophysiological characteristics. Individuals suffering damage to the hippocampal complex (HPC) and surrounding medial temporal lobe structures have severe anterograde and retrograde amnesia for episodic events, although other forms of learning and memory—semantic, perceptual, procedural, and simple forms of conditioning—are spared (for a review, see e.g., Moscovitch, 1982). The highly specialized coding capabilities of this region were revealed most strikingly by the discovery of hippocampal place cells (O'Keefe and Dostrovsky, 1971) and the subsequent discovery of hippocampal long-term potentiation (LTP)—evidence for rapidly induced and long-lasting plasticity in this region (Bliss and Lomo, 1973).

Department of Psychology, Neuroscience, and Behavior, McMaster University, Ontario, Canada
Grant sponsor: Natural Sciences and Engineering Research Council of Canada.
Correspondence to: Suzanna Becker, Department of Psychology, Neuroscience, and Behavior, McMaster University, Bldg 34, Rm 312, 1280 Main St. West, Hamilton, ON, Canada L8S 4K1.
E-mail: becker@mcmaster.ca
Accepted for publication 9 May 2005
DOI 10.1002/hipo.20095
Published online 28 June 2005 in Wiley InterScience (www.interscience.wiley.com).

In 1971, Marr put forward a highly influential theory of hippocampal coding (Marr, 1971), redescribed in a very cogent review paper by Willshaw and Buckingham (1990). Central to Marr's theory were the notions of a rapid, temporary memory store mediated by sparse activations and Hebbian learning, an associative retrieval system mediated by recurrent connections, as well as a gradual consolidation process by which new memories would be transferred into a long-term neocortical store. In the decades since the publication of Marr's computational theory, many researchers have built on these ideas and have simulated memory formation and retrieval in Marr-like models of the hippocampus. Although several recent models have simulated the complete hippocampal circuit (see Discussion), for the most part, modelers have focused on either the CA3 or CA1 fields, using variants of Hebbian learning, for example, competitive learning in the dentate gyrus (DG) and CA3 (Rolls, 1989; McClelland et al., 1995; Hasselmo et al., 1996; McClelland and Goddard, 1996; O'Reilly and Rudy, 2001), Hebbian auto-associative learning (Marr, 1971; McNaughton and Morris, 1987; Rolls, 1989; Treves and Rolls, 1992; Kali and Dayan, 2000b; O'Reilly and Rudy, 2001) or temporal associative learning (Levy, 1996; Gerstner and Abbott, 1997; Wallenstein and Hasselmo, 1997; August and Levy, 1999; Stringer et al., 2002) in the CA3 recurrent collaterals, and Hebbian heteroassociative learning between entorhinal cortex (EC)-driven CA1 activity and CA3 input (Hasselmo and Schnell, 1994; Hasselmo et al., 1996) or between EC-driven and CA3-driven CA1 activity at successive points in time (Levy et al., 1990). In this paper, we propose a principled way of deriving learning equations for all pathways within the hippocampal circuit, by assuming that all regions in the hippocampus are optimizing a common objective function. When combined with local neuroanatomical and physiological constraints, our learning principle leads to surprisingly simple and biologically feasible learning rules (Becker et al., 1999).

Of fundamental importance for computational theories of hippocampal coding is the striking finding of neurogenesis in the adult hippocampus. Although there is now a large literature on neurogenesis in the DG, and it has been shown to be important for at least one form of hippocampal-dependent learning (Shors et al., 2001), surprisingly few attempts have been made to reconcile this phenomenon with theories of hippocampal memory formation. Based on evidence that neurogenesis peaks in early adulthood and declines throughout

life, Kempermann (2002) has suggested that the gradual addition of new neurons into the existing network of the DG could allow the hippocampus to deal with novelty and the concomitant evolution of cortical representations. Nottebohm (2002) suggests that the newly born neurons may be recruited preferentially for storing new memories, thereby protecting old memories from interference. Consistent with this hypothesis, Wiskott, Rasch, and Kempermann (2005) demonstrated in a simple neural network model that the addition of highly plastic new neurons does effectively prevent new learning from interfering catastrophically with older memories. Conversely, Feng et al. (2001) proposed that neurogenesis is important for clearing out older memories once they are consolidated, and several modelers have demonstrated in abstract neural network models that neuronal turnover improves acquisition by helping to discard older memories (Chambers et al., 2004; Deisseroth et al., 2004). Here, we suggest a somewhat different role for new neurons in the DG, namely, in the generation of novel codes for highly similar events. For example, if I drive to work along the same route and park my car in the same parking lot every day, how am I able to recall where I parked today as distinct from the numerous other occasions in which I parked my car? Similarly, we often recollect vivid episodes in memory, and can remember having recollected an event, as distinct from one's memory of the event itself. Here, it is proposed that gradual changes in the internal code of the dentate layer could facilitate the formation of distinct representations for these highly similar episodes.

The above view of neuronal turnover in the DG could also explain why spaced learning is more effective than massed learning for hippocampally dependent tasks. For example, retention in the Morris water maze is superior with spaced learning compared with massed learning, but only on long-term rather than immediate retention tasks (Commins et al., 2003) and in adult but not adolescent rats (Spreng et al., 2002). The standard explanation of the spaced learning effect is that it permits time for consolidation into long-term storage between learning sessions; this would apply to both hippocampally dependent and nonhippocampally dependent learning. We suggest that in addition to this, for hippocampally dependent tasks, spaced learning trials could permit time to develop representations in the DG with greater variability, resulting in multiple overlapping memory traces of an event, and facilitating the recall of details of the event via multiple routes. This view is consistent with the multiple trace theory of Nadel, Moscovitch and colleagues (Nadel et al., 2000; Rosenbaum et al., 2001). We would also predict that the benefit of spaced learning increases with age, as neurogenesis declines. Having a fully implemented model of the complete hippocampal circuit allows us to test the feasibility of this hypothesis, and more generally, to explore basic questions about the consequences of neurogenesis for memory.

A NEW MODEL OF HIPPOCAMPAL LEARNING AND NEUROGENESIS

In this section, we propose a novel model of hippocampal coding. Unlike most previous approaches that have combined a variety of "off-the-shelf" learning principles for different pathways, our approach is to first propose a novel coding principle that is common to all hippocampal regions, and from this one principle, derive learning rules. By incorporating the key anatomical features of the hippocampus as architectural constraints, and the sparse activation levels and high degree of plasticity as processing constraints, we derive learning rules for each of the major pathways within the hippocampus. First, we briefly outline our proposed principle for learning in the hippocampus. We then review some of the important anatomical and electrophysiological data that we draw on to constrain connectivity and coding in the model. In the remainder of this section, we then describe some of the more general details of the model simulated here, namely, the architecture and activation function, and how neurogenesis was simulated. In the next section, we describe a series of simulations of the model's pattern recognition and recall capabilities.

The Learning Principle

The key contribution of the work described here is to propose a novel principle for learning in the hippocampus. We postulate that every region in the hippocampus adapts its synaptic connections according to the same general learning principle. Specifically, we propose that each hippocampal layer should form a neural representation that could be transformed in a simple manner—i.e., linearly—to reconstruct the original activation pattern in the EC. With the addition of biologically plausible processing constraints regarding connectivity, sparse activations, and two modes of neuronal dynamics during encoding versus retrieval, this results in very simple Hebbian learning rules. The key assumptions are as follows (see Appendix for details):

1. During encoding, dentate granule cells are active whereas during retrieval they are relatively silent.
2. During encoding, activation of CA3 pyramidals is dominated by the very strong mossy fiber inputs from dentate granule cells.
3. During retrieval, activation of CA3 pyramidals is driven by direct perforant path inputs from the EC combined with time-delayed input from CA3 via recurrent collaterals.
4. During encoding, activation of CA1 pyramidals is dominated by direct perforant path inputs from the EC.
5. During retrieval, CA1 activations are driven by a combination of perforant path inputs from the EC and Shaffer collateral inputs from CA3.

By combining our learning principle with the above constraints, we obtain Hebbian learning rules for the direct (monosynaptic) pathways from the EC to each hippocampal region, a temporal Hebbian associative learning rule for the CA3 recurrent collateral connections, and interestingly, a form of heteroassociative learning postulated by Hasselmo et al. (1996) for the Shaffer collaterals (the projection from CA3 to CA1).

Why is it that the constant turnover of neurons in the DG, and hence the constant rewiring of the hippocampal memory

circuit, does not interfere with the retrieval of old memories? The answer to this question comes naturally from our assumptions about neuronal dynamics during encoding vs. retrieval. New neurons are added only to the DG, and the DG drives activation in the hippocampal circuit only during encoding, not during retrieval. Thus, the new neurons contribute to the formation of distinctive codes for novel events, but not to the associative retrieval of older memories.

Anatomical Background

The unique multilayered circuitry of the hippocampus, summarized in Figure 1, has intrigued computational neuroscientists interested in unraveling how the hippocampus achieves its unique encoding capabilities. The hippocampus communicates with the rest of the brain via the para-hippocampal region (PHR), a major convergence zone reciprocally connected with widespread cortical and subcortical areas (Witter et al., 1989; Amaral et al., 1990). Within the PHR, the EC is the major input and output region for the hippocampus.[1] In the hippocampus proper, activity passes in turn through the DG and CA3 and CA1 fields—the so-called "trisynaptic circuit"—and back to the EC. The principal cells of the DG greatly outnumber those of the EC (by a factor of five in the rat (Amaral et al., 1990)) while having much lower activity levels (below 0.5 Hz (Barnes et al., 1990; Jung and McNaughton, 1993)). Thus, mapping from the EC to the DG results in a high-dimensional, extremely sparse neural code, likely due to the unique network of principal neurons and inhibitory neurons in the DG and adjacent hilar region (Acsady et al., 2000). The DG in turn projects to the CA3 field via mossy fiber synapses; these synapses are few in number, but are found close to the cell body and are among the largest in the brain, and the activation of only a few Mossy fiber synapses is likely sufficient to activate a CA3 pyramidal cell (Brown and Johnston, 1983). 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 (McNaughton and Morris, 1987; Treves and Rolls, 1992). The CA3 pyramidal cells send a dense and widespread projection of recurrent collaterals across the CA3 field, as well as projecting via the Shaffer collaterals on to the CA1 field. In addition to the trisynaptic circuit through the hippocampus, the EC projects directly via the perforant path to both the CA3 and CA1 fields. Finally, only the CA1 region projects back onto the PHR.

Role of Trisynaptic Circuit in Learning vs. Recall

As in many previous hippocampal models, we have incorporated the key assumption that the hippocampus operates in two distinctly different modes during learning and retrieval. Several lines of evidence support the view that the trisynaptic circuit

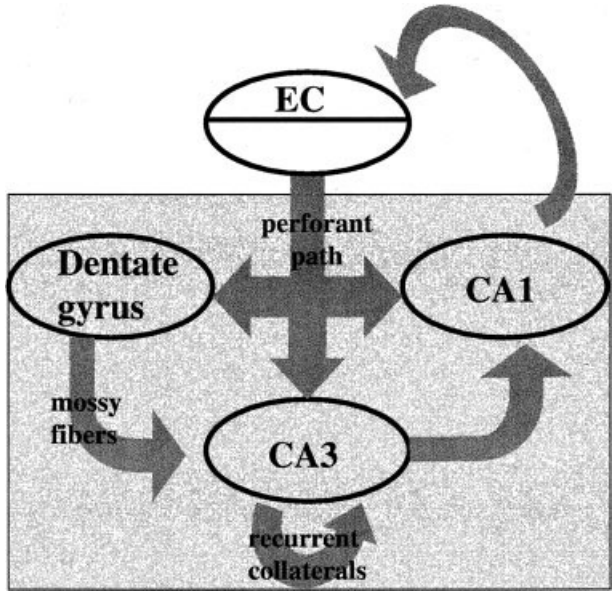

FIGURE 1. The major fields and pathways within the hippocampus included in our model. Additionally, there are inhibitory interneurons that are thought to regulate the overall activity level. Rather than model these explicitly, we normalize the activation levels within each layer by using a *k*-winner-take-all activation function (see text).

plays an active role during encoding, whereas the perforant path connections from the EC to the CA3 and CA1, bypassing the DG, predominate during retrieval. For example, relatively high perforant path stimulating frequencies (5–10 Hz) are required to produce polysynaptic activation of the CA3 and CA1 regions via the DG (Yeckel and Berger, 1990) and are optimal for LTP induction in the DG (Greenstein et al., 1988), whereas low stimulating frequencies (0.2 Hz) are sufficient to excite CA3 and CA1 but not DG cells (Yeckel and Berger, 1990). Correspondingly, the hippocampus spontaneously exhibits oscillatory theta-frequency firing, about 6–8 Hz, when an animal is engaged in exploratory-attentive behaviors (see e.g., Lynch et al., 1991; Buzsaki, 1998). Additionally, selective lesions of DG granule cells by micro-injections of the neuro-toxin colchicine disrupt spatial learning, but do not abolish CA3 and CA1 place specificity (McNaughton et al., 1989), and reversible inactivation of hippocampal mossy fiber synapses disrupts spatial memory formation but not recall in the Morris water maze (Lasalle et al., 2000). Thus, the DG, via its mossy fiber projections to the CA3, may mediate the creation of new, distinctive codes, but may not be necessary for the subsequent reactivation of these codes.

Hasselmo has suggested a mechanism by which the HPC could switch in and out of an acquisition/encoding mode, through the action of the neuromodulator acetylcholine (ACh) (see e.g., Hasselmo, 1999). Microdialysis studies in the hippocampus reveal high levels of acetylchoine during active waking states and in novel environments associated with theta activity, whereas low levels of ACh are observed during quiet behaviors in the awake state and during slow wave sleep (for a review, see

[1]Though other PHR areas also communicate with the hippocampus, most notably the subiculum, they will not be considered further here.

Hasselmo and McGaughy, 2004). As in Hasselmo's model, we further assume that the CA3 recurrent collaterals and CA3-to-CA1 Shaffer collaterals are important for pattern completion and therefore dominate neuronal dynamics during recall, but not during encoding. Thus, during encoding of novel information, these associative pathways should contribute only weakly to neuronal activations, whereas perforant path and mossy fiber inputs dominate in driving CA3 and CA1 activations. Consistent with this view, ACh has been found to suppress neuronal transmission along the CA3 recurrent collaterals (Hasselmo et al., 1995) and Shaffer collaterals (see e.g., Herreras et al., 1988; Hasselmo and Schnell, 1994).

ROLE OF CA3 RECURRENT COLLATERALS IN PATTERN COMPLETION VS. SEQUENCE RECALL

According to Marr's and many other hippocampal theories, a primary function of the hippocampus is pattern completion. That is, once a memory has been stored, it should later be recalled in detail when given an appropriate cue. Recent empirical results with knockout mice having selective loss of NMDA receptors in CA3 are consistent with the pattern completion function of the hippocampal circuit; these mice show normal acquisition and normal place fields in the Morris water maze task, but a loss of spatial selectivity in both areas CA3 and CA1 after visual cue removal (Nakazawa et al., 2002). The CA3 recurrent collaterals in Marr's view were critical to the hippocampal pattern completion capabilities. Following Marr, many models have assumed that the CA3 recurrent collaterals form an attractor network that auto-associates elements of a pattern together (McNaughton and Morris, 1987; Rolls, 1989; Treves and Rolls, 1992; Kali and Dayan, 2000b; O'Reilly and Rudy, 2001). However, we find that the feed-forward pathways between the CA3 and CA1 regions already afford the model with substantial pattern completion capabilities. We therefore favor a more recently proposed role for the CA3 recurrent collaterals, namely, in learning temporal associations and thus forming a continuous attractor network linking spatiotemporal representations (Levy, 1996; Gerstner and Abbott, 1997; Wallenstein and Hasselmo, 1997; August and Levy, 1999; Stringer et al., 2002).

The Learning Principle

An implicit assumption in most, if not all, hippocampal models is that the computational goal of the hippocampus is simply to encode or memorize, and subsequently perform cued recall. In practice, however, most hippocampal modeling efforts have taken a bottom-up approach in meeting this goal, by constructing models out of learning rules selected a priori (typically of the Hebbian form) combined with architectural and processing constraints. The approach taken here is to start with the explicit goal of pattern encoding by proposing a quantitative objective function for the learning. By maximizing this objective function, subject to biological constraints on cell numbers,

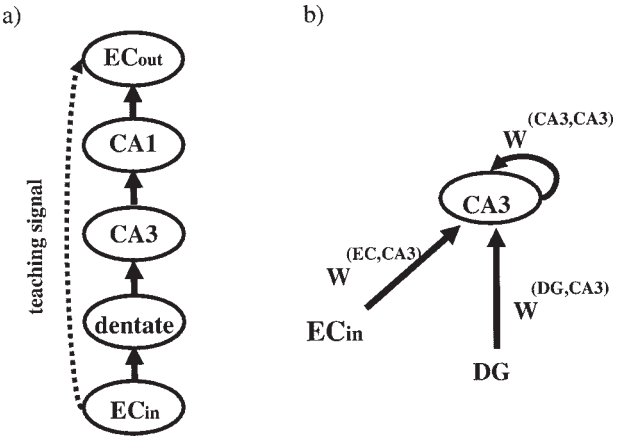

FIGURE 2. (a) An auto-encoder model with entorhinal input and output layers, and hidden layers for each of the hippocampal regions. (b) Part of the circuit for our hippocampal model, showing the notation used here for weights and activations.

activation levels, and connectivity in each region, we can derive learning equations in a principled manner for all pathways within the hippocampal circuit.

One way to model pattern encoding in a multilayer circuit is using the auto-encoder architecture (shown in Fig. 2a). Gluck and Meyers (1993) used the auto-encoder as the basis of their hippocampal model, as a means of generating compressed representations, though it was not intended to be taken literally as a model of hippocampal neural circuitry. If we did want to use the auto-encoder as a model neural circuit, we would have an architecture with three hidden layers as shown in Figure 2a. This model could be trained using the back-propagation algorithm (Rumelhart et al., 1986) to generate output patterns $EC^{(\text{out})}$ identical to the input patterns $EC^{(\text{in})}$, by minimizing the mean squared error,

$$\text{MSE} = \sum_{p} ||EC^{(\text{out})}(p) - EC^{(\text{in})}(p)||^2 \qquad (1)$$

with respect to the weights in the network, where $\|V\|$ is the length (Euclidean norm) of a vector $\boldsymbol{V}$ and p indexes over all training patterns p. Unfortunately, the backward-going connections required to pass derivative terms backwards through the network according to this algorithm are not consistent with hippocampal anatomy. Further, back-propagation learning is very slow, typically requiring thousands of learning iterations through the training patterns—which is not characteristic of animal learning in hippocampally dependent tasks.

To overcome these problems with the auto-encoder model, we propose a simpler, but related model of learning. The key idea is that each layer within the hippocampal circuit should try to encode its perforant path input in a greedy fashion—without regard to what higher layers are doing—by building on the encoding accomplished by previous layers. For example, the CA3 region should find an optimal encoding of the perforant path activations $EC^{(\text{in})}$, given its joint inputs from the perforant path, the DG, and the CA3 recurrent collaterals. Like-

wise, the CA1 region should find an optimal encoding of $EC^{(\text{in})}$ given its joint inputs from the perforant path and the CA3 region.

To make this idea more precise, we propose the following objective for the learning: each hippocampal layer should form a neural representation that could be transformed in a simple manner—i.e., linearly—to reconstruct the original activation pattern in the EC. It is important to note, however, that the model itself is highly nonlinear because of the sparse coding in each region and the multiple stages of processing in the circuit as a whole; the notion of linearity comes in only when we consider the process of *reconstructing* the EC activation pattern from any one region's activities. We need an objective function that measures the quality of each region's representation with respect to how well it reconstructs the EC activations. However, the computation must not require an explicit reconstruction via a set of projection weights, because only the CA1 region sends direct afferent connections to the EC. We can achieve this by assuming that the perforant path connection weights could be used in reverse to reconstruct the EC input pattern. Taking the CA3 layer as an example (see Fig. 2b), we have perforant path input from the $EC^{(\text{in})}$ that is associated with a matrix of weights $W^{(\text{EC,CA3})}$. The CA3 region also receives input connections from the DG with associated weights $W^{(\text{DG,CA3})}$ as well as recurrent collateral input from within the CA3 region with connection weights $W^{(\text{CA3,CA3})}$. We will use the transpose of the perforant path weights, $(W^{(\text{EC,CA3})})^{\text{T}}$, to calculate the CA3 region's reconstruction of the entorhinal input vector:

$$EC^{(\text{reconstructed})} = W^{(\text{EC,CA3})^{\text{T}}} CA3 \tag{2}$$

The goal of the learning is to make this reconstruction as accurate as possible. To quantify this goal, the objective function to be maximized here is the cosine angle between the original and reconstructed activations:

$$\begin{aligned} Perf^{(\text{CA3})} &= \cos(EC^{(\text{in})}, W^{(\text{EC,CA3})^{\text{T}}} CA3) \\ &= \frac{(EC^{(\text{in})})^{\text{T}} \; (W^{(\text{EC,CA3})^{\text{T}}} CA3)}{||EC^{(\text{in})}|| \; ||W^{(\text{EC,CA3})^{\text{T}}} CA3||} \end{aligned} \tag{3}$$

When the original and reconstructed pattern vectors are identical, the cosine angle between the two vectors takes on a maximal value of 1. By rearranging the numerator, and appropriately constraining the activation levels and the weights (see Appendix for details) so that the denominator becomes a constant, we can equivalently maximize the following simpler expression:

$$Perf^{(\text{CA3})} = (W^{(\text{EC,CA3})} EC^{(\text{in})})^{\text{T}} CA3 \tag{4}$$

This says that the incoming weighted input from the perforant path should be as similar as possible to the activation in the CA3 layer. Note that the CA3 activation, in turn, is a function of both perforant path and DG input as well as CA3 recurrent input (see Fig. 2b). We assume that the weights from the DG—representing the mossy fiber synapses—are very sparse, but 100 times larger in magnitude than other weights in the circuit so that the dentate input dominates in the calculation of the CA3 activations during learning. In the present version of the model, we further assume that the mossy fiber pathway is nonplastic. In fact, there is substantial evidence for various forms of LTP in hippocampal mossy fibers, although the functional significance of mossy fiber plasticity for hippocampal-dependent memory has been questioned (Hensbroek et al., 2003).

The objective functions for the dentate and CA1 regions have exactly the same form as Eq. (4), using the DG and CA1 activations and perforant path connection weights, respectively. Thus, we have a computational objective for the learning in each region: to maximize the overlap between the perforant path input and that region's reconstruction of the input. This objective function can be maximized with respect to the connection weights on each set of input connections for a given layer, to derive a set of learning equations.

The Learning Equations

We derived learning equations for updating the weights for each pathway in the circuit shown in Figure 1 by differentiating our objective function with respect to the weights in each pathway. For the weights on the direct projections from the *EC*, $W^{(\text{EC,DG})}$, $W^{(\text{EC,CA3})}$, and $W^{(\text{EC,CA1})}$, under appropriate assumptions (see Appendix), the final learning equations have a "Hebbian" form—a product of presynaptic and postsynaptic activations. For example, the change in weight for the EC-to-CA3 connections is:

$$\Delta W_{jk}^{(\text{EC,CA3})} = lrate \; EC_k \; CA3_j \tag{5}$$

where *lrate* is a learning rate constant (0.5 in the simulations reported here), EC_k is the activation of the kth EC input unit, and $CA3_j$ is the activation of the jth CA3 unit. The weight updates for $W^{(\text{EC,DG})}$, $W^{(\text{EC,CA1})}$ have exactly the same Hebbian form. The assumptions we make, in order for the learning equations to simplify to a Hebbian form, are as follows: During learning, the influence of the DG region dominates over that of the EC input in driving the CA3 region's activations, while the EC region dominates over that of the CA3 in driving the CA1 activations.

For the weights on indirect projections, ($W^{(\text{CA3,CA1})}$ and $W^{(\text{CA3,CA3})}$), the learning rule takes a slightly different form. Rather than being a product of presynaptic and postsynaptic activations, the postsynaptic activation is replaced by the net weighted input from the EC. Thus, for the indirect pathways, the EC activation pattern acts as an implicit teaching signal. Training the CA3-to-CA1 connections in this manner is rather similar to the heteroassociative neural network model for this pathway proposed by Hasselmo et al. (1996). In a heteroassociator, one set of inputs acts as an explicit teaching signal by clamping the output neurons' activations to a target pattern, while the weights on the second set of input connections

undergo Hebbian associative learning. However, under appropriate assumptions (namely, that during learning, the DG provides the dominant input to the CA3 while the EC provides the dominant input to the CA1; see Appendix) we can further simplify the learning equations for the indirect pathways to obtain pure Hebbian rules for the indirect connections, which have exactly the same form as for the perforant path weights. Hasselmo and Schnell (1994) have suggested that ACh selectively suppresses synaptic transmission in the Shaffer collaterals (denoted $W^{(CA3,CA1)}$ in our model) during learning, which would satisfy the assumption that the perforant path provides the dominant input to the CA1 region during learning.

For the CA3 recurrent collaterals, $W^{(CA3,CA3)}$, given that the input to the hippocampus is a time-varying sequence, the CA3 recurrent connections learn to associate together CA3 activation patterns at successive points in time, $t-1$ and t:

$$\Delta W_{jk}^{(CA3,CA3)} = lrate\ CA3_k(t-1)\ CA3_j(t) \qquad (6)$$

The CA3 recurrent collaterals thereby associate together patterns experienced within the same temporal context. In the simulations reported here, when the input is a set of unrelated items or episodes rather than a temporally structured sequence, we assume that activations in the hippocampus are reset to zero at the beginning of each event to be encoded, so that there is no temporal association made across the unrelated episodes by the CA3 recurrent collaterals.

Model Architecture and Activations

The number of model neurons and activity levels in each layer is given in Table 1 for all simulations reported here, except where noted. The DG layer size varied between 25 and 1,000, while the number of active DG units was held fixed at 4. All layers with interconnecting weights (shown in Fig. 1) were fully connected except in the DG → CA3 path in which each DG unit was connected randomly to three CA3 units with fixed weights of 100 to simulate the sparse but very powerful mossy fiber connections.

Sparse coding was achieved using a *k*-winner-take-all activation function (O'Reilly and McClelland, 1994), with *k* chosen to produce progressively less sparse codes for each layer (see Table I). First, each neuron computed its net input, by summing over all incoming connections, the incoming activations multiplied by the weights. Next, the *k* units with the greatest net inputs were activated, and the remaining units' activities were set to zero. Finally, the mean of each layer's activity vector was subtracted from each unit's activation to compute its final output. Using zero-mean activation vectors improves the capacity in sparsely coded networks using Hebb-type learning rules (Willshaw and Dayan, 1990). Note that one could achieve sparse coding in a more biologically realistic neural circuit by approximating the *k*-winner-take-all activation, using a set of inhibitory interneurons to regulate the level of excitation, or by finding an appropriate threshold, for each region, that resulted on average in *k* neurons being active at one time.

TABLE 1.

Number of Neurons in EC and Hippocampal Regions in the Rat[a] and Corresponding Layer Sizes and Activity Levels[b] in the Hippocampal Model Simulations Described Here

	Rat data	Hippocampal model	
Field	No. of cells	No. of units	% Active
CA1	420,000	400	3.0
CA3	330,000	300	3.0
DG	1,000,000	1,000	0.4
EC	200,000	200	10.0

[a]From Amaral et al., 1990.
[b]On the basis of data on the rat hippocampus (Barnes et al., 1990; Jung and McNaughton, 1993).

Neurogenesis

It is difficult to obtain consistent estimates from the experimental literature as to the magnitude of neurogenesis because of major methodological differences between studies. One such factor is the dose of labeling agent used to mark dividing cells. Recent data (Cameron and McKay, 2001) using a high dose of BrdU (300 mg/kg), a nontoxic marker of dividing (S-phase) cells in the adult rat DG, combined with a second S-phase marker (3H-thymidine) suggest that about 9,000 new cells are generated per day, with a survival rate of about 50% within 5–12 days. Given that there are about a million granule cells in the rat DG, this would amount to a daily turnover of about 0.45% of the granule cell population. Unpublished results from Martin Wojtowicz at the University of Toronto (personal communication) in 35-day-old rats (i.e., young adults) indicate that of about 10,000 cells born every 24 h, 70% survive two weeks, suggesting a daily turnover rate of nearly 1%, whereas in 2–3-month-old rats, only 20–30% of newborn cells survive. Note, however, that the survival rates of newly born cells observed in caged rats are likely an underestimate of what would be observed in more natural settings, because many environmental factors including associative learning (Gould et al., 1999a,b), locomotion (van Praag et al., 1999), and environmental enrichment (Kempermann et al., 1997) have been shown to increase the survival rates.

In the simulations reported here, we have modeled the effects of neurogenesis and/or cell death in the DG in two different ways. The first is to vary the number of dentate neurons, and the second is to vary the rate of neuronal turnover while keeping the number of dentate neurons fixed. In the general discussion, we consider the conditions under which these two processes might occur. In our simulations of neuronal turnover in the DG, we have varied the neuronal turnover rate from 25 to 100% per day, which is clearly one or more orders of magnitude too high. However, a more biologically realistic turnover rate of 0.5–1% per day would not produce a detectable effect given the scale of our model, with only 1,000 dentate cells in total. Additionally, it is possible that this slow rate of turnover

would affect the stability of memories over a much longer time scale of weeks or even months. This will be explored in future work.

In the next section, we report the results of three simulations. Simulation 1 evaluates the model's recognition and cued recall for sets of unrelated items. The model's performance is compared with that of a hippocampally lesioned model. Both models' ROC curves for recognition memory are also compared with human data (Yonelinas et al., 2002). Simulation 2 evaluates the effect of shrinking the size of the dentate region on recall of unrelated items, paired associates, and related items. Simulation 3 evaluates the effect of neuronal turnover on recall of unrelated items, paired associates, and related items.

SIMULATION 1, MEMORY CAPACITY: RECOGNITION AND RECALL

A fundamental question about hippocampal coding concerns the unique circuitry of this brain structure: Why are there so many regions? What do they each contribute to the fast, efficient learning for which this structure is so well-known? Simulation 1 begins to address these questions by comparing the memory capacity of the complete hippocampal model with that of a one-layer pattern associator, lacking the DG and CA3 and CA1 regions—thereby simulating a total hippocampal lesion. To study memory capacity, we varied the number of training patterns to be remembered. We expected to observe a degradation in performance with increasing memory load. Two different performance measures were used in this simulation: recognition and cued recall. It was hypothesized that the internal layers of the hippocampus, employing sparse intermediary codes for the EC activation pattern, would increase the memory capacity of the model.

Methods

Weights in all layers of the model, except the mossy fiber pathway, were initialized to random values between 0 and 0.5 and were constrained not to exceed ± 0.5 throughout learning.

The training patterns consisted of sets of unrelated items. Each training pattern was created by generating a random 200-element binary pattern, with each element set to 1 with probability 0.1 and to 0 with probability 0.9, resulting in 10% of units being active. The pattern was then translated to have zero mean. For evaluating recognition memory, we created a test set of new patterns drawn from the same distribution as the training set, but that had not been used during training.

The learning phase was simulated as follows. Each training item was presented to the EC input layer, activation was propagated in a single forward pass through the hippocampal circuit, and the weights were then updated as described in the previous section. Three passes through the training patterns were made, so that the model was trained on each of the patterns three times. During the learning phase, activations in the hippocampal circuit were computed as follows. First, the activations were calculated for the dentate region based on the EC input. Next, activations were calculated for the CA3 region based on the combined perforant path, CA3 recurrent input, and DG inputs. However, because of the large magnitude of the DG-to-CA3 weights (the mossy fibers), the DG inputs dominated in the calculation of the CA3 activations. Also, because we were simulating single item memory rather than memory for sequences, between each pair of successive items, the states of the CA3 layer units were initialized to be all zeros. Therefore, in this simulation, the CA3 recurrent connections underwent no learning. Finally, the CA1 layer activations were calculated based on the combined input from the EC input layer and the CA3 region.

To simulate a hippocampally lesioned model, we created a version of the model shown in Figure 1 containing only the EC input and output layers, but lacking all the hippocampal regions. The two layers were fully interconnected, and the model was trained on each pattern by setting both the input and output activations to the training pattern, and then updating each of the weights according to a Hebbian learning rule:

$$\Delta W_{ij} = lrate\ EC_i^{(\text{in})}\ EC_j^{(\text{in})} \tag{7}$$

After this weights were clipped to lie within $(-2, 2)$. The choice of minimum and maximum weights was determined empirically to give optimal memory performance. The lesioned model was trained and tested on the same patterns as for the intact model, using the same k-winner-take-all activation function at the EC output layer. Thus, this model can be described as a feed-forward nonlinear Hebbian pattern associator with one layer of adaptive weights.

To evaluate recognition memory, a test item was presented to the EC input layer, activations were calculated throughout the hippocampal circuit, and finally, the EC output layer's activation was compared with that of the input layer. Activation was propagated in a single pass through the hippocampal circuit from the EC input layer to the EC output layer as follows. First, the activations were calculated for the CA3 region via the perforant path connections (but not via the DG pathway). The CA1 layer activations were then calculated based on the combined input from the EC input layer and the CA3 region. Finally, the EC output layer activations were calculated based on their input from the CA1 region. As was done during the learning phase, between each pair of items, the states of the CA3 layer units were initialized to be all zeros. The CA3 recurrent connections, therefore, had no effect on the processing of successive items in this simulation. An item was considered to be recognized correctly if the activation pattern generated at the EC output layer had at least 95% of the pattern elements correct.

To test the recall of each item, the network was cued with a degraded version of each item with 50% of the active inputs turned off, activations were computed in the same manner as earlier, and the EC output layer was tested on its ability to complete the item. An item was considered to be recalled correctly if the completed pattern had at least 95% of its elements correct.

Recognition and recall were evaluated for the two models on varying sized pattern sets, ranging from 20 to 500 items in

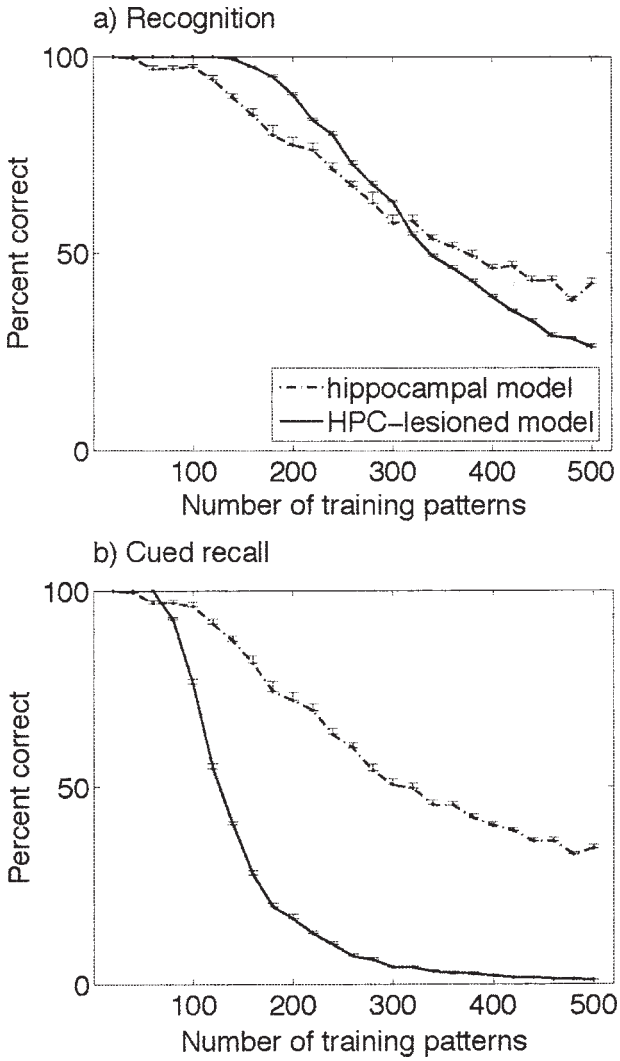

FIGURE 3. Mean recall and recognition scores, averaged across 10 networks (run from different random initial weights and with different randomly generated training patterns), for the hippocampally lesioned and intact model.

increments of 20 patterns. For each training set size, 10 different runs were simulated. On each run, the model started with a different set of initial random weights and a different set of training patterns.

ROC curves for recognition memory for the two models were evaluated using four different training set sizes: 500, 1,000, 1,500, and 2,000 patterns, for one run each. ROC curves were generated by applying a sliding criterion to the reconstruction measure (described previously) for the EC layer, and at each criterion level, computing the probability of hits (correctly recognizing an old or previously learned item as old) and false alarms (incorrectly recognizing a new item as old).

Results and Discussion

The average recognition memory performance of the full hippocampal model and of the hippocampally lesioned model, for varying numbers of training patterns, is shown in Figure 3a. Both models perform at ceiling up to about 100–200 patterns, after which point the performance declines with increasing pattern set sizes. The lesioned model's recognition memory capacity is superior for the smaller training set sizes: performance stays at ceiling up to about 200 patterns, whereas the intact model's capacity only stays at ceiling up to about 100 patterns. However, the recognition performance of the intact model drops off much more gradually, so that it overtakes that of the lesioned model for pattern sets larger than 300 patterns.

The recognition memory results shown in Figure 3a apply to one arbitrary criterion for correct recognition. We could instead have used a lower criterion and thereby artificially boosted performance. A better measure of recognition memory would compare hits and false alarms on old and new items, respectively, across a range of criteria. The ROC curve does just this. The superiority of the intact model's recognition memory for very large pattern sets can be seen even more clearly in the ROC curves, shown in Figures 4b and 4c. On training sets of 2000 patterns, the intact model's ability to discriminate old from new patterns is still quite good, whereas the ROC curve for the lesioned model is nearly flat. It is interesting to compare the ROC curves of the intact and hippocampally lesioned model with humans with intact and lesioned hippocampi. Although it is often claimed that hippocampal amnesics have intact recognition memory, a careful analysis using ROC curves by Yonelinas et al. (2002) has shown that this is not the case; these data are reproduced in Figure 4a. The ROC curves generated by our two models at the largest training set sizes are most similar to the human data from Yonelinas et al. (2002) (shown in Fig. 4a); the performance of our lesioned model is similar to that of patients with large medial temporal lobe lesions (denoted H+ in Fig. 4a), while our full model performs similar to controls.

In contrast to recognition memory, recall performance (Fig. 3b) differs markedly for the two models: the lesioned model's performance drops very sharply, as the training set size increases beyond 100 patterns, whereas that of the full hippocampal model declines very gradually beyond this point.

Note that the intact hippocampal model lacks direct connections from the EC input to the EC output layer. This may explain the advantage the lesioned model has for smaller pattern set sizes. We deliberately omitted direct within-EC connections to force the hippocampal model to make use of its internal layers and sparse coding to attain higher capacity because we were interested here in investigating the unique contributions made by the hippocampus proper. However, the hippocampus does have connections from the deep layers of the EC to the superficial layers, and so a more realistic model would incorporate processed output from the hippocampal circuit as part of its input.

In conclusion, the memory performance of our hippocampal model fits the general patterns in the human data quite well. The recognition memory capacity is extremely high, and is reduced by hippocampal lesions to a degree comparable with that seen in human data, while the recall performance is

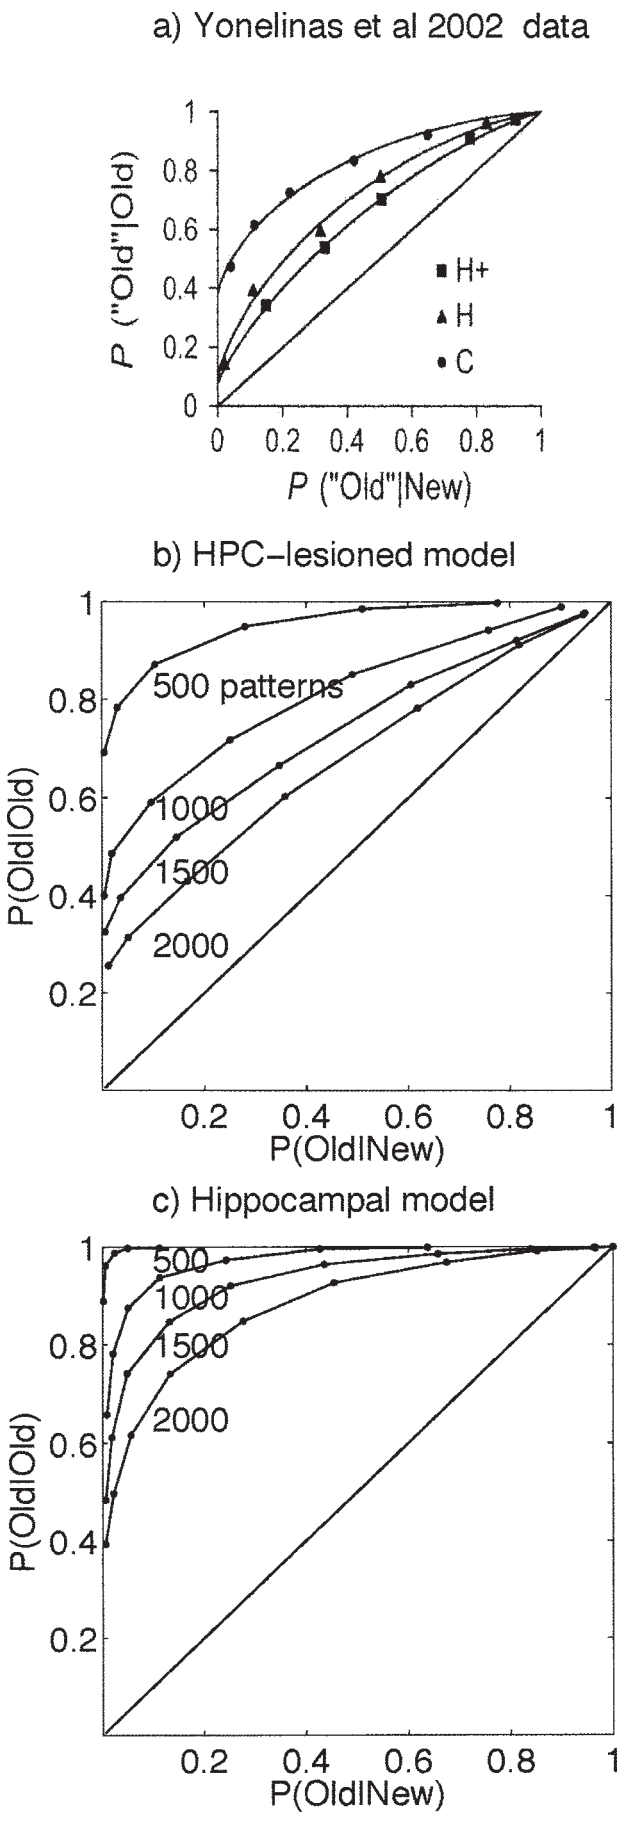

severely impaired after hippocampal damage. In interpreting the recall and recognition performance of the simulated models, it must be emphasized that the absolute numbers should not be taken as literal predictions—it is the qualitative rather than the quantitative pattern in these results that is of interest, given the small scale of the network relative to the real hippocampal circuit.

SIMULATION 2: EFFECT OF DG SIZE ON MEMORY CAPACITY

In Simulation 2, we investigated the contribution of the DG to recall in the intact model. Models with 125, 250, 500, and 1,000 dentate layer neurons were compared on three different pattern sets: (a) 100 unrelated items, (b) 10 paired associates, and (c) 12 related items. For the paired associates, the role of the CA3 recurrent collaterals was critical in associating each pair of items across time. The cued recall test, in this case, consisted of cuing the model with only the first item in each pair. It was hypothesized that the size of the DG would have a direct impact on the memory capacity of the model, for all three pattern sets.

Methods

The model's weights were initialized in the same manner as in Simulation 1. In the full scale model with 1,000 dentate units, as indicated in Table I, the activity level was 0.4%, and thus, 4 of the 1,000 units were active. In all models with fewer dentate units, the total number of active units in the dentate region was held constant at 4 to permit a fair comparison across the different sized models.

The training patterns in the three conditions consisted of (a) 100 unrelated items created in the same manner as in Simulation 1, (b) 10 paired associates, where the items in each pair were unrelated and created in the same manner as in simulation 1, and (c) 12 related items. A set of related items was generated as follows. First, three category prototypes were randomly created in the same manner as for the unrelated items. Next, from each prototype, four training exemplars were generated by randomly swapping the state of each input element with that of another with probability 0.1. We thereby created a set of 12 patterns consisting of 3 categories with 4 members from each category. Patterns from the same category were highly similar, and therefore highly confusable. We deliberately chose patterns with a very high degree of overlap for Simulations 2 and 3 that the model would find extremely challenging, so as to illustrate the full impact of the size of the dentate layer

FIGURE 4. Recognition memory ROC curves, (a) for patients with mild hippocampal lesions (H) due to hypoxic-ischemic events, more extensive temporal lobe lesions (H+), and controls (C), (b) and (c) for typical runs of our hippocampal lesioned and intact models, respectively, when tested on 500 to 2000 items. Part (a) was reproduced with permission from Yonelinas et al. (Nature Neuroscience, 2002, Volume 5, Number 11, Pages 1236–1241, Fig. 4b).

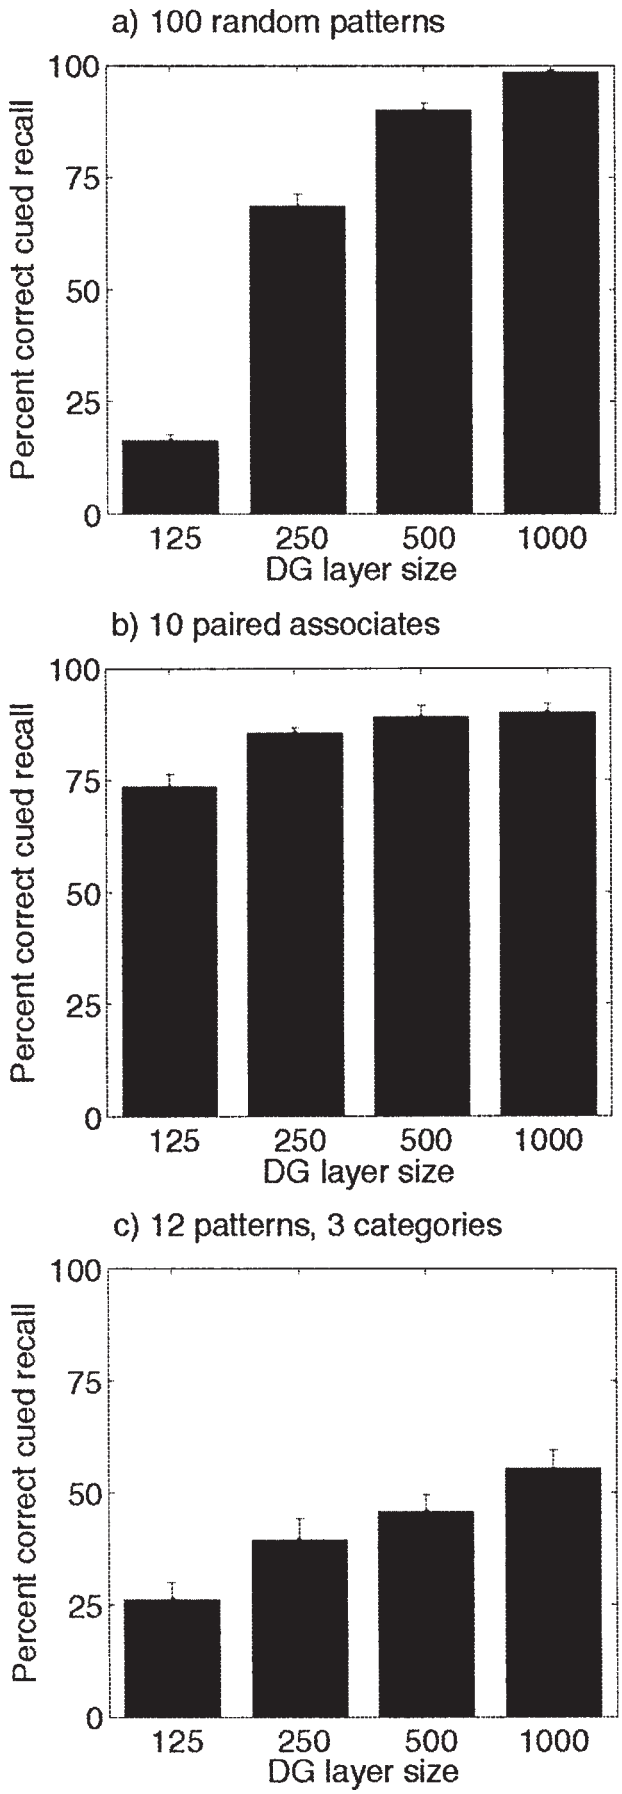

(Simulation 2) and of neuronal turnover (Simulation 3). With less overlap, or fewer patterns of the same category, the model's performance was much better, but the effects of neurogenesis were consequently much smaller.

The model was trained in the same manner as in Simulation 1, except that for the paired associate patterns, the CA3 units' states were not set to zero between items within the same pair. Both items in each pair, therefore, contributed to learning on the CA3 recurrent collaterals. For each training set size and model size, 20 runs were simulated starting from different initial random weights and with different pattern sets.

To test recall of each item, the network was cued with a degraded version of the item with 50% of the active inputs turned off, and tested on its ability to complete the item. For testing recall of paired associates, the network was cued with the first item in the pair, and tested on its ability to complete the second item in the pair, using the CA3 recurrent collaterals to complete the paired associate across time. In all cases, an item was considered to be recalled correctly if the completed pattern had at least 95% of its elements correct.

Results and Discussion

We found that smaller dentate sizes produced a substantial drop in the recall performance of the hippocampal model, as shown in Figure 5. As the dentate size approached 1,000—the size of our full model—performance approached ceiling levels for the unrelated items and was only slightly worse for paired associates. For highly confusable items, performance was much worse, but showed the same pattern of dramatic improvement with increasing size of the dentate layer.

Note that with the smallest sized dentate region, performance is even worse than our simulated hippocampally lesioned model (Simulation 1). This is because the full model lacks the direct connections from the EC input to EC output layers, and is therefore forced to encode patterns strictly using its "shrunken" dentate region, which likely causes some confusions between patterns.

SIMULATION 3, NEUROGENESIS: EFFECTS OF NEURONAL TURNOVER

In Simulation 3, we investigated the effect of neuronal turnover in the DG on cued recall of unrelated items, paired associates, and related items. We have proposed a functional role for neurogenesis in the DG, namely, to create distinct memory traces for highly similar items. We, therefore, hypothesized that the related, highly confusable items would benefit most from neuronal turnover.

Methods

The model's weights were initialized in the same manner as in Simulations 1 and 2. The three sets of training patterns—

FIGURE 5. Effect of number of dentate cells on recall, averaged across 20 runs (with standard error bars) for (a) unrelated items, (b) unrelated paired associates, and (c) related items.

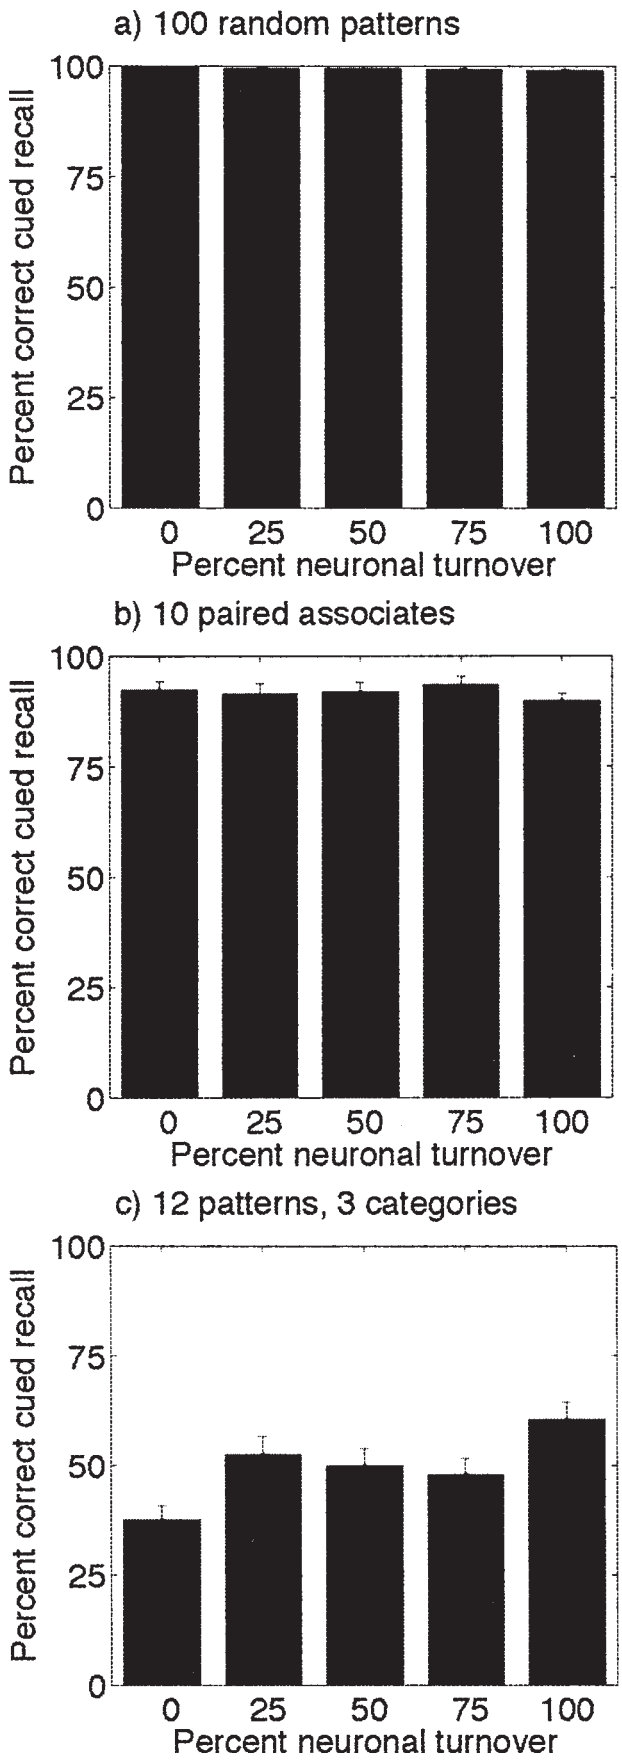

unrelated items, paired associates and related items—were created in the same manner as in Simulation 2.

The model was trained in the same manner as in the previous simulations, except that learning proceeded for 10 repetitions through the training set, to simulate learning over 10 successive days. On each day, the model was trained once on each pattern in the training set. Between days, we simulated the passage of time by training on a further set of 20 novel, unrelated items, after which the dentate layer underwent neurogenesis. Neuronal turnover was simulated by randomly selecting a fixed percentage of the dentate layer neurons, and rerandomizing their incoming weights from the EC, and reconnecting them randomly to a different subset of CA3 cells. The percentage of "new neurons" created in this manner was either 0, 25, 50, 75, or 100. After the 10th simulated day, testing proceeded as in Simulation 2.

Results and Discussion

The effect of neuronal turnover on performance in the three types of patterns was rather different, as shown in Figure 6. In contrast to the lack of an effect of neuronal turnover on unrelated items and unrelated pairs, there was a substantial effect on related items. Superimposed on the positive effect of neuronal turnover was a negative effect because of the interference of the novel, unrelated items learned on each day. Thus, it appears that repeated learning sessions interleaved with the opportunity for neurogenesis can be of benefit, most particularly, in the encoding of highly confusable items.

GENERAL DISCUSSION

The main result of this paper is a theoretical one: we can view the job of the hippocampus as encoding the incoming activations from the cortex, and recreating those same activation patterns during recall. We formalize the notion of encoding by proposing that different parts of the hippocampus are all greedily optimizing the same performance measure: the similarity between the original and reconstructed inputs. Thus, each hippocampal region should learn an invertible code. If successive layers of the model have relatively stable weights, then each successive layer can build on the learning that has taken place in previous layers. However, even when learning is fairly rapid, only after 2–3 repetitions, recall is still very accurate in this model and far superior to a one-layer pattern associator. Previous models of individual regions of the hippocampus have been useful in demonstrating the utility of important computational principles thought to be employed by the hippocampus, such as sparse coding (O'Reilly and McClelland, 1994), pattern completion (McNaughton and Morris, 1987; Rolls, 1989; Treves and Rolls, 1992; Kali and Dayan, 2000b),

FIGURE 6. Effect of neuronal turnover on recall, averaged across 20 runs (with standard error bars), for (a) unrelated items, (b) unrelated paired associates, and (c) related items.

and temporal associations (Levy, 1996; Gerstner and Abbott, 1997; Wallenstein and Hasselmo, 1997; August and Levy, 1999; Stringer et al., 2002). Nonetheless, it is also important to demonstrate in a simulation of the entire multilayer circuit that these principles still hold up. At the very minimum, the model should outperform a simple one-layer pattern association network, which is confirmed by our simulations. Our model also accounts for deficits in recognition memory and recall after hippocampal lesions. This topic is explored in much greater depth in simulations by Norman and O'Reilly (2003). Finally, we have presented novel results on the consequences of neurogenesis for hippocampal coding.

The learning principle of invertibility, combined with sparse coding and other biological constraints, leads to a set of simple Hebb-like learning rules. This work therefore has broader implications for models that employ Hebbian learning rules. For example, Rumelhart and Zipser's competitive learning model (Rumelhart and Zipser, 1986) is based on a variant of Hebbian learning combined with a sparse coding constraint in the form of a winner-take-all activation function. The learning in this model minimizes the error between each unit's weight vector and the input patterns that unit is selective for. Similar to the multilayer hippocampal model proposed here, competitive learning can also be viewed as constructing invertible codes subject to a sparseness constraint.

Related Work

Since Marr published his computational theory of hippocampal coding, there have been many models in the literature embodying the coding principles he advocated. Interestingly, a multilayer model of associative memory was proposed by Baum et al. (1988) that encapsulates many of Marr's proposed features: pattern encoding, internal layers with sparse coding, and recurrent connections for pattern completion. Although this model bears a striking similarity to many of the models discussed here, it was worked out completely independently of considerations about Marr's theory or the hippocampus, and emerged purely from theoretical coding considerations (John Moody, personal communication).

Several other modelers have simulated the complete hippocampal circuit (e.g., O'Reilly and McClelland, 1994; Hasselmo et al., 1996; Hasselmo and Wyble, 1997; Fellenz and Taylor, 2003; Norman and O'Reilly, 2003). The models developed by Hasselmo and colleagues emphasize sparse coding in the DG, auto-associative recall in CA3, and a role for cholingeric modulation in implementing a heteroassociative memory function in the CA1 region. The models by O'Reilly and colleagues (e.g., O'Reilly and McClelland, 1994; Norman and O'Reilly, 2003) emphasize the differences between cortical and hippocampal modes of learning, namely, incremental learning of distributed codes vs. rapid, low-interference learning with sparse codes, as first described by McClelland et al. (McClelland et al., 1995). The learning rules employed in O'Reilly's more recent hippocampal models (e.g., O'Reilly and Rudy, 2001) are a combination of Hebbian and error-driven learning terms. A recently proposed model by Fellenz and Taylor (2003) takes a previously described model of associative memory with hidden layers and adapts it to fit hippocampal circuitry. Interestingly, they propose a randomly generated code in the internal layers that is similar to the effect of neuronal turnover in our model.

Memory for Spatiotemporal Events

There is ample evidence for the crucial role played by the hippocampus in memory for sequences. For example, hippocampal lesions impair rats' memory for the sequential ordering of a sequence of odors (Fortin et al., 2002). A number of models of hippocampal memory for sequences have been proposed (e.g., Levy, 1996; August and Levy, 1999; Lisman, 1999). The model proposed here, as with most other hippocampal models, discretizes time to simplify the problem; the hippocampus can then simply associate items at successive discrete time points. The cost function proposed here could in principle be applied to models with continuous time dynamics. However, the discrete time approximation used to derive the learning equations in our model may not be unreasonable given that hippocampal neurons fire in synchrony with the theta rhythm during behavioral states associated with novel encoding. Further, it has been found that LTP in CA1 pyramidals occurs when stimulation coincides with the peak of the theta cycle (Hyman et al., 2003), evidence that plasticity is phase-locked with the theta rhythm.

An important unsolved question regarding the encoding of spatio-temporal episodes is how exactly does the hippocampus bind together a set of discontiguous events comprising a unitary episode? Some modelers (Wiebe et al., 1997; Lisman, 1999) have taken into account the detailed circuitry of the DG and CA3, and predict that CA3 is able to sustain reverberatory activation patterns for quite long periods of time, on the order of seconds, which could help to solve this problem. Further, Salin et al. (1996) have found that mossy fiber frequency potentiation decays over a period of about 20–40 s, long enough to allow CA3 cells to integrate spikes from the DG over very long time periods, even though dentate cells spike at extremely low firing rates.

Hippocampal-Cortical Interactions

The model described here treats the hippocampus in isolation, whereas in reality, it is a key component in a great many cortical circuits. The hippocampus is thought to be a convergence zone for information from virtually everywhere else in the brain, and it projects back to most, if not all, of the areas it receives input from. Some models have addressed how the hippocampus interacts with the neocortex, for example, in memory consolidation (McClelland et al., 1995; McClelland and Goddard, 1996; Murre, 1999; Kali and Dayan, 2000a). Other theories have focused on the role of the hippocampus as a binding site for storage and reactivation of widely distributed memories (e.g., Bibbig et al., 1995; Moll and Miikkulainen, 1997). Recently, some modelers have begun to address the

complexities of what controls memory storage and retrieval (e.g., Becker and Lim, 2003).

USE IT OR LOSE IT?

Our results of varying the size of the dentate layer suggest the size of the rat hippocampus in near optimal, in the sense that a larger sized DG does not provide much of a gain in memory capacity, if we can extrapolate from our admittedly very small scale model. On the other hand, when the size of the dentate was decreased by 50% or more, retrieval suffered substantially. Thus, the ability to enlarge one's hippocampus at times of greater memory demand might not be very practical, as it would appear to require an increase of an extremely large magnitude to have much of an impact on learning.

It is interesting to note that in avian species, hippocampal volume correlates with memory demands, as measured by degree of food-storing behavior (Healy and Krebs, 1996). And recent human data show that extensive training over a period of years on a very hippocampally demanding task such as driving a taxi in London might perhaps be sufficient to cause changes in hippocampal volume: Maguire and colleagues have reported London taxi drivers to have larger hippocampal volume than controls (Maguire et al., 2000), whereas hippocampal volume was found not to correlate with navigational expertise in non-taxi drivers (Maguire et al., 2003). Of course, volumetric changes could be due to factors other than or in addition to neurogenesis, including changes in dendritic spine densities.

It is more likely that neuronal replacement, rather than neuronal addition, occurs on a very large scale measurable over a time course of days. Our simulations show that neuronal turnover has a significant effect on memory for related items. This is consistent with our prediction that neuronal turnover, which increases the diversity of dentate layer codes across successive learning trials, should be of benefit in spaced learning trials when items are very similar and there is maximal potential for interference. There is substantial evidence in the literature to support the notion that neurogenesis is important for learning and memory. For example, when the number of newly born dentate cells is reduced pharmacologically, a form of hippocampal-dependent learning, trace conditioning, is impaired, while simple conditioning without a temporal gap is not (Shors et al., 2001). Further, neurogenesis is enhanced by learning (Shors et al., 2001), environmental enrichment (Kempermann et al., 1997), and locomotion (van Praag et al., 1999).

On the other hand, stress-related psychiatric illnesses may be associated with a decrease in survival of newly born hippocampal cells. Neurogenesis is suppressed by adrenal hormones (Gould et al., 1992), and stress (Gould et al., 1998; Lemaire et al., 2000; McEwen and Magarinos, 2001) as well as in depression (Jacobs et al., 2000), which is associated with a loss of hippocampal volume (MacQueen et al., 2003). The consequences for memory of cell death in stress-related disorders certainly merit further study.

Acknowledgments

The author thanks Rich Zemel, Ron Racine, Neil Burgess, Allen Cheung, Laurenz Wiskott, and the anonymous reviewers for helpful comments. Ted Meeds and Adeline Chin contributed to the development of the Matlab software used for the simulations reported here.

APPENDIX: DERIVATION OF THE LEARNING EQUATIONS

In each hippocampal layer, weights were adapted to maximize the objective function, the cosine angle between the actual and reconstructed input. Taking the CA3 region as an example (see Fig. 2b), the objective function for this region is

$$\frac{EC^{(\mathrm{in})^{\mathrm{T}}}(W^{(\mathrm{EC,CA3})^{\mathrm{T}}}CA3)}{||EC^{(\mathrm{in})}||\ ||W^{(\mathrm{EC,CA3})^{\mathrm{T}}}CA3||} = \frac{(W^{(\mathrm{EC,CA3})}EC^{(\mathrm{in})})^{\mathrm{T}}CA3}{||EC^{(\mathrm{in})}||\ ||W^{(\mathrm{EC,CA3})^{\mathrm{T}}}CA3||} \tag{A1}$$

If we constrain the activity vector *CA3* and the rows of the weight matrix $(W^{(\mathrm{EC,CA3})})^{\mathrm{T}}$ to have constant lengths, then the denominator in the above equation becomes a constant; it is then equivalent to maximize the numerator

$$(W^{(\mathrm{EC,CA3})}EC^{(\mathrm{in})})^{\mathrm{T}}\ CA3 \tag{A2}$$

The activity constraint was enforced by using the *k*-winner-take-all activation rule described previously. The weight constraint, in its exact form, involves normalizing the rows of $(W^{(\mathrm{EC,CA3})})^{\mathrm{T}}$, which requires each EC unit to have access to its outgoing weights to the hippocampus. This would not be very biologically realistic; instead, we approximate the length constraint by imposing an upper bound on the length of the weight vectors during learning. We do so by forcing each of the weights in the matrix, $W_{jk}^{(\mathrm{EC,CA3})}$, to stay within a fixed range $(-k, k)$ (here we used $k = 0.5$). Thus, after each learning iteration, each unit's incoming weights were clipped to lie within the range $(-0.5, 0.5)$. The weights in all weight matrices (except the mossy fibers) were normalized in this manner.

Our learning principle states that we should maximize the objective function given by Eq. (A2), when applied to each hippocampal region. We can do so by the method of steepest descent (or rather, steepest *ascent* in this case because we are maximizing rather than minimizing an objective function), as is done to derive the learning rules for the back-propagation learning procedure (Rumelhart et al., 1986). Optimization by steepest ascent involves repeatedly calculating the gradient vector—the derivative of the objective function with respect to each of the weights—and updating the weights by a small amount in the direction of this gradient vector.

In preliminary simulations, we used a differentiable, continuous activation function that approximates the *k*-winner-take-all function and permits an exact calculation of the gradient. However, we found that the *k*-winner-take-all activation func-

tion was much faster, led to much simpler learning equations, and worked equally well. To obtain learning equations with the *k*-winner-take-all activation function requires making some approximations in calculating this gradient.

To derive the weight update rules for each pathway in the circuit, again taking the CA3 region weights as an example, we differentiate Eq. (A2) with respect to each incoming weight, w_{jk} (which could be either a perforant path or inter- or intra-layer weight), to each CA3 unit, where $CA3_j$ is the activation of the *j*th unit:

$$\Delta w_{jk} = lrate \frac{\partial((W^{(EC,CA3)} EC^{(in)})^T CA3)}{\partial w_{jk}} \quad (A3)$$

Applying the chain rule to the above expression, we obtain a sum of two terms:

$$\begin{aligned}\Delta w_{jk} &= lrate \left(CA3_j \frac{\partial(W_{jk}^{(EC,CA3)} EC_k^{(in)})}{\partial w_{jk}} \right. \\ &\quad \left. + (W_j^{(EC,CA3)} EC^{(in)}) \frac{\partial CA3_j}{\partial w_{jk}} \right) \\ &= lrate \left(CA3_j \frac{\partial(W_{jk}^{(EC,CA3)} EC_k^{(in)})}{\partial w_{jk}} \right. \\ &\quad \left. + netECinput_j^{(CA3)} \sigma'(net_j^{(CA3)}) \frac{\partial net_j^{(CA3)}}{\partial w_{jk}} \right) \quad (A4)\end{aligned}$$

where *lrate* is a learning rate constant, $net^{(CA3)}$ is the net input from all incoming connections to this unit, $netECinput_j^{(CA3)} = W_j^{(EC,CA3)} EC^{(in)}$ is the net input from the EC to this unit, and $\sigma'(net^{(CA3)})_j$ is the derivative of the nonlinear activation function. The final term, $(\partial net_j^{(CA3)})/(\partial w_{jk})$ is equal to the incoming activation on the *k*th connection associated with weight w_{jk}.

Perforant Path Weights

For the perforant path weights, Eq. (A4) becomes

$$\begin{aligned}\Delta W_{jk}^{(EC,CA3)} &= lrate(CA3_j\, EC_k^{(in)} \\ &\quad + netECinput_j^{(CA3)} \sigma'(net_j^{(CA3)}) EC_k^{(in)} \quad (A5)\end{aligned}$$

Under appropriate assumptions, we can greatly simplify Eq. (A5) by dropping the second term to obtain a pure Hebbian learning rule:

$$\Delta W_{jk}^{(EC,CA3)} = lrate\ CA3_j\ EC_k^{(in)} \quad (A6)$$

and we obtain rules of the same form for $W^{(EC,DG)}$ and $W^{(EC,CA1)}$. This simplified learning rule was found to produce comparable results with Eq. (A5) in simulations. The simplifying assumptions are as follows. For the CA3 region, during training, the mossy fiber pathway dominates in generating the activations, whereas the other pathways, whose weights are bounded within a small range, have a negligible effect. Therefore, the derivative of the CA3 activations with respect to the perforant path weights is negligible relative to the first term in Eq. (A5) and can be set to zero. For the DG and CA1 regions, we assume during training that the EC input dominates over other sources of input in generating activations in that region; if this assumption holds, $netECinput_j^{CA1}$ will be highly correlated with $CA1_j$, and similarly, $netECinput_j^{DG}$ will be highly correlated with DG_j. The two terms in Eq. (A5) will be highly correlated, and therefore, highly redundant, and so the second term can be presumed approximately equal to the first. The assumption that the EC dominates over other sources of input is clearly valid for the DG where there are no other sources of input. For the CA1 region, this assumption is valid provided that the activity level in the CA3 region is sufficiently low relative to that of the EC.

The weights from CA1 back to the EC output layer were simply set to equal the transpose of the perforant path weights to the CA1.

Nonperforant Path Weights

For pathways other than the direct EC input connections, which we simply refer to here as the "indirect weights," the first term in Eq. (A4) is zero. This applies to the CA3 recurrent collaterals and the CA3-to-CA1 pathway (but not the DG-to-CA3 pathway because the mossy fibers do not undergo any learning in our model). Equation (A4) is now

$$\Delta w_{jk} = lrate\ netECinput_j^{(CA3)}\ \sigma'(net_j^{(CA3)})\ \frac{\partial net_j^{(CA3)}}{\partial w_{jk}} \quad (A7)$$

For the CA3 recurrent collaterals, $W^{(CA3,CA3)}$, if the input is assumed to be steady-state, then $(\partial CA3_j)/(\partial W_{jk}^{(CA3,CA3)}) = 0$, and so no learning takes place. If the input is assumed to be a temporal sequence, as in the simulations reported here, then the output of the CA3 layer is a function of the input from the EC and DG at the current time step and the CA3 output from the previous time step $CA3(t-1)$; in this case, the final term in the learning equation is $(\partial net_j^{(CA3)}(t))/(\partial W_{jk}^{(CA3)}) = CA3(t-1)$. For the CA1 region, the final term in Eq. (A7) is $(\partial net_j^{(CA1)})/(\partial w_{jk}) = CA3_k$.

We used two different methods of approximating the derivative of the nonlinear activation function in Eq. (A7). The first method linearizes the activation function, and so $\sigma'(net_j)$ becomes a constant and therefore drops out of the learning equation. In simulations using this linear approximation, we found that the hippocampal model performed reasonably well on all conditions tested except paired associate learning, which relies crucially upon the CA3 recurrent collateral connections to form associations across time. The second method assumes that the CA3 and CA1 activation patterns are relatively stable and optimal at reconstructing the EC pattern, given the EC and DG inputs. We can then equivalently make the goal of the learning on the indirect

pathways be to reconstruct the CA3 (or CA1) activation pattern as well as possible. We can now obtain learning equations for the indirect weights by analogy to our derivation for the direct perforant path weights, yielding simple Hebbian learning rules. For example, for the CA3 recurrent collaterals we have

$$\Delta W_{jk}^{(\mathrm{CA3,CA3})} = lrate\ CA3_j(t)\ CA3_k(t-1) \qquad \text{(A8)}$$

and for the Shaffer collaterals we have

$$\Delta W_{jk}^{(\mathrm{CA3,CA1})} = lrate\ CA1_j\ CA3_k \qquad \text{(A9)}$$

Provided that the network is exposed to the same pattern for more than one learning iteration, the assumption that the CA3 and CA1 activations have converged to stable values for each pattern is a reasonable one. In practice, we find that 2–3 iterations of learning on the same set of patterns is sufficient for the learning to converge to stable performance levels. This method produced good results in all simulations, including paired associate learning.

The simulations reported here therefore employed learning rules given by Eqs. (A6), (A8), and (A9).

REFERENCES

Acsady L, Katona I, Martinez-Guijarro F, Buzsaki G, Freund T. 2000. Unusual target selectivity of perisomatic inhibitory cells in the hilar region of the rat hippocampus. J Neurosci 20:6907–6919.

Amaral DG, Ishizuka N, Claiborne B. 1990. Neurons, numbers, and the hippocampal network. Prog Brain Res 83:1–11.

August D, Levy W. 1999. Temporal sequence compression by an integrate-and-fire model of hippocampal area ca3. J Comput Neurosci 6:71–90.

Barnes CA, McNaughton BL, Mizumori SJY, Leonard BW. 1990. Comparison of spatial and temporal characteristics of neuronal activity in sequential stages of hippocampal processing. Prog Brain Res 83:287–300.

Baum EB, Moody J, Wilczek F. 1988. Internal representations for associative memory. Biol Cybern 59:217–228.

Becker S, Lim J. 2003. A computational model of prefrontal control in free recall: strategic memory use in the california verbal learning task. J Cogn Neurosci 15:347–374.

Becker S, Chin A, Meeds T. 1999. Modelling episodic memory: a global cost function that leads to fast, local, high-capacity learning. In: Proceedings of the Learning workshop, Snowbird, Utah (abstract).

Bibbig A, Wennekers T, Palm G. 1995. A neural network model of the cortico-hippocampal interplay and the representation of contexts. Behav Brain Res 66:169–175.

Bliss T, Lomo T. 1973. Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path. J Physiol (Lond) 232:331–356.

Brown T, Johnston D. 1983. Voltage-clamp analysis of mossy fiber synaptic input to hippocampal neurons. J Neurophysiol 50:487–507.

Buzsaki G. 1998. Memory consolidation during sleep: a neurophysiological perspective. J Sleep Res Suppl 1:17–23.

Cameron HA, McKay RD. 2001. Adult neurogenesis produces a large pool of new granule cells in the dentate gyrus. J Comp Neurol 435:406–417.

Chambers R, Potenza M, Hoffman R, Miranker W. 2004. Simulated apoptosis/neurogenesis regulates learning and memory capabilities of adaptive neural networks. Neuropsychopharmacology 29:747–758.

Churchland P, Sejnowski T. 1994. The computational brain. Cambridge, MA: MIT Press.

Commins S, Cunningham L, Harvey D, Walsh D. 2003. Massed but not spaced training impairs spatial memory. Behav Brain Res 139:215–223. CominsEtA103.pdf.

Deisseroth K, Singla S, Toda H, Monje M, Palmer T, Malenka R. 2004. Excitation-neurogenesis coupling in adult neural stem/progenitor cells. Neuron 42:535–552.

Fellenz W, Taylor J. 2003. The hidden-layer model of hippocampus. Neurocomputing 50:31–50.

Feng R, Rampon C, Tang YP, Shrom D, Jin J, Kyin M, Sopher B, Miller MW, Ware CB, Martin GM, Kim SH, Langdon RB, Sisodia SS, Tsien JZ. 2001. Deficient neurogenesis in forebrain-specific presenilin-1 knockout mice is associated with reduced clearance of hippocampal memory traces. Neuron 32:911–926.

Fortin NJ, Agster KL, Eichenbaum HB. 2002. Critical role of the hippocampus in memory for sequences of events. Nat Neurosci 5:458–462. FortinAgsterEichenbaum02.pdf.

Gerstner W, Abbott L. 1997. Learning navigational maps through potentiation and modulation of hippocampal place cells. J Comput Neurosci 4:79–94.

Gluck MA, Myers CE. 1993. Hippocampal mediation of stimulus representation: a computational theory. Hippocampus 3:491–516.

Gould E, Cameron H, Daniels D, Woolley C, McEwen B. 1992. Adrenal hormones suppress cell division in the adult rat dentate gyrus. J Neurosci 12:3642–3650.

Gould E, Tanapat P, McEwen B, Flugge G, Fuchs E. 1998. Proliferation of granule cell precursors in the dentate gyrus of adult monkeys is diminished by stress. Proc Nat Acad Sci USA 95:3168–3171.

Gould E, Beylin A, Tanapat P, Reeves A, Shors T. 1999a. Learning enhances adult neurogenesis in the adult hippocampal formation. Nat Neurosci 2:260–265.

Gould E, Tanapat P, Hastings N, Shors T. 1999b. Neurogenesis in adulthood: a possible role in learning. Trends Cognitive Neurosci 3:186–192.

Greenstein Y, Pavlides C, Winston J. 1988. Long-term potentiation in the dentate gyrus is preferentially induced at theta rhythm periodicity. Brain Res 438:331–334.

Hasselmo M. 1999. Neuromodulation: acetylcholine and memory consolidation. Trends Cogn Sci 9:351–359.

Hasselmo M, McGaughy J. 2004. High acetylcholine levels set circuit dynamics for attention and encoding and low acetylcholine levels set dynamics for consolidation. Prog Brain Res 145:207–231.

Hasselmo M, Schnell E. 1994. Laminar selectivity of the cholinergic suppression of synaptic transmission in rat hippocampal region CA1: computational modeling and brain slice physiology. J Neurosci 14:3898–3914.

Hasselmo M, Schnell E, Barkai E. 1995. Dynamics of learning and recall at excitatory recurrent synapses and cholinergic modulation in rat hippocampal region CA3. J Neurosci 15(7) (Part 2):5249–5262.

Hasselmo M, Wyble B. 1997. Simulation of the effects of scopolamine on free recall and recognition in a network model of the hippocampus. Behav Brain Res 89:1–34.

Hasselmo M, Wyble B, Wallenstein G. 1996. Encoding and retrieval of episodic memories: role of cholinergic and gabaergic modulation in the hippocampus. Hippocampus 6:693–708.

Healy S, Krebs J. 1996. Food storing and the hippocampus in Paridae. Brain Behav Evol 47:195–199.

Hensbroek R, Kamal A, Baars A, Verhage M, Spruijt B. 2003. Spatial, contextual and working memory are not affected by the absence of mossy fiber long-term potentiation and depression. Behav Brain Res 138:215–223.

Herreras O, Solis J, Herranz A, Martin del Rio R, Lerma J. 1988. Sensory modulation of hippocampal transmission. II. Evidence for

a cholinergic locus of inhibition in the Schaffer-CA1 synapse. Brain Res 461:303–313.
Hyman J, Wyble B, Goyal V, Rossi C, Hasselmo M. 2003. Stimulation in hippocampal region ca1 in behaving rats yields ltp when delivered to the peak of theta and ltd when delivered to the trough. J Neurosci 23:11725–11731.
Jacobs B, van Praag H, Gage F. 2000. Depression and the birth and death of brain cells. Am Sci 88:340.
Jung M, McNaughton B. 1993. Spatial selectivity of unit activity in the hippocampal granular layer. Hippocampus 3:165–182.
Kali S, Dayan P. 2000a. Hippocampally-dependent consolidation in a hierarchical model of neocortex. In Proceedings of Neural Information Processing Systems 2000. p 24–30.
Kali S, Dayan P. 2000b. The involvement of recurrent connections in area CA3 in establishing the properties of place fields: a model. J Neurosci 20:7463–7477.
Kempermann G. 2002. Why new neurons? Possible functions for adult hippocampal neurogenesis. J Neurosci 22:635–638.
Kempermann G, Kuhn H, Gage F. 1997. More hippocampal neurons in adult mice living in an enriched environment. Nature 386: 493–495.
Lasalle J, Bataille T, Halley H. 2000. Reversible inactivation of the hippocampal mossy fiber synapses in mice impairs spatial learning, but neither consolidation nor memory retrieval, in the morris navigation task. Neurobiol Learn Mem 73:243–257.
Lemaire V, Koehl M, Le Moal M, Abrous DN. 2000. Prenatal stress produces learning deficits associated with an inhibition of neurogenesis in the hippocampus. Proc Nat Acad Sci USA 97:11032–11037.
Levy WB. 1996. A sequence predicting CA3 is a flexible associator that learns and uses context to solve hippocampal-like tasks. Hippocampus 6:579–590.
Levy WB, Colbert CM, Desmond NL. 1990. Elemental adaptive processes of neurons and synapses: a statistical/computational perspective. In: Gluck MA, Rumelhart DE, editors. Neuroscience and connectionist models. Hillsdale, New Jersey: Lawrence Erlbaum Associates. p 187–235.
Lisman J. 1999. Relating hippocampal circuitry to function: recall of memory sequences by reciprocal dentate-CA3 interactions. Neuron 22:233–242.
Lynch G, Larson J, Staubli U, Granger R. 1991. Variants of synaptic potentiation and different types of memory operations in hippocampus and related structures. In: Squire L, Weinberger N, Lynch G, McGaugh J editors. Memory: organization and locus of change. Oxford University Press. p 330–363.
MacQueen G, Campbell S, McEwen B, Macdonald K, Amano S, Joffe R, Nahmias C, Young L. 2003. Course of illness, hippocampal function, and hippocampal volume in major depression. Proc Nat Acad Sci USA 100:1387–1392.
Maguire EA, Gadian DG, Johnsrude IS, Good CD, Ashburner J, Frackowiak RS, Frith CD. 2000. Navigation-related structural change in the hippocampi of taxi drivers. Proc Nat Acad Sci USA 97:4398–4403.
Maguire EA, Spiers H, Good C, Hartley ,T, Frackowiak RS, Burgess N. 2003. Navigation expertise and the human hippocampus: a structural brain imaging analysis. Hippocampus 13:250–259.
Marr A. 1971. Simple memory: a theory for archicortex. Philos Trans R Soc Lond B Biol Sci 262:23–81.
McClelland J, Goddard N. 1996. Considerations arising from a complementary learning systems perspective on hippocampus and neocortex. Hippocampus 6:654–665.
McClelland JL, McNaughton BL, O'Reilly RC. 1995. 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 102:419–457.
McEwen B, Magarinos A. 2001. Stress and hippocampal plasticity: implications for the pathophysiology of affective disorders. Hum Psychopharmacol 16:S7–S19.
McNaughton BL, Morris RGM. 1987. Hippocampal synaptic enhancement and information storage within a distributed memory systems. Trends Neurosci 10:408–415.
McNaughton BL, Barnes CA, Meltzer J, Sutherland R. 1989. Hippocampal granule cells are necessary for normal spatial learning but not for spatially-selective pyramidal cell discharge. Exp Brain Res 76:485–496.
Moll M, Miikkulainen R. 1997. Convergence-zone episodic memory: analysis and simulations. Neural Netw 10:1017–1036.
Moscovitch M. 1982. Multiple dissociations of function in amnesia. In: Cermak L editor. Human memory and amnesia. Hillsdale, NJ: Lawrence Erlbaum.
Murre J. 1999. Interaction of cortex and hippocampus in a model of amnesia and semantic dementia. Rev Neurosci 10:267–278.
Nadel L, Samsonovich A, Ryan L, Moscovitch M. 2000. Multiple trace theory of human memory: computational, neuroimaging, and neuropsychological results. Hippocampus 10:352–368.
Nakazawa K, Quirk M, Chitwood R, Watanabe M, Yeckel M, Sun L, Kato A, Carr C, Johnston D, Wilson M, Tonegawa S. 2002. Requirement for hippocampal CA3 NMDA receptors in associative memory recall. Science 297:211–218.
Norman KA, O'Reilly R. Modeling hippocampal and neocortical contributions to recognition memory: a complementary learning systems approach. Psychol Rev (in press).
Nottebohm F. 2002. Why are some neurons replaced in adult brain? J Neurosci 22:624–628.
O'Keefe J, Dostrovsky J. 1971. The hippocampus as a spatial map: preliminary evidence from unit activity in the freely-moving rate. Brain Res 34:171–175.
O'Reilly R, McClelland J. 1994. Hippocampal conjunctive encoding, storage, and recall: avoiding a tradeoff. Hippocampus 4:661–682.
O'Reilly R, Rudy J. 2001. Conjunctive representations in learning and memory: principles of cortical and hippocampal function. Psychol Rev 108:311–345.
Rolls E. 1989. Functions of neural networks in the hippocampus and neocortex in memory. In: Byrne JH, Berry WO editors. Neural models of plasticity: theoretical and empirical approaches. San Diego, CA: Academic Press. p 240–265.
Rosenbaum R, Winocur G, Moscovitch M. 2001. New views on old memories: reevaluating the role of the hippocampal complex. Behav Brain Res 127:183–197.
Rumelhart DE, Zipser D. 1986. Feature discovery by competitive learning. In: Rumelhart DE, McClelland JL, the PDP research group editors. Parallel distributed processing: explorations in the microstructure of cognition. Vol. I. Cambridge, MA: Bradford Books.
Rumelhart DE, Hinton GE, Williams RJ. 1986. Learning internal representations by back-propagating errors. Nature 323:533–536.
Salin P, Scanziani M, Malenka R, Nicoll R. 1996. Distinct short-term plasticity at two excitatory synapses in the hippocampus. Proc Nat Acad Sci USA 93:13304–13309.
Shors T, Miesegaes G, Beylin A, Zhao M, Rydel T, Gould E. 2001. Neurogenesis in the adult is involved in the formation of trace memories. Nature 410:372–376.
Spreng M, Rossier J, Schenk F. 2002. Spaced training facilitates long-term retention of place navigation in adult but not in adolescent rats. Behav Brain Res 128:103–108. SprengRossierSchenk02.pdf.
Stringer S, Rolls E, Trappenberg T, de Araujo I. 2002. Self-organizing continuous attractor networks and path integration: two-dimensional models of place cells. Network 13:429–446.
Treves A, Rolls ET. 1992. Computational constraints suggest the need for two distinct input systems to the hippocampal CA3 network. Hippocampus 2:189–200.
van Praag H, Christie B, Sejnowski T, Gage F. 1999. Running enhances neurogenesis, learning, and long-term potentiation in mice. Proc Nat Acad Sci USA 96:13427–13431.

Wallenstein GV, Hasselmo ME. 1997. GABAergic modulation of hippocampal place cell activity: sequence learning, place field development, and the phase precession effect. J Neurophysiol 78:393–408.
Wiebe S, Staubli U, Ambros-Ingerson J. 1997. Short-term reverberant memory model of hippocampal field CA3. Hippocampus 7:656–665.
Willshaw D, Dayan P. 1990. Optimal plasticity from matrix memories: What goes up must come down. Neural Comput 2:85–93.
Willshaw DJ, Buckingham JT. 1990. An assessment of Marr's theory of the hippocampus as a temporary memory store. Philos Trans R Soc (Lond) B Biol Sci 329:205–215.
Wiskott L, Rasch M, Kempermann G. 2005. What is the functional role of adult neurogenesis in the hippocampus? Cognitive Sciences EPrint Archive (CogPrints) 4012 (http://cogprints.org/4012/).
Witter M, Groenewegen H, Lopes da Silva F, Lohman A. 1989. Functional organization of the extrinsic and intrinsic circuitry of the parahippocampal region. Prog Neurobiol 13:161–253.
Yeckel M, Berger T. 1990. Feedforward excitation of the hippocampus by afferents from the entorhinal cortex: redefinition of the role of the trisynaptic pathway. Proc Nat Acad Sci USA 87:5832–5836.
Yonelinas A, Kroll N, Quamme J, Lazzara M, Sauve M, Widaman K, Knight R. 2002. Effects of extensive temporal lobe damage or mild hypoxia on recollection and familiarity. Nat Neurosci 5:1236–1241.