The Race to Learn: Spike Timing and STDP Can Coordinate Learning and Recall in CA3

Christopher R. Nolan,[1,2]* Gordon Wyeth,[1] Michael Milford,[1,2] and Janet Wiles[1,2]

ABSTRACT: The CA3 region of the hippocampus has long been proposed as an autoassociative network performing pattern completion on known inputs. The dentate gyrus (DG) region is often proposed as a network performing the complementary function of pattern separation. Neural models of pattern completion and separation generally designate explicit learning phases to encode new information and assume an ideal fixed threshold at which to stop learning new patterns and begin recalling known patterns. Memory systems are significantly more complex in practice, with the degree of memory recall depending on context-specific goals. Here, we present our spike-timing separation and completion (STSC) model of the entorhinal cortex (EC), DG, and CA3 network, ascribing to each region a role similar to that in existing models but adding a temporal dimension by using a spiking neural network. Simulation results demonstrate that (a) spike-timing dependent plasticity in the EC-CA3 synapses provides a pattern completion ability without recurrent CA3 connections, (b) the race between activation of CA3 cells via EC-CA3 synapses and activation of the same cells via DG-CA3 synapses distinguishes novel from known inputs, and (c) modulation of the EC-CA3 synapses adjusts the learned versus test input similarity required to evoke a direct CA3 response prior to any DG activity, thereby adjusting the pattern completion threshold. These mechanisms suggest that spike timing can arbitrate between learning and recall based on the novelty of each individual input, ensuring control of the learn-recall decision resides in the same subsystem as the learned memories themselves. The proposed modulatory signal does not override this decision but biases the system toward either learning or recall. The model provides an explanation for empirical observations that a reduction in novelty produces a corresponding reduction in the latency of responses in CA3 and CA1. © **2010 Wiley-Liss, Inc.**

KEY WORDS: spike timing; STDP; modulation; pattern separation; pattern completion

INTRODUCTION

The ability to switch between learning new information and using known information is of fundamental importance to memory systems. Such systems must learn new information when possible to aid in making future decisions and also use previously acquired knowledge to make ongoing decisions in real time. When an animal explores an environment for the first time, it learns the layout of the space by integrating landmark features and determines the site of any useful goals, such as a food source. On returning to a familiar location, the animal should be able to navigate using the previously learned cues and more efficiently attain its goals. In the real world, however, the appearance of landmarks and other features is often in a state of flux, and the animal may approach the landmarks from different perspectives or at different times of the day. Some features may be changed permanently, whereas others may appear transiently different to the way they were originally perceived. The animal's memory system must be capable of determining which details of the location are new, requiring that a distinct memory be formed (pattern separation), and which details should be recalled from incomplete or noisy information (pattern completion).'

Pattern separation and completion in neural systems has been the subject of many theoretical and computational modeling studies (Marr, 1971; Hopfield, 1982; Treves and Rolls, 1992; O'Reilly and McClelland, 1994; McClelland et al., 1995; Rolls, 1996; Fuhs and Touretzky, 2000), however, not all studies define separation and completion the same way. In the current study, consistent with O'Reilly and McClelland (1994) and Fuhs and Touretzky (2000) among others, we define pattern separation to be the process by which an input with a particular similarity to a reference input generates an output of a lower similarity to the corresponding reference output, and pattern completion to be an increase in output similarity with respect to input similarity (Fig. 1). An alternative definition for pattern completion is that it is the ability for a noisy or partially cued input to produce an output arbitrarily close to 100% similarity to a reference output, such as in the classic Hopfield (1982) model. In this latter sense, pattern completion is most

[1] **School of Information Technology and Electrical Engineering, The University of Queensland, Brisbane, Queensland, Australia;** [2] **Queensland Brain Institute, The University of Queensland, Brisbane, Queensland, Australia**

Additional Supporting Information may be found in the online version of this article.

Grant sponsors: Tony Murphy Postgraduate Scholarship (to CN), Australian Research Council (Special Research Initiative: Thinking Systems); Grant number: TS0669699.

Abbreviations used: DG, dentate gyrus; EC, entorhinal cortex; LTP, long-term potentiation; O-LM, oriens-lacunosum moleculare; STDP, spike-timing-dependent-plasticity; STSC, spike-timing separation and completion.

*Correspondence to: Christopher Nolan, School of Information Technology and Electrical Engineering, The University of Queensland, St Lucia, QLD 4072, Australia. E-mail: cnolan@itee.uq.edu.au

Accepted for publication 4 January 2010
DOI 10.1002/hipo.20777
Published online in Wiley InterScience (www.interscience.wiley.com).

© *2010 WILEY-LISS, INC.*

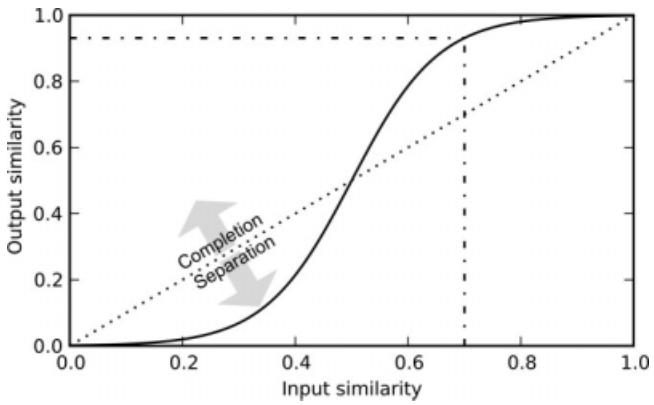

FIGURE 1. Pattern separation and completion. The graph shows the transfer function from input (x-axis) to output (*y*-axis). Each axis measures the similarity of a pattern to a reference pattern. The $y = x$ line (dotted) shows when the output similarity would be the same as the input similarity. When the transfer function (solid line) is below the $y = x$ line, then output patterns will be less similar than their input patterns (separation). When the transfer function is above the $y = x$ line, then output patterns will be more similar than their input patterns (completion). For example, for the transfer function shown, an input similarity of 0.7 generates an output similarity greater than 0.9 (see dot-dash line).

commonly achieved using recurrent autoassociative networks; however, heteroassociative networks with feedback from the output back to the input are similarly capable. Neither autoassociative nor heteroassociative feedback is a computational requirement for pattern completion, in the sense in which we use it in this study.

Many of these theoretical models are similar in network structure to that of the hippocampus, some by design (Marr, 1971; Treves and Rolls, 1992; Rolls, 1996) and others seemingly by chance (Hopfield, 1982). This similarity, along with the observations that the hippocampus appears to be required for episodic memory formation in humans (Scoville and Milner, 1957) and for complex spatial memory tasks (locale navigation) in animals (Morris, 2007), supports the hypothesis that the region may be performing separation and completion. One mechanism allowing separation and completion in a single neural region—initially proposed by Treves and Rolls (1992) and based on known hippocampal anatomy—suggests that the disynaptic (indirect) pathway from Layer II of the entorhinal cortex (EC2) to CA3 via dentate gyrus (DG) and the monosynaptic (direct) pathway from EC2 to CA3 fulfill pattern separation and completion roles respectively. Substantial experimental data now support this proposal: activity in DG suggests a role of pattern separation for the region (Leutgeb et al., 2007) and activity in CA3 suggests that inputs sufficiently similar to a previously learned input evoke the same learned response (completion), but that below some input similarity threshold, the response rapidly decorrelates from the learned response (separation) (Vazdarjanova and Guzowski, 2004; Leutgeb et al., 2007). Furthermore, disruption of the indirect pathway interferes with the encoding of new memories but not the retrieval of existing memories, and disruption of the direct pathway interferes with retrieval but not encoding (Lee and Kesner, 2004). Downstream CA1 activity is similar to that in CA3 but with less environmental specificity (more completion) (Leutgeb et al., 2004; Vazdarjanova and Guzowski, 2004). Although this evidence is largely from rodent studies, new imaging studies suggest that a similar functional distinction between regions may be present in humans (Bakker et al., 2008). Together, these data begin to explain what processes are facilitating memory storage and where these processes occur, but it remains unclear when each process is occurring, that is, when new details should be learned and when they should be recalled.

Determining when to learn and when to recall requires a mechanism to detect when an input is novel. Intuitively, it seems as though such a decision should depend on the patterns stored within the hippocampus itself, thereby requiring that the structure mediate its own novelty signal. Treves and Rolls (1992) identified the mechanistic consequences of novelty on learning and recall that were required for their aforementioned dual pathway hypothesis. During learning, the indirect input should be strong in comparison with other CA3 inputs (for example, the direct input and recurrent connections) to allow new associations to be learned. During retrieval, the indirect input should be rendered effectively inactive to avoid spurious separation noise—input noise amplified by the separation properties of the DG—interfering with retrieval and potentially disrupting learned memories. Although these requirements are recognized, their underlying biological mechanisms remain unclear. Experimental data from Hasselmo et al. (1995) implicate the septo-hippocampal cholinergic system in downregulation of the CA3-CA3 recurrent synapses during learning; however, this modulation operates at a multisecond timescale (Hasselmo and Fehlau, 2001). This multisecond time scale may be too slow for the purposes of online learning and recall, leading Hasselmo et al. (2002) to propose that pattern-by-pattern learning may be accounted for by explicitly modulating between completion and separation in separate phases of the hippocampal theta rhythm. However, with an explicit learning phase, separation noise interferes even with well-known patterns. None of these models internally generate a pattern-by-pattern novelty signal and use that signal as a basis for switching between learning and recall. The question arises, what mechanism supports such a signal, and could it be generated within the hippocampus itself?

We sought to establish a mechanism capable of deciding when to learn a novel input pattern and when to recall by completing to a previously learned pattern, using a decision criterion based on patterns currently stored in the system. Here, we demonstrate a new computational model of the dual-pathway circuit formed by EC2, DG, and CA3 that implicitly performs pattern-by-pattern novelty detection. We call this model the spike-timing separation and completion (STSC) model. This model incorporates spike time as an integral aspect of neurons, which enables it to distinguish between patterns that have been previously seen, and patterns that are new and should be learned. Spike timing not only provides the ability to discriminate between known and unknown patterns through spike-timing-dependent-plasticity (STDP) (Levy and Steward, 1983; Markram et al., 1997; Song et al., 2000) but also inher-

ently drives the mechanism to effect or suppress learning. This pattern-by-pattern suppression ensures that even in unfamiliar situations, already known patterns are not relearned; whereas in familiar situations, unknown patterns can be learned. We begin by outlining a three-neuron network to illustrate the principle of using spike-timing as a correlate of learning in a dual-pathway network. We then present the complete STSC model, showing that (i) a larger network based on EC2-DG-CA3 anatomical details correlates learning and timing, (ii) the network learns patterns that are both reliable and distinct, and (iii) modulation can tune the network to achieve an appropriate balance between pattern separation and pattern completion. Finally, we examine the predictions this model generates regarding both the timing of activity in the hippocampus and the required neuromodulatory effects.

SPIKE TIMING AND THE DUAL-PATHWAY HYPOTHESIS: AN ILLUSTRATIVE MODEL

STDP is a biological manifestation of Hebb's theoretical learning rule (Song et al., 2000). In its most commonly observed form, synapses connected to a cell are strengthened when they fire in a time window prior to the cell firing, and are weakened when they fire in a time window after the cell fires. This temporal dependence means that afferent synapses that predict a cell's firing will have a greater postsynaptic effect on their next activation. If a cell is activated after (but independently of) the activation of a particular set of synapses, the cell will gradually spike earlier with respect to the activation of the synapses (Guyonneau et al., 2005). Furthermore, assuming this synaptic set is sufficiently large; the cell will eventually fire in response to the predictive synapses alone. This behavior was elegantly demonstrated by Izhikevich (2007) using a conditioning paradigm, in which the cells originally activated by an unconditioned stimulus were eventually activated by synapses receiving input because of a conditioned stimulus. In Izhikevich's model, spike timing encodes whether a response has been learned with a shorter latency between input and response, indicating a well-known input.

STDP requires that the postsynaptic cell be active to learn active inputs. In the conditioning scenario of Izhikevich (2007), the unconditioned stimulus activated the cells that the conditioned stimulus predicted. In the context of generalized memory, the same sensory stimuli can be considered as both the target to be learned (conditioned stimulus) and the cue that is triggering learning (unconditioned stimulus). A mechanism is thus required whereby a novel input can evoke a pattern of activity in a target cell population that can then be associated with the input. The dual-pathway architecture (already discussed) provides such a mechanism: the indirect pathway does not undergo associative plasticity and enables novel input to evoke baseline activity; the direct pathway is subject to associative long-term potentiation (LTP), thus pairing presynaptic and postsynaptic activity. We investigated the effect of the temporal dynamics of STDP in a simplified dual-pathway model. Consistent with previous dual-pathway models, synapses from the two pathways differ in strength and plasticity; however, in the current model, they also differ in the latency of their initial activation following a given input stimulus. Using this model, we demonstrate that the distinction between a fixed and learned response is retained in the response of the target cell if the temporal dynamics are considered.

The network consists of three neurons, each one is a representative of one region (Fig. 2a). Individual neurons are modeled as point devices, following the model presented by Izhikevich (2003, 2004). Two-state variables are required for each neuron, the membrane potential v and the recovery current u. The rates of change of v and u are given by

$$\dot{v} = 0.04v^2 + 5v + 140 - u + I \quad (1)$$

$$\dot{u} = a(bv - u) \quad (2)$$

with a reset condition for both variables given by

$$\text{if } v \geq 30\ mV, \text{then} \begin{cases} v \leftarrow c \\ u \leftarrow u + d \end{cases} \quad (3)$$

where I is the sum of synaptic input currents, parameters a and b determine whether the neuron is an integrator or a resonator, and parameters c and d model the postspike transient behavior because of high-threshold voltage-gated channel dynamics (parameter values set as suggested by Izhikevich (2003) for achieving these behaviors; see the parameter table in the Supporting Information).

Each STDP synapse is modeled as a single-state variable I_s, representing the current injection to its postsynaptic neuron

$$\dot{I}_s = -I_s/\tau_s + \delta(t - \mathbf{t}_s)w_s \quad (4)$$

where τ_s is the decay time constant of synapse s, δ is the Dirac delta function, $\mathbf{t}_s$ is the vector of presynaptic spike times, t is time, and w_s is the current synaptic weight. Synapses subject to STDP follow the model of Song et al. (2000). For each pair of presynaptic and postsynaptic spikes, the change in synaptic weight w is given by

$$\Delta w = \max(\min(F(\Delta t), w_{\max} - w), -w) \quad (5)$$

where $F(\Delta t)$ is a nonlinear function of the post–pre interspike interval Δt given by

$$F(\Delta t) = \begin{cases} A_+ \exp(-\Delta t/\tau_+) & \text{if } \Delta t \geq \tau_{\text{gap}} \\ -A_- \exp(\Delta t/\tau_-) & \text{if } \Delta t < -\tau_{\text{gap}} \end{cases} \quad (6)$$

where A_+ and A_- determine the maximum modification for a positive and negative change, respectively, τ_+ and τ_- are time constants that determine the time range under which synaptic strengthening and weakening occur, and τ_{gap} is the minimum time difference required for any weight change. Only neighboring presynaptic and postsynaptic pairs are considered, and Eq. (5) ensures that the weight of each synapse is non-negative, with an upper bound $w_{\max}$.

The synaptic current for static synapses I_s is given by

$$I_s = \sum_{r \epsilon \mathbf{R}} g_r(v - E_r) \quad (7)$$

where R is a set of receptor types, in this case containing only AMPA (α-amino-3-hydroxyl-5-methyl-4-isoxazole-propionate), v is the membrane potential of the postsynaptic neuron, E_r is

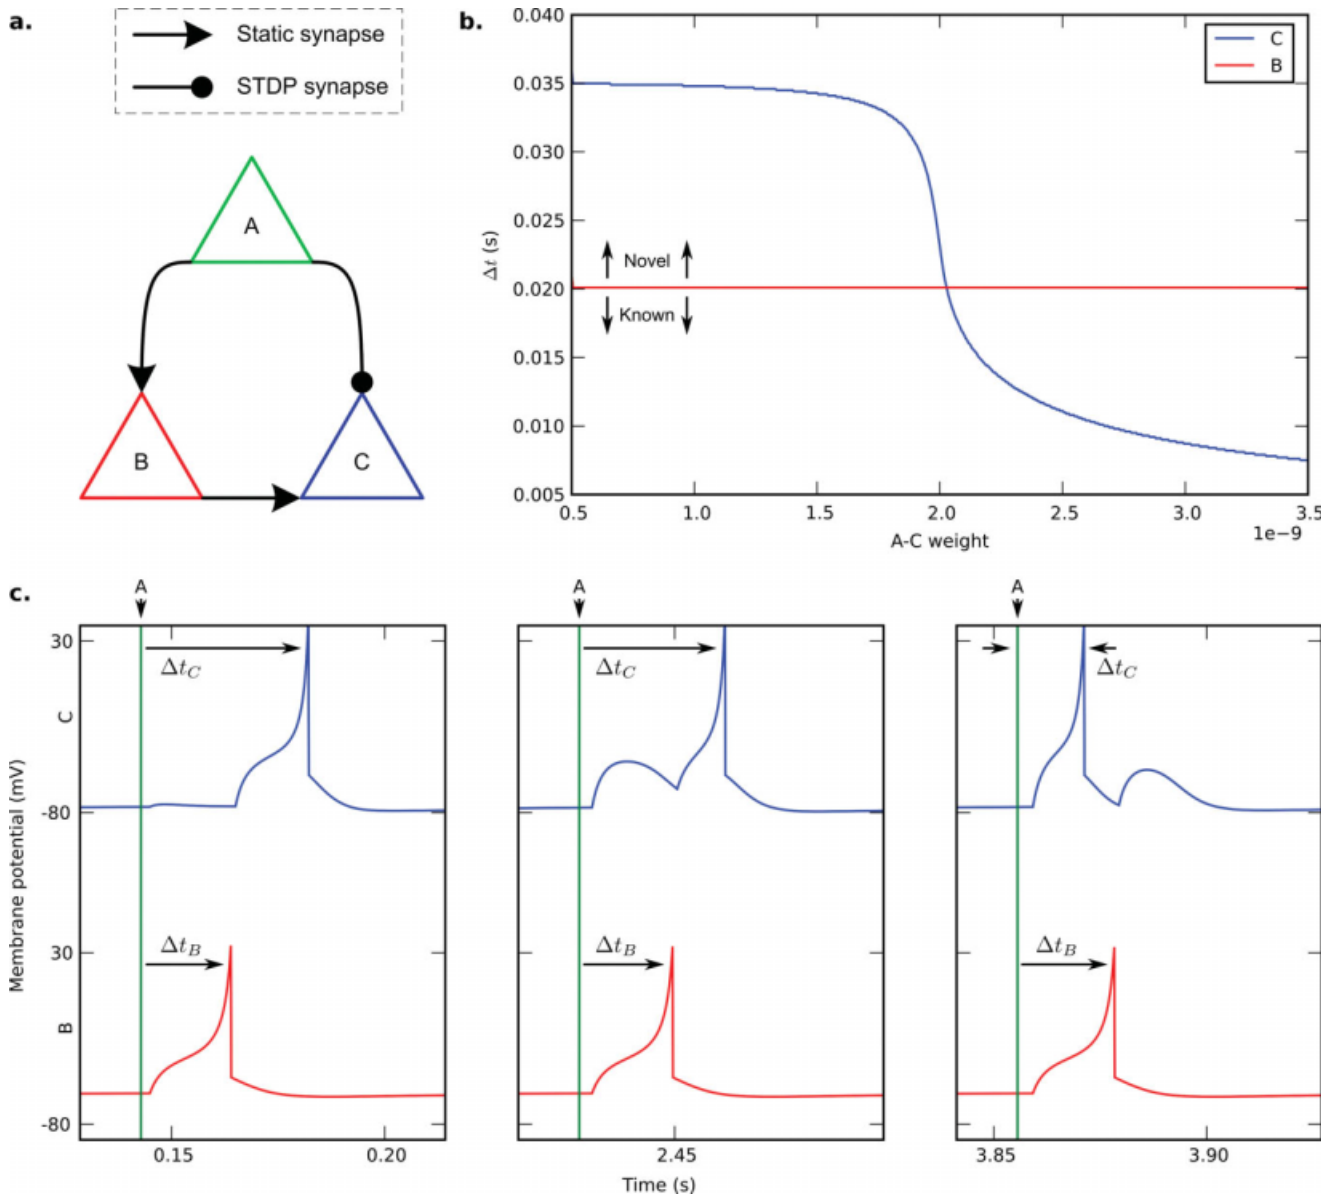

FIGURE 2. How timing can signify the distinction between novel and familiar inputs. (a) Schematic illustration of a three-field network with two paths from A to C, an indirect path A-B-C and a direct path A-C. (b) The latency of the response of cell B relative to cell A is fixed (see horizontal dotted line), the response of cell C relative to cell A has a latency that is a function of the strength of the A-C synapse (see solid line). With a novel input (A-C weight < 2.03×10^{-9}), the signal travels via the indirect path A-B-C and the latency of C is longer than that of B (the solid line above the dotted line indicates that A is novel). With a known input (A-C weight > 2.03×10^{-9}), the signal travels via the direct path A-C and C's latency is shorter than that of B (the solid line below the dotted line indicates that A has previously been learned). Between these two conditions (A-C weight = 2.03×10^{-9}), the two pathways contribute equally to the timing of cell C's action potential. The latency, therefore, indicates which path, A-B-C or A-C, is dominant. (c) Three panels show snapshots of the behavior for novel, near-threshold and known weights during a learning trial: for low A-C weights, the response of C depends entirely on the input from B (c, left). As the A-C weight approaches a threshold (2.03×10^{-9}), less input from B is required as the membrane potential of C is still elevated due to the direct input from A, and thus the time of C's response is closer to that of B (c, middle). Beyond this threshold, cell C is responding before cell B, showing that the direct A-C synapse is solely responsible for C's response and signifying that the input is known (c, right). [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

the reversal potential of the receptor, and g_r is the current conductance of the receptor given by

$$\dot{g} = -g/\tau_r + \delta(t - \mathbf{t}_s)w_s \quad (8)$$

in which τ_r is the decay time constant of the receptor (Dayan and Abbott, 2001). See the parameter Tables in the Supporting Information for parameter values.

In the example ABC-AC network, synaptic weights are set such that a single spike from neuron A is capable of causing a spike in neuron B, and in turn a single spike from cell B is capable of causing a spike in neuron C. The weights from A to B and from B to C are static, that is, unaffected by activity. The direct synapse from A to C has a small initial weight, but

is modified according to the above STDP rule, such that a fully potentiated synapse can evoke an action potential in C.

The simulation protocol involved stimulating input neuron A 28 times at 7 Hz and recording membrane voltages of neurons B and C (Figs. 2b,c). The 7-Hz stimulation frequency was chosen as a typical theta rhythm frequency in the hippocampus and the EC. Initially, the direct input had little impact on the membrane potential of cell C, which spiked only after it received input from neuron B. Over repeated presentations, the direct synapse was potentiated, and eventually became strong enough to directly evoke an action potential in cell C. At this stage, when the corresponding input from neuron B arrived the cell was still partially hyperpolarized, as a result of which the indirect input was no longer sufficient to cause a postsynaptic potential.

The timing of the spikes of cell C relative to those of cell A indicates whether the postsynaptic action potential is caused by the direct stimulation through the A-C synapse or the indirect stimulation through the B-C synapse. If cell C fires after the input from cell B arrives, the stimulus can be considered unknown, and thus requires the input from B to enable learning. If cell C fires before cell B, the stimulus is known, and any input via the indirect pathway can be ignored. Furthermore, when cell C fires before cell B it does not also fire after cell B (due to hyperpolarization), demonstrating that the circuit intrinsically selects between learning and recall for each individual input. The relative timing of firing between cell A and cell C inherently coordinates two functions: (1) it encodes the novelty of the association between cells A and C and (2) it controls whether the A-C synapse is modified by the B-C input. These results indicate that the dual-pathway architecture, when implemented with a spiking neural input and STDP, provides a conceptually simple, integrated mechanism to distinguish novel from known input. However, to similarly signal when to learn based on a network of inputs, this mechanism must be expanded to detect novelty of not just a single cell input, but a pattern of active cells. In the next section, we scale up the three-cell network to model the EC-DG-CA3 circuit, demonstrating how spike timing can control when subsets of CA3 neurons learn new patterns, and when they recall known patterns.

THE STSC MODEL

Like the hypothesized dual-pathway architecture of Treves and Rolls (1992), anatomical data suggest that cells in CA3 receive external excitatory input from EC2 both directly via the perforant path and indirectly through DG and its mossy fiber system (Amaral, 1993). The projections from DG to CA3 are sparse: each dentate granule cell forms synapses on ~14 CA3 neurons (Claiborne et al., 1986). Dynamic facilitation in the mossy synapses on CA3 cells allows a burst from even a single dentate granule cell to evoke a postsynaptic action potential (Henze et al., 2002). The direct synapses from EC2 to CA3 number in the thousands per CA3 cell, but each individual synapse has a much smaller effect on the postsynaptic potential than the DG inputs. Plasticity characteristics also differ between the two synaptic paths. Although LTP is evident in the DG to CA3 synapses, they show little evidence of associative modification (Chattarji et al., 1989), whereas associative LTP is evident in perforant path inputs to both DG (Bliss and Lømo, 1973) and to CA3 (Do et al., 2002). The CA3 cells are also self-connected through the recurrent collateral pathway. Like the perforant path, recurrent collateral synapses exhibit associative LTP (Chattarji et al., 1989; Martinez et al., 2002), and although numerous, are individually weak.

Pattern separation requires the prewired, sparse connections in the DG-CA3 pathway to provide an index into the CA3 population. Activation in EC2 evokes a response in DG, which in turn causes strong activation of a small number of CA3 cells. Hebbian-like direct synapses connecting active EC2 cells to the most active neurons are potentiated, as are recurrent links between active CA3 neurons. The EC2-CA3 connections learn to associate EC2 stimuli with patterns in CA3, while the recurrent CA3-CA3 connections learn to self-associate patterns in CA3. The recurrent connections provide pattern completion for partial or noisy stimuli from EC2. Fuhs and Touretzky (2000) demonstrated in a computational model that a similar effect can be achieved with competitive inhibition in CA3 instead of excitatory recurrence.

In the previously described models, the prewired DG-CA3 connections bias CA3 activity, both when learning new patterns and when recalling familiar patterns. The activity of the CA3 population contains no information about the contribution of the direct EC2-CA3 path relative to that of these prewired DG-CA3 connections. Hasselmo et al. (2002) hypothesized that theta half cycles might alternately facilitate learning and recall in CA1. This idea has since been explored in more detailed modeling scenarios in both CA3 (Kunec et al., 2005) and in CA1 (Cutsuridis et al., in press).

The ABC-AC network of the previous section suggests an alternative solution to the problem of mediating between learning and recall in the hippocampal circuit. If the three regions A, B, and C are considered analogous to EC2, DG, and CA3, respectively, the two pathways (indirect EC2-DG-CA3 and direct EC2-CA3) could act to simultaneously learn and cue memories by virtue of their relative timing. If the direct pathway evokes a sufficient response, the indirect response can be safely ignored. If instead the direct pathway fails to evoke a direct response, the indirect pathway can activate a novel set of CA3 neurons, which can be associated with the input. The dual pathway guarantees not only that a set of neurons will be activated in CA3 but also that the relative timing of the response determines whether the active set corresponds to a familiar pattern that is being reactivated or a novel memory that needs to be learned. The network thus acts simultaneously as a memory system and as a novelty detector.

In the following sections, we illustrate our STSC model of the hippocampus that incorporates many of the known ana-

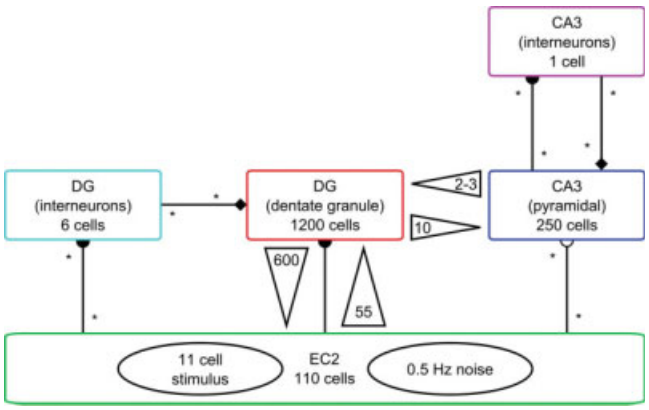

FIGURE 3. Network architecture showing main regions, cell types, and the fan in-fan out connectivity (* indicates full connectivity). Open semicircles indicate plastic excitatory synapses, filled semicircles indicate fixed excitatory synapses and filled diamonds indicate fixed inhibitory synapses. The CA3 region is the primary component of associative memory storage in the system. It receives input from EC2 via two pathways: one is a direct excitatory path, the other is indirect via DG. [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

tomical and functional details of the EC-DG-CA3 circuit to investigate the role of spike timing in learning new patterns and recalling known patterns. We address three key questions:

- Does a larger network based on EC2-DG-CA3 anatomical details show the same clear correlation between learning and timing that was seen in the three-cell network?
- Can the network learn patterns that are both reliable and distinct? Reliability ensures that pattern completion is stable, and distinctiveness provides a measure of the network's discriminative power.
- Can the network be modulated to implement an appropriate balance between pattern separation (finding distinct new patterns to learn) and pattern completion (recalling the learned pattern from an incomplete or noisy input)?

Model Description

The architecture of our STSC model is similar to that of the existing models in terms of the functional roles of each region (Treves and Rolls, 1992; Fuhs and Touretzky, 2000; Kali and Dayan, 2000) (Fig. 3). Regional cell numbers are proportional to estimated numbers of cells in the adult rat brain (Amaral et al., 1990), reduced by a factor of 1,000 (110 EC2, 1,200 DG, 250 CA3). Network input was provided via direct stimulation of the EC2 neurons. Each stimulus activated a specific subset of the EC2 neurons (11/110) at a random time within a 5-ms window. Stimuli were presented at theta frequency (7.0 Hz). In addition to the activity generated by this stimulation, the EC2 input neurons were activated probabilistically at each time step (probability calculated for an average baseline activity of 0.5 Hz).

Connectivity between the regions was assigned so as to achieve sparse activation of the CA3 cell population via the DG-CA3 mossy fiber system. For modeling convenience, all synapses on the indirect pathway (EC2-DG and DG-CA3) were nonplastic. Each CA3 cell was connected to 10 DG cells, randomly distributed but ensuring that each DG cell was connected to two or three CA3 cells. The DG-CA3 synapses had static weights, set such that two synapses active in a 5-μs window provided sufficient input to evoke a postsynaptic action potential.

Excitatory connectivity for the indirect pathway was calculated to ensure that for greater than 99% of possible stimuli as described above, in the absence of noise, at least one CA3 neuron would respond (calculated via the combinatorics of neural population numbers and average input activity). The corresponding minimum active DG cell numbers per stimulus was calculated to achieve this CA3 response. To ensure that this required minimum number of DG cells responded, each dentate granule cell received weighted input from 50% (55) of the EC2 cells, such that a minimum number $\left(n_{\text{req}}^{\text{DG}}\right)$ of active synapses caused any particular DG cell to fire (in this simulation, the minimum number of active synapses was nine).

Each EC neuron was connected to all the CA3 neurons. Initially, the efficacies of the direct EC2 synapses were sufficiently low that none of the presented patterns in EC2 caused activity in CA3. Perforant path inputs were subject to STDP, such that if any one EC2 neuron consistently fired before a particular CA3 cell, then the synapse directly connecting them was potentiated. As the CA3 network is small, any particular novel stimulus may activate only a few CA3 cells. In this situation, all active cells must be capable of learning the input. The 100% connectivity for EC2-CA3 is necessary because of these scaling issues.

Two different types of inhibitory neurons were included. Six inhibitory cells in DG received excitatory input from all EC2 cells. Each of these cells was calibrated to fire when greater than a given number of inputs $\left(n_{\text{req}}^{\text{DGi}}\right)$ were active within a 5-ms window, where each cell required a different minimum number of active inputs ranging from 11 to 16. The effect of this graded inhibition was to restrict the activity of DG despite the varying levels of EC activity in the presence of noise. An inhibitory cell in CA3 received input from the local pyramidal cells, and acted to suppress further CA3 network activity after approximately three action potentials in a 5-ms window $\left(n_{\text{req}}^{\text{CA3i}}\right)$. The purpose of the CA3 inhibitory cell was twofold. First, it ensured that small variations in DG activity (variations not normalized by inhibition in DG) did not cause large variation in the CA3 activation levels (which would then contribute to synaptic plasticity). Second, it acted as a threshold. If sufficient CA3 cells responded to a pattern directly, the CA3 inhibitory cell would signal the pattern as known. The CA3 inhibitory cell was designed to act as a regional equivalent to the hyperpolarization of cell C in the ABC-AC network: if it fired before the indirect pathway caused a response, then it attenuated any such response. The effect of this inhibition enabled the circuit to again intrinsically select between learning and recall for each individual input.

DG and CA3 cells were modeled as spiking neurons with parameters chosen to match the cell properties characterized in a recent review by Spruston and McBain (2007), who describe that a somatic current injection into CA3 pyramidal cells elicits a bursting response, whereas a similar stimulus in a dentate granule cell causes spike trains exhibiting spike frequency adaptation. For the current model, we implemented regular spiking dynamics for the dentate granule cells, intrinsically bursting dynamics for CA3 pyramidal cells and fast-spiking dynamics for all inhibitory cells (neuronal models were as described for the three-neuron model, neuronal parameters were chosen according to the suggestions of Izhikevich (2003); see Supporting Information for parameter values).

As in the three-neuron model (ABC-AC), both static and plastic excitatory synapses were modeled with a single time constant in the range of AMPA channel dynamics [Eq. (7)], and plastic excitatory synapses were modified according to STDP Eq. (4). Inhibitory synapses were modeled according to Eq. (7) but with time constants and reversal potentials in the range of $GABA_A$ (γ-Aminobutyric acid) channel dynamics, and the CA3 inhibitory feedback interneuron had an additional channel with a longer effect akin to the time course of $GABA_B$. This second time course was intended to ensure that, like the C cell in the three-neuron model, all CA3 neurons were hyperpolarized upon arrival of the indirect input if a known pattern was presented.

Timing as a Signal of Learning

The timing between the stimulus and response provides a clear marker of learning for a single neuron. To investigate the relationship between learning and timing in a network structure, we trained the STSC system to respond to varied and noisy stimuli. Each unique EC2 stimulus (as described previously) was presented 28 times sequentially at theta frequency (7.0 Hz). Action potentials from the CA3 network were recorded for 10 sets of 10 different stimuli. As the number of learned stimuli increases, the potential for crosstalk between the weights for these different stimuli increases, and ultimately affects classification performance. Our study was not intended to evaluate network capacity, thus each of 10 networks learned only 10 stimuli each. The difference between the time of each CA3 spike and the mean time of the most recent input stimulus was recorded and grouped for each per-stimulus trial number, that is, for the first instance of each stimulus, the second instance of each stimulus and so on.

Qualitative results of the network being trained to the first three stimuli over 10 s demonstrate the typical behavior of the network (see the raster plot of Fig. 4a). During the first second, no stimuli are presented, and without coordinated input activity, the DG and CA3 networks are silent. In subsequent seconds, the stimulus is applied to EC2. The DG response clearly precedes the CA3 response at the first presentation of each stimulus (Fig. 4b), showing that DG input is initially required to cause CA3 activity. By the final presentation of the same stimulus (Fig. 4c), the temporal order has reversed and the response of the CA3 region precedes the DG response.

The progressive change in spike timing during repeated presentation of a stimulus follows a characteristic pattern (Fig. 4d). Across the range of trials, the variance of the responses initially increases. Between trials 10 and 14 there is a median shift showing the transition between post-DG and pre-DG response. The variance of the response timing—increasing as the response is first learned then gradually decreasing—suggests a consolidation of neural responses. Noise causes variation in the activity of the dentate granule cells, which initially causes corresponding variation in the instantaneous CA3 response. Over several presentations, synapses between the most consistently responsive EC2 and CA3 neurons are strengthened enough to directly effect a postsynaptic action potential. When few CA3 neurons respond directly $\left(n_{\text{resp}}^{\text{CA3}} < n_{\text{req}}^{\text{CA3i}}\right)$, the variance in response times is high. As more CA3 neurons respond directly and reliably to the EC2 stimulation, the variance in response times falls.

Distinctiveness and Reliability of Learned Responses

Although the preceding results show that learning reduces the time of the CA3 response to EC2 stimulation, they do not show whether the patterns that are learned are in fact distinct and reliably reactivated: ideally, the network should demonstrate a capability to learn distinct responses for different input, and furthermore activate these responses reliably when presented with sufficiently similar input.

To investigate the distinctiveness and reliability of the learned responses, the network was trained with a set of three distinct stimuli, each of which activated 10% of EC2 (i.e., 11 out of 110 EC2 neurons). During training, each stimulus was presented 28 times. After training, synaptic plasticity was disabled. Distinct test input patterns were generated for each of the trained patterns, such that none of the neurons active in the trained pattern were also active in the test patterns. A set of 200 distinct test input patterns were chosen in this fashion for each of the three stimuli. To vary the similarity of the input pattern to the trained patterns, testing consisted of stimulating the network with a combination of the trained pattern and the current test pattern. Initially, the whole trained pattern was presented, giving 100% similarity to the trained EC2 stimulus. On each subsequent trial, one neuron from the trained pattern was removed from the active stimulus and replaced by a neuron from the current test pattern, reducing the similarity to the trained input until all neurons corresponded to the test pattern giving 0% similarity. This process was repeated for each of the 600 test patterns (three stimuli $\times$ 200 test patterns each). All input patterns were also subject to noise (as described previously), creating further variation in similarity of the input pattern with the trained pattern.

The separation of responses was determined by comparisons between the sets of responsive cells in CA3 for different input stimuli in EC2. The EC2 stimulus and the CA3 responses can be represented as vectors, where each element in the vector corresponds to a neuron in the respective regions. Elements in the vector were set to one if the associated neuron fired between

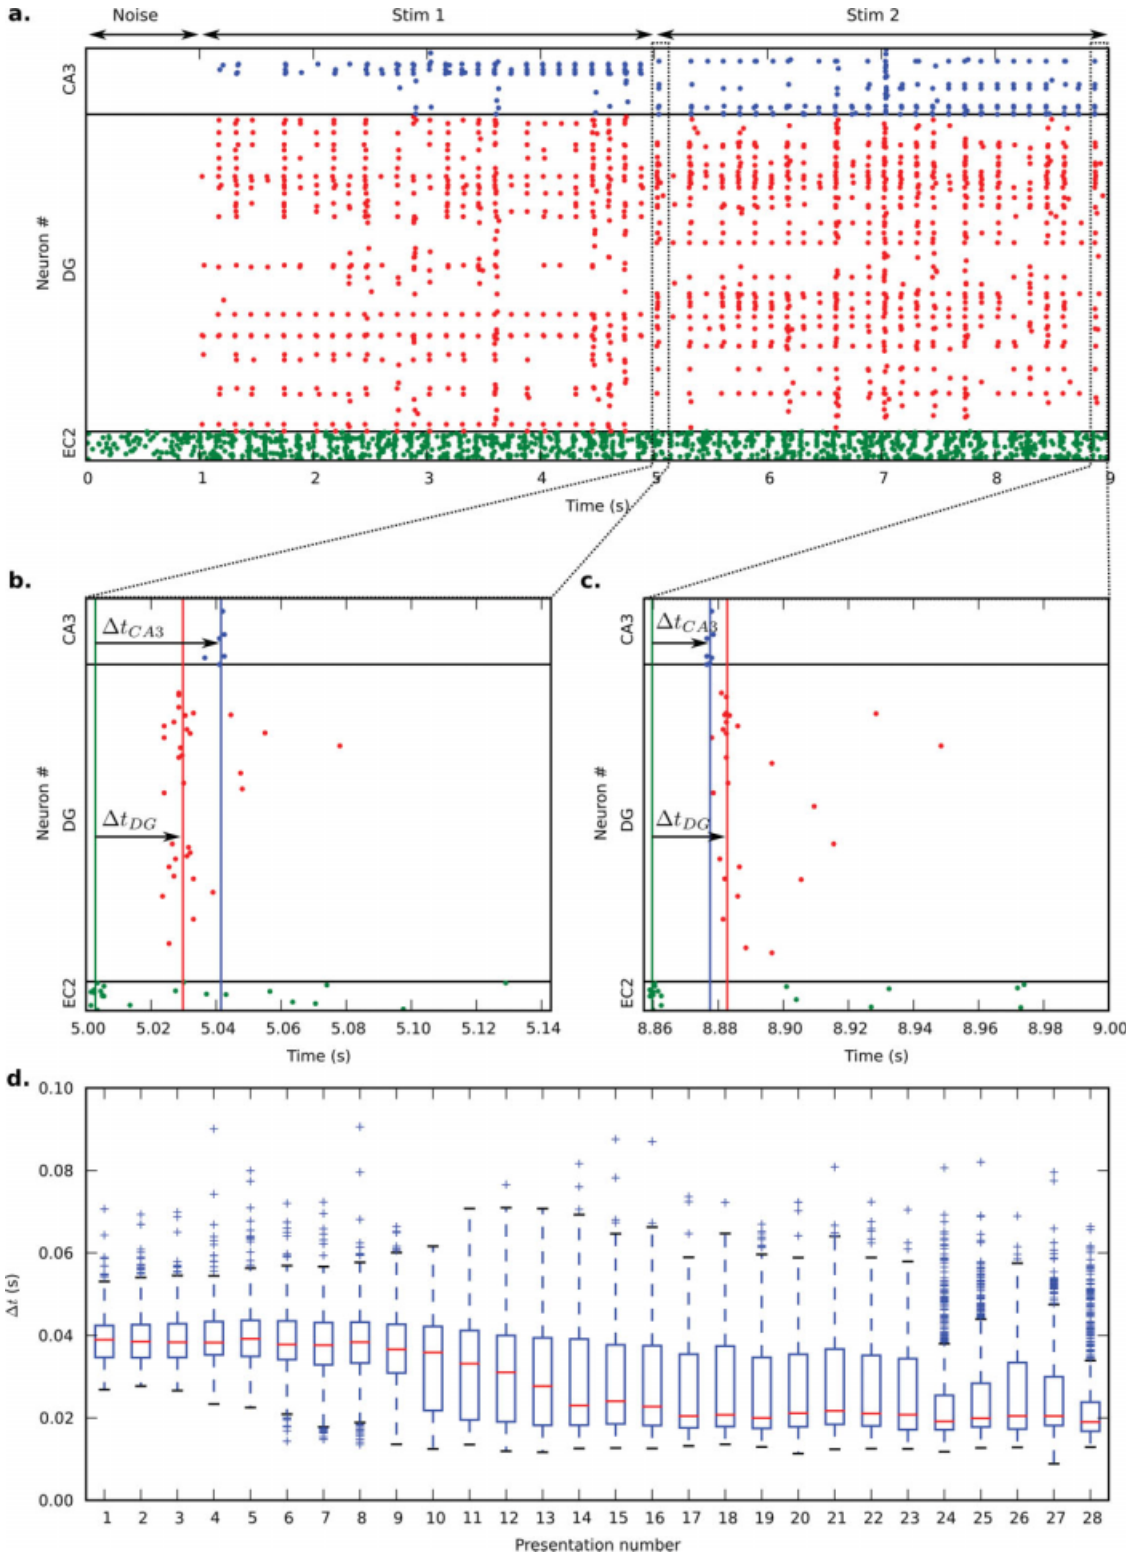

FIGURE 4. **CA3 response timing distinguishes learned from novel inputs: At the start of a presentation, the input is novel, and CA3 cells are slow to activate, not responding until activated by the indirect link through DG. After repeated presentations, CA3 cells start responding faster, activated directly by the EC2 input and eventually respond before the arrival of the indirect signal. (a) Operation of the network over 9 seconds. During the first second, noise was injected. Three distinct input patterns were chosen and stimulated 28 times sequentially (7 Hz for 4 s) in the EC cells with 5 ms jitter in spike timing. (b, c) Boxed areas from a, showing the first and last presentation of the third pattern. While the CA3 cells initially respond only to the indirect DG pathway, after multiple presentations the CA3 cells respond directly to the EC2 input prior to the DG response. The left-most vertical line shows the mean time of the jittered input stimulus, the two right-most vertical lines show the median of the CA3 and DG responses within the theta cycle. (d) Over the course of 28 presentations, the median response reduces from 39 ms to 19 ms, indicating that CA3 cells initially activated by the EC2-DG-CA3 pathway so as to respond to the direct EC2 input. [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]**

one stimulus and the next, and set to zero otherwise. The similarity between learned and test responses was determined by correlating the respective response vectors (calculated as the normalized dot product).

For each level of EC2 similarity, the similarity of the CA3 responses was averaged (Fig. 5). The response curve shows that the EC2-DG-CA3 network is acting as a strong pattern separator, and only performs limited pattern completion. For input patterns with less than 50% input similarity to the trained patterns, the CA3 response had a low similarity to the learned response (<5%). However, once the input pattern was sufficiently similar to a training pattern the CA3 response matched the learned response. The improvement of separation and completion when filtering the results according to CA3 spike latency shows that ignoring inputs after learned responses prevents the highly separated input (which also amplifies noise) from interfering with already known memories. Such improvement also shows that the self-inhibition of CA3 is not fulfilling this role perfectly.

This result demonstrates a core property of a memory system: the CA3 region accurately reactivates its previous response when a familiar pattern is presented, and activates a new and distinct pattern in response to a novel input. However, in many instances, a memory system must perform pattern completion with only partially complete input information. In the following section, we show how the network can be tuned to suit the situation by varying its similarity response curve between pattern separation and pattern completion.

Modulating the Threshold

The STSC model thus far has required the test input to exceed a certain degree of similarity to the learned input to complete the learned output representation. In existing models, this similarity threshold was fixed and as low as 30% (Marr, 1971; Treves and Rolls, 1992). To investigate the effect of varying the threshold between completion and separation, we introduced a modulatory input to the perforant path synapses. This input directly modifies the efficacy of the synapse, such that I_s is calculated as in Eq. (4) but with the postsynaptic effect of a presynaptic spike also multiplied by the current modulatory value:

$$\dot{I}_s = -I_s/\tau_s + \delta(t - \mathbf{t}_s)w\mu \qquad (9)$$

where τ_s is the decay time constant of the synapse, δ is the Dirac delta function, $\mathbf{t_s}$ is the vector of presynaptic spike times, t is the current time, and μ is the current modulation value.

At baseline efficacy ($\mu = 1.0$), the synapses act as in the previous simulations [as Eqs. (4) and (9) are equal]. Increasing μ increases the postsynaptic potential evoked by presynaptic spikes, thereby decreasing the total number of active synapses required to evoke an action potential in CA3. Consequently the specificity required to activate any particular neuron also decreases. By varying μ, we sought to test whether the network could correctly complete test input patterns with lower similarity to learned input patterns.

The stimulation and recording procedure closely followed that of the previous experiment. The network was trained with a set of three stimuli presented 28 times and tested with 200 distinct sets of input patterns and related variations for each of the stimuli. The testing phase was executed five times, starting at baseline modulation but increasing μ on each repetition such that one fewer fully potentiated synapse would be required to cause a postsynaptic action potential. The specific values of μ were determined by finding the threshold of single neurons with input similar to the EC2 input in the STSC model (for the exact μ values see Supporting Information). Using the normalized dot product to compare network responses, input/output overlap relations were generated for each test stimulus against the network response to the appropriate trained stimulus. Results were grouped by modulatory input level.

The network response at baseline modulation was no different from that in the previous tests of distinctiveness, that is, the network was highly selective (Fig. 6). As the modulatory value was increased, the level at which the output correlation exceeded the input correlation decreased, that is, the network completed more patterns.

DISCUSSION

Our goal in this research was to explore a mechanism to determine when to learn and when to recall in the hippocampus. We have shown that spike timing can act as an indicator of novelty to distinguish the two functions in the CA3 network (exemplified in the three-neuron model and demonstrated in the STSC model) and that modulation of synapses can alter the boundary between pattern separation (needed for learning) and pattern completion (needed for recall) (STSC model). In the following sections, we consider the significance of these timing and modulation results in the context of functional memory architectures, and describe the predictions the model makes regarding the biological systems on which it is based.

Novelty is Encoded by Spike Timing to Distinguish Learning and Recall

Our STSC model shows that the elapsed time from the stimulation of EC2 to a response in CA3 can act as an indicator of the novelty of the stimulation pattern. The system has functionally different behaviors for familiar and novel stimuli: EC2 patterns that are highly similar to a familiar pattern will directly evoke patterns in CA3 that reflect the high input similarity; EC2 patterns that have low or moderate similarity are insufficient to evoke a direct response in CA3, and instead activate CA3 through the indirect EC2-DG-CA3 pathway. The DG input to CA3 itself is active for both novel and familiar patterns; however, a sufficiently familiar EC2 pattern activates CA3 via the direct path before the corresponding DG input arrives, separating learned and novel responses temporally and allowing CA3 to suppress the DG input (in the STSC model, this is through $GABA_B$ inhibition of CA3). Unlike previous models that modulate between learning and recall over numerous pattern presentations with a multisecond timescale

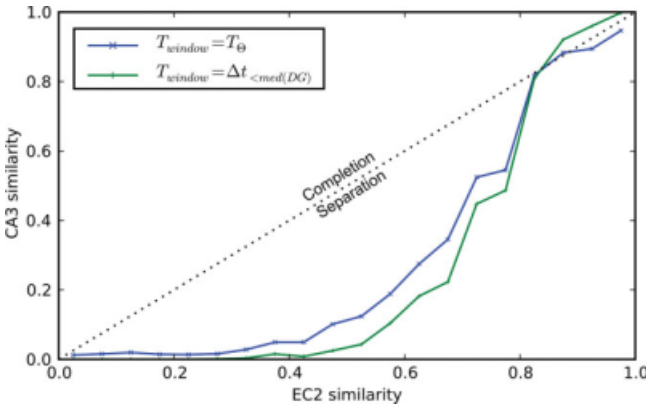

FIGURE 5. Learning and recall of patterns with and without consideration of timing. The dotted line ($y = x$) indicates the boundary between separation and completion: the farther below the line, the greater the separation; the farther above the line, the greater the completion. EC2 input vs. CA3 similarity comparison in the simulated hippocampal network shows that for low EC2 similarity, CA3 similarity approaches zero; only for the highest levels of EC2 similarity does CA3 similarity significantly rise above zero. The cross markers show the CA3 response overlap which includes all CA3 cell responses within a single theta cycle, with CA3 response to the indirect input only attenuated by the action of the inhibitory interneuron. Both separation and completion are improved if the timing of the response is taken into account: the plus markers show the average CA3 similarity responses which were active earlier than the median DG response within each theta cycle. Short latency responses provide reliable completion for known inputs before indirect stimulation, such that subsequent attenuation of the indirect input can prevent spurious separation noise. [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

(Hasselmo et al., 1995; Kali and Dayan, 2000), each presentation of a pattern in the STSC model is learned or recalled depending on its individual novelty, a process taking a fraction of each theta cycle.

Once a pattern is known, the network reliably and directly generates a response in CA3 before indirect activation via DG. The transition from a novel to a known input is facilitated through STDP at the CA3 synapses; if a CA3 neuron fires shortly after a presynaptic EC2 neuron, the synapse connecting them is potentiated. At the tested rate of one presentation per 7-Hz theta cycle, patterns are learnt by the network after 10–20 presentations of a stimulus (Fig. 4), corresponding to 1.5–3 s of real time. This time frame (1–3 s) is critical for many psychological phenomena including learning, memory, and attention.

The 7-Hz stimulation frequency ensures that there is no significant learning due to the response of the CA3 cells to one input and any subsequent stimulations of EC2 cells (due to the STDP time window and the maximum relative time of CA3 responses to EC2 stimuli, as shown in the ABC-AC case in Fig. 2b). As the STSC model does not incorporate the recurrent connections of CA3 and hence requires no particular minimum stimulation frequency to facilitate CA3 to CA3 synaptic plasticity, lower stimulation frequencies would be unlikely to modify the model's behavior. Although we did not explore simulations with greater stimulation frequencies, rates beyond ~12 Hz would cause significant interference in CA3 responses by subsequent inputs.

The STSC model does not incorporate plasticity in the synapses of the indirect path—neither the associative LTP from EC2 to DG nor the nonassociative LTP from DG to CA3.

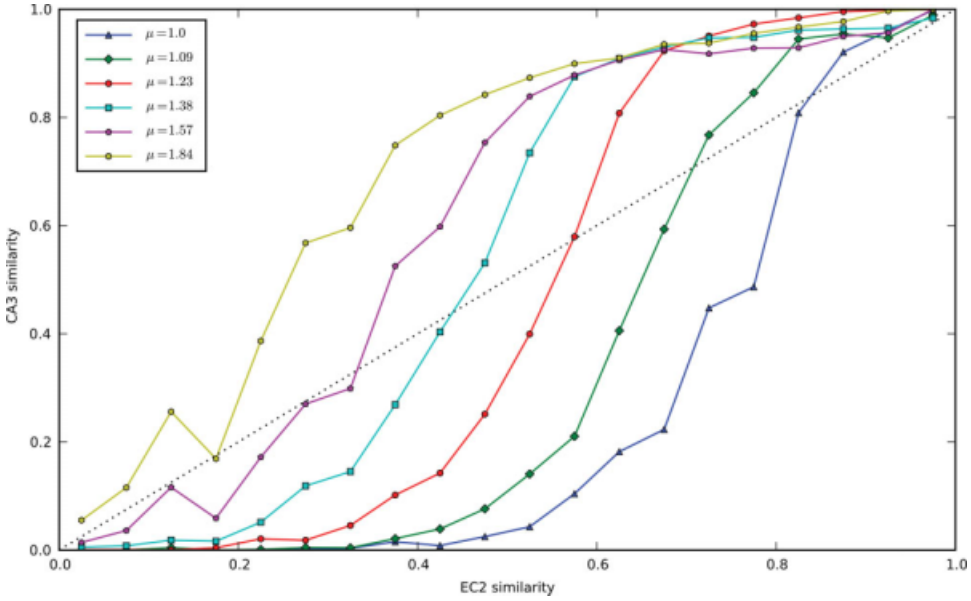

FIGURE 6. Modulation of the EC2-CA3 synapses affects the balance between pattern separation and completion. At baseline modulation (μ = 1.0), separation dominates over completion, resulting in a network that is highly selective. As the modulatory input is increased (μ = 1.09), the balance shifts from separation towards completion. At intermediate values (μ = 1.23, 1.38), the functions are approximately equally balanced, with areas equal above and below the separation–completion boundary. At the highest level of modulation (μ = 1.84), the network is strongly biased toward recall. [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

Although the former in particular may result in a latency reduction along the indirect route, similar to the operation of STDP on the direct route, the delay introduced by the extra cell (the DG cell) in the indirect path makes it unlikely that a known input will reach a CA3 cell via the indirect path before the direct input causes a CA3 response. Importantly, the prerequisite for the correct isolation of novel and known responses is the relative timing of the first responses of CA3 cells and the synaptic input from DG; so long as the first CA3 responses precede the DG input, the CA3 cells have some capability to "ignore" the DG input. In the simple ABC-AC network, once cell C learns the input from A, hyperpolarization of the cell after the direct response ensures that the input from B is ignored, resulting in a response only to the learned pathway—recall. In the STSC model, a response is known if a small subset of CA3 neurons respond directly and cause an inhibitory feedback neuron to attenuate further incoming signals for a short time, also resulting in recall by using only the learned pathway. Although the dynamics of the fast acting global inhibition do achieve this suppression in the current model, it seems likely that attenuation of the indirect pathway would occur at the mossy fiber synapses rather than solely via $GABA_B$ inhibition of the CA3 cells (as modeled here), so as to allow other inputs to CA3 to drive the network's behavior (e.g., the CA3 recurrent connections) until the following EC2 input.

On the basis of the timing produced by the STSC model, we would predict that in vivo recordings from DG and CA3 in response to stimulation of EC2 would show a population of spikes in CA3 preceding DG activity (representing recall of cues) and a population of spikes in CA3 following DG activity (representing cues yet to be learnt). A study by Yeckel and Berger (1990) recorded the responses of DG, CA3, and CA1 to the stimulation of axons of EC in an anesthetized rabbit. Responses in CA3 were shown to occur at times both before and after responses in DG, with the majority of responses in CA3 occurring with low latency before a response in DG. Yeckel and Berger commented that the responses in CA3 with longer latency than the latency in DG "did not follow more than the first few trains of continuous stimulation" of the perforant path, although this effect was not quantified in the paper. The results and the observation are consistent with predictions made by our model, although the exact latencies to the two CA3 responses differ. Based on the assumption that artificial stimulation of EC is an unfamiliar pattern for the system, the STSC model predicts that the initial stimulation will produce a long latency CA3 response through DG. After a few trains of stimulation, the learned connections from EC2 would result in a lower latency response in CA3. As the magnitude of the earlier direct response increases, the magnitude of the later indirect response should decrease—an effect seen clearly in the STSC model (Fig. 4d).

Our method of combining learning and recall is also consistent with recent data demonstrating that the phase in theta of pyramidal activity in CA1 increases with environmental novelty (Lever et al., 2010). Lever et al. interpret their data as support for the hypothesis that half phases of theta are alternately used for encoding and retrieval. However, the data is equally consistent with an interpretation that the initial longer latency responses correspond to the untrained stimuli through the trisynaptic pathway (EC2-DG-CA3-CA1), and the shorter latencies result from the learned stimuli through the monosynaptic pathway (EC3-CA1). The 33° theta phase shift observed by Lever et al. would correspond to approximately half a gamma cycle, which could be due to the different delays in the pathways. Assuming a theta rhythm of ~6 Hz, the Lever et al. data is consistent with that of Yeckel and Berger (1990). In contrast to the theta phase account of separation and completion given by Hasselmo et al. (2002), our STSC model makes a strong prediction about how learning would affect the phase shift in CA regions relative to EC: namely the STSC model predicts that, relative to the phase shift between the EC inputs, the magnitude of the phase shift due to learning would be lower in CA3 than in CA1.

Modulating Uncertainty: Roles for Acetylcholine and Dopamine

Our DG model acts as a delayed pathway for the purpose of novelty detection, but it also performs an orthogonalization role similar to previous models (Treves and Rolls, 1992). Differences between EC2 stimulation patterns are enhanced by the indirect pathway. Small differences between EC2 stimuli lead to larger differences between baseline CA3 responses before learning. As a consequence, the CA3 patterns that are evoked by the indirect pathway are inherently distinct (as shown in Fig. 5).

Because small differences in the EC2 layer are further separated in DG, input noise is also amplified. Computational models to date have worked around this issue via different means, for example, by using explicit learning and recall modes (Kali and Dayan, 2000), by designating separate subphases of the theta rhythm for learning and recall (Hasselmo et al., 2002), or by permitting remapping of known inputs in partially known environments (Hasselmo et al., 1995). Our STSC model selectively suppresses the effect of highly separated input from DG for known EC2 patterns on a pattern-by-pattern basis, intrinsically ensuring that any patterns that are sufficiently well known will not be relearned. This intrinsic novelty detection operates without regard to the level of input pattern to learned pattern correlation, or to the corresponding correlation of evoked responses in DG, thus modifying the correlation level required to elicit a learned response will not negate the separation–suppression effect of the novelty detection process.

We illustrated a mechanism for tuning the rate of perceived novelty of a pattern, and hence the ratio of learning over recall by modulating the synaptic strength between EC2 and CA3. The results show that increased synaptic efficacy in CA3 will cause CA3 cells to fire more frequently before DG input reaches the region, increasing the rate of pattern completion, which in turn will prevent the encoding of a new pattern (Fig. 6). An equivalent modulatory effect is visible even at a single CA3 synapse, where modulation affects the weight at which that synapse will become strong enough to evoke a postsynaptic

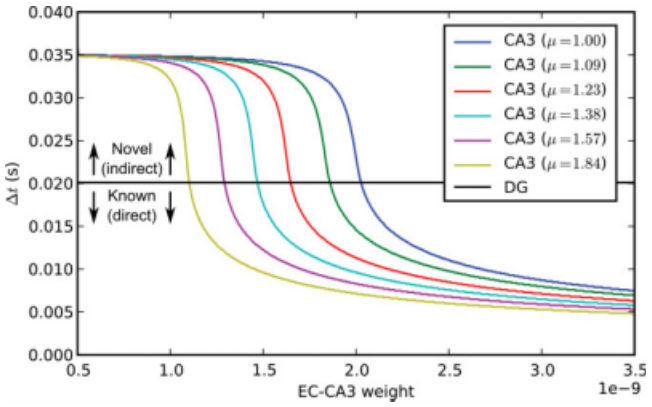

FIGURE 7. Modulation in the response latency of a single neuron. For low EC2-CA3 weights the delay is high. Initially the modulation has little effect; at a synaptic weight of 5.0×10^{-10} the CA3 response time is at its ceiling (~35 ms), a value determined by latency of the indirect path. For weights between $\sim 1.0 \times 10^{-9}$ and 2.1×10^{-9}, the CA3 response time depends strongly on the modulation level of the synapses, which determines the EC2-CA3 weight at which the CA3 cell responds to the direct stimulation before the indirect stimulation arrives. For weights beyond $\sim 2.1 \times 10^{-9}$, the modulatory value has little functional difference as the direct path evokes responses irrespective of modulation level, and the response latency decreases monotonically with the weight, ensuring that so long as μ is equal across all cells, the cells latencies are ordered according to their strength (and hence how close the input is to that cell's learned input). [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

spike (Fig. 7). As modulation is increased over baseline ($\mu > 1$), the mean EC2-CA3 synaptic weight required for a CA3 cell to respond to the EC2 input decreases. Completion occurs when a CA3 cell responds with a lower latency than the response of DG. Thus, the greater the modulation value, the lower the required similarity between the current input and the learned pattern. However, as the strongest learned weights still respond most quickly ($d\Delta t/dw$ is always negative), and the CA3 response is attenuated once a certain number of CA3 cells respond, then irrespective of the level of modulation, only the memory that best matches the input will be activated.

ACh has previously been proposed as a modulator of the novelty detection threshold—the threshold at which an input is similar enough to a previously learned input to be considered known (Hasselmo and Schnell, 1994; Hasselmo and Wyble, 1997). In support of this theory, Rogers and Kesner (2003) have shown that an infusion of scopolamine (a cholinergic antagonist) in CA3 prevents encoding of new patterns but does not impair retrieval during a rat navigation task. Evidence suggests that associational/commissural (A/C) connections in the stratum oriens of CA3 are attenuated in the presence of ACh (Vogt and Regehr, 2001), potentially ensuring that the pattern separation mechanism through the indirect EC2-DG-CA3 path is free from interference during encoding (Hasselmo et al., 1995). Furthermore, in high-ACh environments, LTP is elevated in the stratum oriens of CA1 (Huerta and Lisman, 1995; Ovsepian et al., 2004; Shinoe et al., 2005), although there appear to be no experiments to date testing for a similar effect in CA3.

Dopamine is similarly implicated in hippocampal function, leading to a proposal that it may likewise modulate the novelty detection threshold in the region (Lisman and Otmakhova, 2001). Like ACh in the stratum oriens, elevated concentrations of dopamine appear to selectively attenuate the effect of perforant path inputs in the stratum lacunosum-moleculare (LM) of CA1 (Otmakhova and Lisman, 1999). Dopaminergic neurons, including those in the region afferent to the hippocampus, are highly active for unpredicted reward (Schultz, 2002) and LTP is elevated at perforant path synapses to CA1 pyramidal cells with dopamine receptor activation (Otmakhova and Lisman, 1996). Although dopamine is generally implicated in reward learning, evidence suggests it is also involved in the attentional or motivational scenarios required of a signal for novelty detection (Lisman and Otmakhova, 2001). It thus appears that ACh and dopamine selectively attenuate and up-modulate learning in internal and external inputs, respectively.

We agree here with Hasselmo et al. (1995) that ACh release in CA3 may not directly signal learning, but rather that neuromodulation of CA3 via both ACh and dopamine may encode uncertainty by affecting the probability that any afferent stimulation will be perceived as novel and hence will be learned. Increased neuromodulatory release would imply greater uncertainty: decreasing the strength of the direct EC2-CA3 "recall" path, thereby increasing the similarity threshold required to reactivate previously learned representations and resulting in more patterns being classed as novel and therefore being learned. Decreased modulatory release would imply less uncertainty and have the inverse effect, resulting in fewer novel experiences and less learning.

A prediction from the STSC model is that neuromodulation of CA3 under novelty should, in addition to its effect on the A/C fibers, also have some effect on the perforant path inputs to CA3. Presynaptic stimulation of the perforant path should have a postsynaptic somatic effect that varies with the concentration of some neuromodulator. Given the evidence of dopaminergic attenuation of perforant path synapses in CA1, we suggest dopamine as the most likely candidate for this modulation in CA3. In our model, the efficacy of perforant path synapses increased with the modulatory parameter μ. We would then expect μ to vary inversely with dopamine concentration, such that as the concentration of dopamine rose, the efficacy of the perforant path would fall, and potentially the rate of STDP would rise [an effect omitted in the present work but one that could operate as described by Izhikevich (2007)]. Modulation could alternatively be achieved by direct cholinergic effect or by indirect means similar to the GABAergic innervation suggested by Vogt and Regehr (2001) at mossy fiber synapses. For example, the stratum oriens-LM (O-LM) interneurons are enhanced by activation of M1 and M3 muscarinic ACh receptors (Lawrence et al., 2006) and their axons terminate among the CA3 dendrites receiving perforant path stimulation (Gulyas et al., 1993).

Future Directions

Pattern separation and completion together provide a functional memory for individual input patterns. The resemblance of

the recurrent collateral network in CA3 to classic associative networks in artificial neural network literature has led many to propose that the CA3 network functions as an autoassociative attractor network that enables the completion of patterns (Treves and Rolls, 1992; Kali and Dayan, 2000). The characteristic recurrent collaterals of the CA3 network have been omitted in the current model; we have demonstrated that a heteroassociative form of separation and completion is possible in a network without recurrence. For example, our results show that with modulation level $\mu = 1.84$, an input pattern with 40% similarity to a previously learned input pattern leads to 80% similarity in the corresponding CA3 responses (Fig. 6). In part, this replicates existing studies (O'Reilly and McClelland, 1994; Fuhs and Touretzky, 2000); however, we have added a key capability: spurious separation noise, corrected in many recurrent autoassociative networks by their attractor dynamics, can be suppressed by a short latency response when the input is known.

One question left unanswered in this work is that of the network's capacity. To avoid pattern crosstalk during the experiments, a maximum of 10 patterns were learned by the network at any time. Capacity in autoassociative networks has been well characterized (Abu-Mostafa and St. Jacques, 1985). Interestingly, modeling results suggest that adding biological detail to heteroassociative networks, and in particular considering spike timing, increases network capacity (Lytton, 1998). We are currently investigating this topic further.

In addition to pattern completion and separation, the hippocampus has been implicated in episodic memory (Tulving and Markowitsch, 1998). An alternative hypothesis of the role of the CA3 recurrent collaterals is that they encode temporal relations between memories, forming a cognitive graph (Muller et al., 1996; Lisman and Otmakhova, 2001). Sensory snapshots or even cortical replay of hippocampal inputs could thus evoke temporally distant but linked patterns, and similarities could be extracted based not only on sensory input but also on temporal links. Although our STSC model uses two gamma cycles to learn or recall memories, the data provided by Yeckel and Berger (1990) shows that more realistic timing may in fact be a single gamma cycle. If there are ~7–10 gamma cycles per theta cycle (Bragin et al., 1995), the remaining gamma cycles could be used (after learning) for traversal of temporally linked patterns via the recurrent network. In this manner, we envisage the CA3 network facilitating learning and retrieval of patterns unbiased by sensory similarity and interconnected solely by temporal proximity.

REFERENCES

Abu-Mostafa Y, St. Jacques J. 1985. Information capacity of the Hopfield model. IEEE Trans Inf Theory 31:461–464.

Amaral DG. 1993. Emerging principles of intrinsic hippocampal organization. Curr Opin Neurobiol 3:225–229.

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

Bakker A, Kirwan CB, Miller M, Stark CEL. 2008. Pattern separation in the human hippocampal CA3 and dentate gyrus. Science 319:1640–1642.

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

Bragin A, Jando G, Nadasdy Z, Hetke J, Wise K, Buzsaki G. 1995. Gamma (40–100 Hz) oscillation in the hippocampus of the behaving rat. J Neurosci 15:47–60.

Chattarji S, Stanton PK, Sejnowski TJ. 1989. Commissural synapses, but not mossy fiber synapses, in hippocampal field CA3 exhibit associative long-term potentiation and depression. Brain Res 495:145–150.

Claiborne BJ, Amaral DG, Cowan WM. 1986. A light and electron microscopic analysis of the mossy fibers of the rat dentate gyrus. J Comp Neurol 246:435–458.

Cutsuridis V, Cobb S, Graham BP. Encoding and retrieval in a model of the hippocampal CA1 microcircuit. Hippocampus 10.1002/hipo.20661.

Dayan P, Abbott LF. 2001. Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems, Cambridge, MA: MIT Press.

Do VH, Martinez CO, Martinez JL, Derrick BE. 2002. Long-term potentiation in direct perforant path projections to the hippocampal CA3 region in vivo. J Neurophysiol 87:669–678.

Fuhs MC, Touretzky DS. 2000. Synaptic learning models of map separation in the hippocampus. Neurocomputing 32:379–384.

Gulyas AI, Miles R, Hajos N, Freund TF. 1993. Precision and variability in postsynaptic target selection of inhibitory cells in the hippocampal CA3 region. Eur J Neurosci 5:1729–1751.

Guyonneau R, VanRullen R, Thorpe SJ. 2005. Neurons tune to the earliest spikes through STDP. Neural Comput 17:859–879.

Hasselmo ME, Fehlau BP. 2001. Differences in time course of ACh and GABA modulation of excitatory synaptic potentials in slices of rat hippocampus. J Neurophysiol 86:1792–1802.

Hasselmo ME, 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 ME, Wyble BP. 1997. Free recall and recognition in a network model of the hippocampus: Simulating effects of scopolamine on human memory function. Behav Brain Res 89:1–34.

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

Hasselmo ME, Hay J, Ilyn M, Gorchetchnikov A. 2002. Neuromodulation, theta rhythm and rat spatial navigation. Neural Netw 15:689–707.

Henze DA, Wittner L, Buzsaki G. 2002. Single granule cells reliably discharge targets in the hippocampal CA3 network in vivo. Nat Neurosci 5:790–795.

Hopfield JJ. 1982. Neural networks and physical systems with emergent collective computational abilities. Proc Natl Acad Sci USA 79:2554–2558.

Huerta PT, Lisman JE. 1995. Bidirectional synaptic plasticity induced by a single burst during cholinergic theta oscillation in CA1 in vitro. Neuron 15:1053–1063.

Izhikevich EM. 2003. Simple model of spiking neurons. IEEE Trans Neural Netw 14:1569–1572.

Izhikevich EM. 2004. Which model to use for cortical spiking neurons? IEEE Trans Neural Netw 15:1063–1070.

Izhikevich EM. 2007. Solving the distal reward problem through linkage of STDP and dopamine signaling. Cereb Cortex 17:2443–2452.

Kali S, Dayan P. 2000. The involvement of recurrent connections in area CA3 in establishing the properties of place fields: A model. J Neurosci 20:7463–7477.

Kunec S, Hasselmo ME, Kopell N. 2005. Encoding and retrieval in the CA3 region of the hippocampus: A model of theta-phase separation. J Neurophysiol 94:70–82.

Lawrence JJ, Statland JM, Grinspan ZM, McBain CJ. 2006. Cell type-specific dependence of muscarinic signalling in mouse hippocampal stratum oriens interneurones. J Physiol 570:595–610.

Lee I, Kesner RP. 2004. Encoding versus retrieval of spatial memory: Double dissociation between the dentate gyrus and the perforant path inputs into CA3 in the dorsal hippocampus. Hippocampus 14:66–76.

Leutgeb S, Leutgeb JK, Treves A, Moser M-B, Moser EI. 2004. Distinct ensemble codes in hippocampal areas CA3 and CA1. Science 305:1295–1298.

Leutgeb JK, Leutgeb S, Moser M-B, Moser EI. 2007. Pattern separation in the dentate gyrus and CA3 of the hippocampus. Science 315:961–966.

Lever C, Burton S, Jeewajee A, Wills TJ, Cacucci F, Burgess N, O'Keefe J. 2010. Environmental novelty elicits a later theta phase of firing in CA1 but not subiculum. Hippocampus 20:229–234.

Levy WB, Steward O. 1983. Temporal contiguity requirements for long-term associative potentiation/depression in the hippocampus. Neuroscience 8:791–797.

Lisman JE, Otmakhova NA. 2001. Storage, recall, and novelty detection of sequences by the hippocampus: Elaborating on the SOCRATIC model to account for normal and aberrant effects of dopamine. Hippocampus 11:551–568.

Lytton WW. 1998. Adapting a feedforward heteroassociative network to Hodgkin-Huxley dynamics. J Comput Neurosci 5:353–364.

Markram H, Lubke J, Frotscher M, Sakmann B. 1997. Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. Science 275:213–215.

Marr D. 1971. Simple memory: A theory for archicortex. Philos Trans R Soc Lond B Biol Sci 262:23–81.

Martinez CO, Do VH, Martinez JL, Derrick BE. 2002. Associative long-term potentiation (LTP) among extrinsic afferents of the hippocampal CA3 region in vivo. Brain Res 940:86–94.

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.

Morris R. 2007. Theories of hippocampal function. In: Andersen P, Morris RGM, Amaral DG, Bliss TVP, O'Keefe J, editors. The Hippocampus Book. New York, NY: Oxford University Press. pp 581–713.

Muller RU, Stead M, Pach J. 1996. The hippocampus as a cognitive graph. J Gen Physiol 107:663–694.

O'Reilly RC, McClelland JL. 1994. Hippocampal conjunctive encoding, storage, and recall: Avoiding a trade-off. Hippocampus 4:661–682.

Otmakhova NA, Lisman JE. 1996. D1/D5 dopamine receptor activation increases the magnitude of early long-term potentiation at CA1 hippocampal synapses. J Neurosci 16:7478–7486.

Otmakhova NA, Lisman JE. 1999. Dopamine selectively inhibits the direct cortical pathway to the CA1 hippocampal region. J Neurosci 19:1437–1445.

Ovsepian SV, Roger A, Michael JR. 2004. Endogenous acetylcholine lowers the threshold for long-term potentiation induction in the CA1 area through muscarinic receptor activation: In vivo study. Eur J Neurosci 20:1267–1275.

Rogers JL, Kesner RP. 2003. Cholinergic modulation of the hippocampus during encoding and retrieval. Neurobiol Learn Mem 80:332–342.

Rolls ET. 1996. A theory of hippocampal function in memory. Hippocampus 6:601–620.

Schultz W. 2002. Getting formal with dopamine and reward. Neuron 36:241–263.

Scoville WB, Milner B. 1957. Loss of recent memory after bilateral hippocampal lesions. J Neurol Neurosurg Psychiatry 20:11–21.

Shinoe T, Matsui M, Taketo MM, Manabe T. 2005. Modulation of synaptic plasticity by physiological activation of M1 muscarinic acetylcholine receptors in the mouse hippocampus. J Neurosci 25:11194–11200.

Song S, Miller KD, Abbott LF. 2000. Competitive Hebbian learning through spike-timing-dependent synaptic plasticity. Nat Neurosci 3:919–926.

Spruston N, McBain C. 2007. Structural and functional properties of hippocampal neurons. In: Andersen P, Morris RGM, Amaral DG, Bliss TVP, O'Keefe J, editors. The Hippocampus Book. New York, NY: Oxford University Press. pp 133–242.

Treves A, Rolls ET. 1992. Computational constraints suggest the need for two distinct input systems to the hippocampal CA3 network. Hippocampus 2:189–199.

Tulving E, Markowitsch HJ. 1998. Episodic and declarative memory: Role of the hippocampus. Hippocampus 8:198–204.

Vazdarjanova A, Guzowski JF. 2004. Differences in hippocampal neuronal population responses to modifications of an environmental context: Evidence for distinct, yet complementary, functions of CA3 and CA1 ensembles. J Neurosci 24:6489–6496.

Vogt KE, Regehr WG. 2001. Cholinergic modulation of excitatory synaptic transmission in the CA3 area of the hippocampus. J Neurosci 21:75–83.

Yeckel MF, Berger TW. 1990. Feedforward excitation of the hippocampus by afferents from the entorhinal cortex: Redefinition of the role of the trisynaptic pathway. Proc Natl Acad Sci USA 87:5832–5836.