Available online at www.sciencedirect.com
SCIENCE DIRECT
Neural Networks 18 (2005) 1265–1279
Neural Networks
www.elsevier.com/locate/neunet

2005 Special issue

Cortico-hippocampal interaction and adaptive stimulus representation: A neurocomputational theory of associative learning and memory

Mark A. Gluck*, Catherine Myers, Martijn Meeter

Center for Molecular and Behavioral Neuroscience, Rutgers University-Newark, Newark, NJ 07102, USA

Abstract

Computational models of the hippocampal region link psychological theories of associative learning with their underlying physiological and anatomical substrates. Our approach to theory development began with a broad description of the computations that depend on the hippocampal region in classical conditioning (Gluck & Myers, 1993, 2001). In this initial model, the hippocampal region was treated as an Information-processing system that transformed stimulus representations, compressing (making more similar) representations of inputs that co-occur or are otherwise redundant, while differentiating (or making less similar) representations of inputs that predict different future events. This model led to novel predictions for the behavioral consequences of hippocampal-region lesions in rodents and of brain damage in humans who have amnesia or are in the earliest stages of Alzheimer's disease. Many of these predictions have, since been confirmed by our lab and others. Functional brain imaging studies have provided further supporting evidence. In more recent computational modeling, we have shown how some aspects of this proposed information-processing function could emerge from known anatomical and physiological characteristics of the hippocampal region, including the entorhinal cortex and the septo-hippocampal cholinergic system. The modeling to date lays the groundwork for future directions that increase the depth of detail of the biological modeling, as well as the breadth of behavioral phenomena addressed. In particular, we are working now to reconcile these kinds of incremental associative learning models with other models of the hippocampal region that account for the rapid formation of declarative memories.

The hippocampus and associated structures, often termed the hippocampal region, have long been thought to play a critical role in memory, although there has traditionally been very little consensus on what precisely that role is. In humans, the hippocampal region lies within the medial temporal lobe, and humans with bilateral medial temporal damage exhibit anterograde amnesia, a profound deficit for learning new 'declarative' memories for facts and auto-biographical events. These amnesic patients can, however, learn new 'non-declarative' memories such as simple habits, skills, and conditioned reflexes (Squire, 1987). In the past, researchers have often adopted the simple rule of thumb that declarative learning depends on the hippocampus but non-declarative learning does not.

This simple dichotomy breaks down quickly, however, when one looks at more complicated types of non-declarative learning. For example, a canonical form of non-declarative learning is classical eyeblink conditioning, in which an animal or human receives an airpuff to the eye that evokes a protective eyeblink response. If this airpuff (called the unconditioned stimulus or US) is repeatedly preceded by a cue such as a tone (called the conditioned stimulus or CS), the animal learns to respond to the cue by giving an anticipatory eyeblink (called the conditioned response or CR), so that the eyelid is closed at the time of expected airpuff arrival (Gormezano, Kehoe, & Marshall, 1983). A basic form is delay eyeblink conditioning, in which the CS and US overlap and co-terminate. Delay eyeblink conditioning is spared in animals with hippocampal-region damage (Fig. 1(A), left; Schmaltz & Theios, 1972) and in humans with medial temporal amnesia (Fig. 1(A), right; Gabrieli, McGlinchey-Berroth, Carrillo, Gluck, Cermack and Disterhoft, 1995). However, in slightly more complicated forms of eyeblink conditioning, the hippocampus does play a role. For example, latent inhibition is a phenomenon whereby prior exposure to the CS alone slows subsequent acquisition of a CS–US association (Lubow, 1973); hippocampal-region damage abolishes the effect in rabbits, so that exposed and non-exposed animals

* Corresponding author.
E-mail address: gluck@pavlov.rutgers.edu (M.A. Gluck).

0893-6080/$ - see front matter © 2005 Published by Elsevier Ltd.
doi:10.1016/j.neunet.2005.08.003

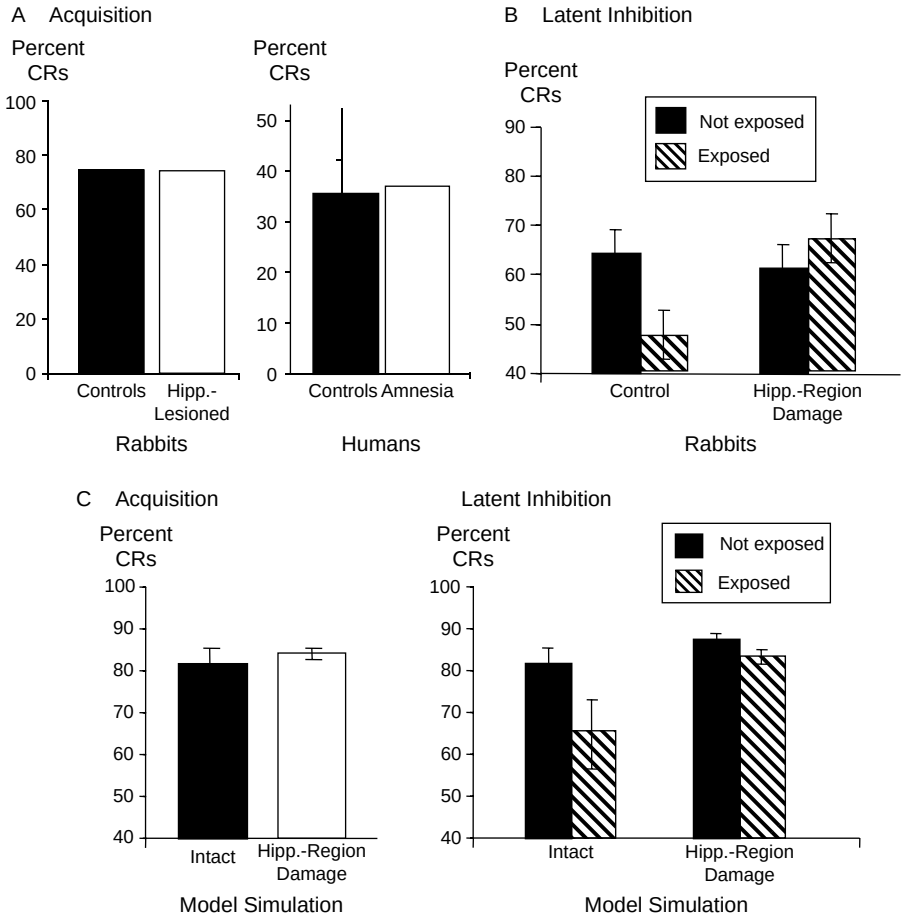

Fig. 1. Hippocampal-region damage spares simple eyeblink conditioning in (A) animals (after Schmaltz & Theios, 1972) and humans (after Gabrieli et al., 1995). However, hippocampal-region damage in rabbits (B) abolishes latent inhibition in the eyeblink conditioning paradigm (after Solomon & Moore, 1975; see also Shohamy et al., 2000). The Gluck and Myers (1993) cortico-hippocampal model correctly accounts for data showing that (C) hippocampal lesion produces no impairment in 'simple' delay conditioning but (D) does abolish latent inhibition.

learn at the same speed (Fig. 1(B); Solomon & Moore, 1975).

In an effort to link psychological theories of learning with underlying neural substrates, we began with a broad description of the information-processing computations that appear to depend on the hippocampal region in classical conditioning (Gluck & Myers, 1993, 2001). As reviewed below, this model led to several novel predictions for the behavioral consequences of hippocampal-region lesions in rodents and of brain damage in humans who have amnesia or are in the earliest stages of Alzheimer's disease. Many of these predictions have, since been confirmed by our lab and others while brain imaging studies have provided further supporting evidence.

The modeling to date lays the groundwork for future directions that increase the depth of detail of the biological modeling, as well as the breadth of behavioral phenomena address. Several of these future directions are described at the end of this paper, along with a summary of preliminary progress. In particular, we are working to reconcile models of incremental associative learning (including classical and instrumental conditioning) with models of the hippocampal region that account for the rapid formation of declarative memories.

1. 'Top-down' model of cortico-hippocampal function in associative learning

What information-processing role does the hippocampus play in classical conditioning? To address this, we developed a computational model of cortico-hippocampal interaction based on connectionist theories of learning and representation (Gluck & Myers, 1993, 1995, 1996, 2001). The model conceptualizes the brain as a series of interacting modules, each implementing the information-processing functions subserved by a particular brain region.

For eyeblink conditioning, the cerebellum is the necessary and sufficient substrate for the storage and expression of learned CS–US associations (Thompson, 1986). CS information arrives as input to the cerebellum, as does information about the US (airpuff); cerebellar output drives the behavioral CR (eyeblink). Learning occurs via plasticity at neuronal synapses in cerebellum that associate

the CS with the US. This biological learning can be formalized as the least-mean square (LMS) algorithm (Widrow & Hoff, 1960), which has relations both to biological plasticity mechanisms such as long-term potentiation and depression (e.g. Levy, Brassel, & Moore, 1983) and to mathematical models of conditioning (Rescorla & Wagner, 1972).

We formalized the cerebellar substrates of eyeblink conditioning in a connectionist network, shown in Fig. 2(A) (Gluck, Myers, & Thompson, 1994; Thompson, 1986). This network learns to map from inputs specifying the presence of CSs and contextual cues, to a pattern of activation in an internal layer of nodes, via a layer of weighted connections. This activation pattern constitutes a re-mapping or re-representation of the input, which is then mapped to output driving the behavioral CR via a second layer of weighted connections. On each trial, the system 'error' is the difference between the actual response (CR) and the desired response (US). The LMS rule can be used to modify the weights between the internal-layer and output-layer nodes, proportional to this error. However, no such error measure is defined for the internal-layer nodes, and so LMS cannot be used to modify the weights on the inputs to the internal-layer. As a result, no learning takes place in the weights between the input-layer and internal-layer nodes. That is to say, the 'representation' is fixed. Nevertheless, for many simple problems (such as delay eyeblink conditioning) this system with only one layer of modifiable weights can generate outputs that capture the behavior of an animal learning a conditioned response.

What does the hippocampal-region add to this? We proposed that the hippocampal-region develops new representations that encode important stimulus–stimulus regularities in the environment (Gluck & Myers, 1993). In particular, if two CSs reliably co-occur or are otherwise redundant, their representations become compressed, or more similar. Conversely, the representations for two CSs that predict different USs become differentiated, or highly dissimilar. Gluck and Myers (1993) suggested that the hippocampal-region performs this kind of redundancy compression and predictive differentiation.

We implemented this theory in a connectionist network model shown in Fig. 2(B) (Gluck & Myers, 1993, 2001). Hippocampal-region processing is implemented via a predictive autoencoder (Baldi & Hornik, 1989; Hinton, 1989), which learns to map CS inputs, through an internal node layer, to outputs that reconstruct those inputs and also predict the US. This network, unlike the cerebellar network, is able to modify both layers of weighted connections through a learning algorithm such as error backpropagation (Rumelhart, Hinton, & William, 1986). In the process, internal-layer nodes form a representation of the input that tends to compress redundant information while preserving and differentiating information that predicts the US, just as required by our hypothesis.

This hippocampal-region network then provides these new representations to the cerebellar network. A random recoding of the hippocampal-region network's internal-layer activations becomes the 'desired output' for the internal layer of the cerebellar network, and the error is the difference between this and the output of the internal layer of the cerebellar model. The cerebellar network then uses the LMS error-correcting rule to adapt the input-to-internal layer weights, just as it uses LMS to adapt internal-to-output layer weights. Over time, representations develop in the internal-layer nodes of the cerebellar network that are linear recombinations of representations developed by the hippocampal region network.

Within this model framework, broad hippocampal-region damage is simulated by disabling the hippocampal-

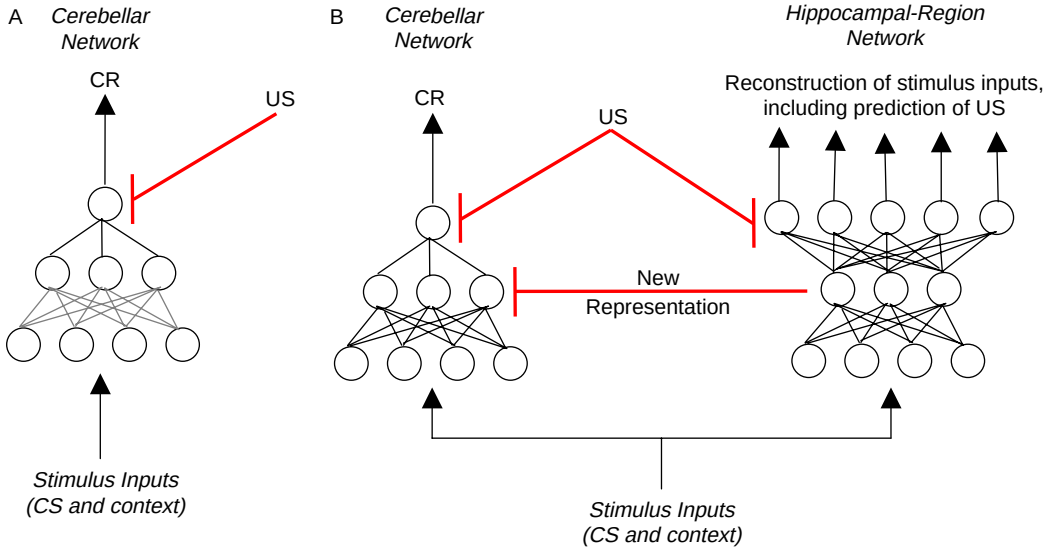

Fig. 2. (A) A connectionist model of the cerebellar substrates of motor-reflex conditioning. (B) A connectionist model of hippocampal region function in conditioning.

region network (Fig. 2(A)). In this lesioned model, no new hippocampal-dependent representations are formed, and the training signal to the cerebellar network internal layer is silenced. As a result, the cerebellar network cannot adopt any new representations, although it can still learn to map from its existing representations to new behavioral responses by modifying its upper layer of weights.

This computational model of hippocampal-region function correctly accounts for data showing that hippocampal-region damage does not impair simple delay conditioning (Fig. 1(C); compare Fig. 1(A)) but does impair more complex behaviors such as latent inhibition (Fig. 1(D); compare Fig. 1(B)). Similarly, the model accounts for many other trial-level conditioning phenomena in both intact and hippocampal-lesioned animals (Gluck & Myers, 1993; Myers & Gluck, 1994; Myers, Gluck, & Granger, 1995), and has even been applied to other domains such as human category learning (Gluck, Oliver & Myers, 1996) and rodent odor discrimination (Myers & Gluck, 1996). Several novel predictions of the computational model have since been tested and confirmed in animal studies of hippocampal-region lesion (e.g. Allen, Chelius, & Gluck, 2002a; Allen, Padilla, & Gluck, 2002b; Shohamy, Allen, & Gluck, 2000). These and other results confirm the basic predictions of our model, and suggest that that the hippocampus and related structures are intimately involved in even 'simple' procedural learning tasks such as classical conditioning.

2. Applications to memory disorders in humans

Learned irrelevance is a conditioning phenomenon, in which prior uncorrelated exposure to the CS and US slows subsequent learning of the CS–US association (Mackintosh, 1973). It is similar to the more well-known latent inhibition paradigm (Lubow, 1973), in which an animal is pre-exposed to CS alone trials (rather than uncorrelated CS and US trials as in learned irrelevance), in that both produce slower learning of a subsequence CS–US association. One novel prediction of our cortico-hippocampal model of Fig. 2(B) was that latent inhibition and learned irrelevance should depend on the hippocampal-region. This prediction was later confirmed by our laboratory in a study of eyeblink conditioning in rabbits: hippocampal-region lesions abolished the learned irrelevance effect of uncorrelated prior exposure to the CS and US (Fig. 3(A); Allen et al., 2002a,b). One immediate question was whether this applies to humans as well. To test this, we developed a paradigm that embeds the logical structure of learned irrelevance within a computer-based learning task. In this task, people had to learn that a color change (conceptually similar to a CS) predicted a salient screen event (similar to a US). Learning this association was markedly slowed in controls who were previously exposed to the CS and US uncorrelated with each other (Fig. 3(B)). This learned irrelevance effect was abolished in amnesic individuals with bilateral medial temporal damage (Myers et al., 2000). In other words, learned irrelevance appears to depend on the hippocampal system in humans just as in rabbits.

Another implication of our model for human learning comes from two-phase studies in which initial learning is followed by a transfer test (Myers et al., 2002a,b; Myers, Shohamy et al., in press; see also Eichenbaum, Mathews, & Cohen, 1989). One example is a concurrent discrimination task, in which subjects see pairs of colored shapes and must choose the correct member of each pair (Fig. 4(A)). This simple form of learning does not depend on hippocampal mediation in the computational model (Myers & Gluck,

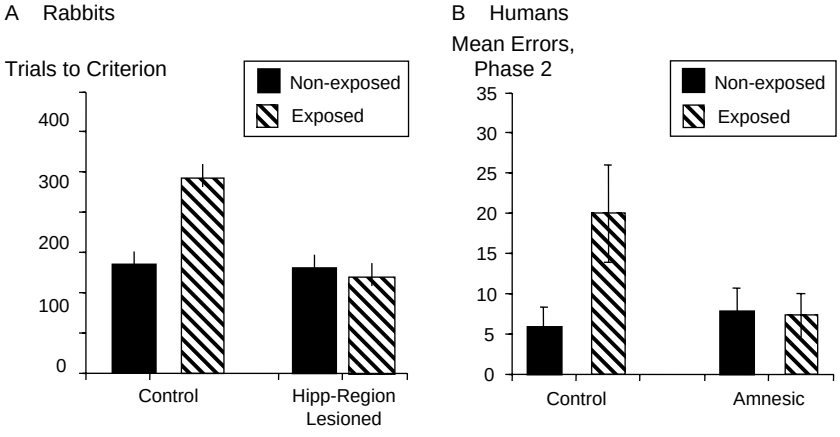

Fig. 3. (A) Rabbits in an eyeblink conditioning task were trained on the learned irrelevance paradigm. In phase 1, the exposed group of rabbits saw the CS and US presented uncorrelated while the non-exposed group did not. In phase 2, both groups were trained that the CS predicts the US using a standard delay conditioning paradigm. In control rabbits, the exposed group took significantly longer to learn the association in phase 2. In contrast, rats with hippocampal-region lesions (specifically to the entorhinal region), did not show this effect. Data adapted from Allen, Chelius, and Gluck (2002). (B) Data from a computer-based 'learned irrelevance' task. In healthy controls, the Exposed Group is significantly slower than the Non-Exposed Group to learn the cue-outcome association in phase 2. However, individuals with medial temporal amnesia learn the association equally quickly regardless of prior exposure condition. Not only is learned irrelevance abolished, but the amnesics in the Exposed Group actually learn faster than the controls in the Exposed Group! (Fig. adapted from Myers, McGlinchey-Berroth et al., 2000).

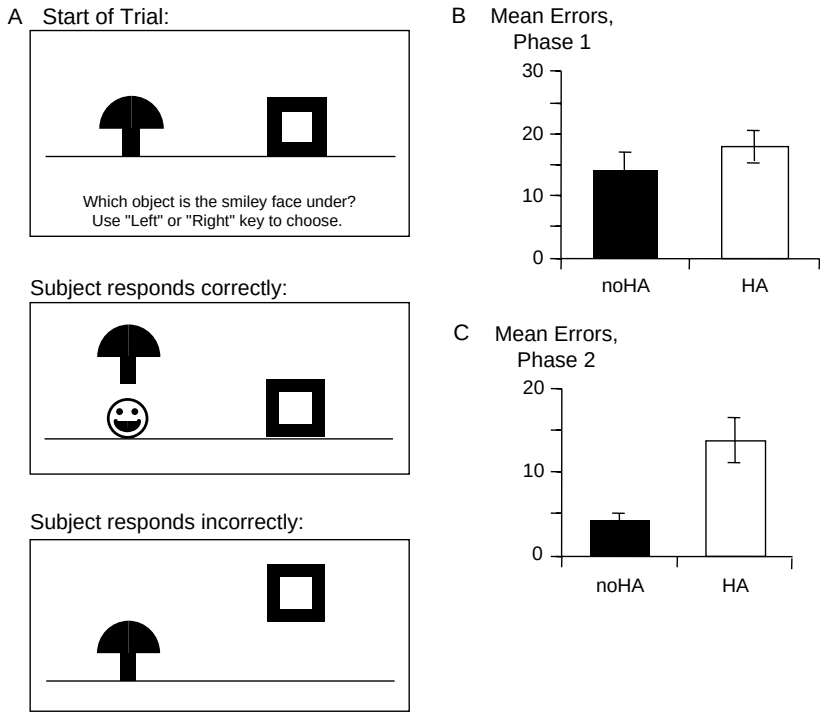

Fig. 4. (A) A computer-based concurrent discrimination task. On each trial, subjects see two colored objects (Top) that differ in color or shape but not both. Here, for example, color (black) is redundant but shape is not. Subjects choose one object; if correct (Middle) there is a smiley face underneath; if incorrect (Bottom), there is no smiley face. Subjects learn eight such pairs concurrently. Healthy subjects appear to base learning on the relevant features while ignoring redundant ones (here, mushroom beats frame regardless of color), while individuals with hippocampal damage appear to learn based on all features (black-mushroom beats black-frame). Either is a valid approach to initial learning. (B) Healthy elderly subjects learn the initial discriminations at the same speed regardless of whether or not they have hippocampal atrophy (HA) revealed on neuroimaging. (C) Initial learning is followed by a transfer phase in which the familiar stimuli are presented with novel irrelevant features. Thus, the example in (A) might be changed to a green mushroom vs. a green frame. Now, the style of learning in phase 1 is critical: healthy individuals (noHA) transfer nearly perfectly, since their learning was not based on the irrelevant features anyway. By contrast, individuals with HA transfer poorly; in fact, they take about as long to learn the correct answers in phase 2 as they did on the initial learning phase—indicating little if any information from phase 1 transfers to phase 2. (Fig. adapted from Myers et al., 2002) (For interpretation of the reference to colour in this legend, the reader is referred to the web version of this article).

1996). However, we predicted that during this simple learning, the hippocampal region sets up representations that will allow subsequent generalization later, when task demands change. As a result, the model predicts that hippocampal damage may appear to spare initial learning, but the learning will be qualitatively different, and thus there will be impairments on subsequent transfer.

To test this idea, we administered our colored-shape discrimination task to a population of healthy elderly individuals, some of whom had hippocampal atrophy (HA) visible on magnetic resonance images (MRI) of the brain. As shown in Fig. 4(B), both the HA and no-HA individuals could learn well (Myers et al., 2002a,b). But this initial learning was followed by a transfer phase in which irrelevant features changed. For example, having learned to choose a black mushroom over a black frame (as in Fig. 4(A)), the transfer might include a choice between a green mushroom and a green frame. In other cases, shape might be the irrelevant feature: having learned to choose a red diamond over a yellow diamond, people might be challenged to choose between a red circle and a yellow circle. Under these circumstances, healthy elderly with no HA generalized very well, making few errors (Fig. 4(C)); by contrast, elderly with even relatively mild HA were impaired at generalizing (Myers et al., 2002a,b). Apparently, although, they had learned the initial task quickly, they had done so without the benefit of hippocampal representations that support subsequent transfer. The same kind of spared learning but impaired transfer has also been observed in elderly individuals with HA using an acquired equivalence task, in which individuals first learn a number of associations, and then have to generalize when the familiar items are presented in novel recombinations. Under these conditions, healthy elderly tend to generalize well, while those with HA generalize poorly (Myers, Shohamy, Gluck, Grossman, Kluger and Ferris, 2003a). Since mild hippocampal atrophy is a risk factor for subsequent development of Alzheimer's disease, even in individuals who are not yet showing any behavioral impairments (de Leon, George, Stylopoulos, Smith, & Miller, 1989), transfer tasks such as this, which are sensitive to

hippocampal atrophy in otherwise healthy individuals, may have some clinical utility as a diagnostic tool.

On the other hand, we expect a different pattern of results in patients with Parkinson's disease. Parkinson's disease results in neuronal death and dysfunction of the basal ganglia, a brain area involved in learning new stimulus-response associations based on feedback. In these patients, the initial learning may be slow, but—since, these patients have an intact hippocampal system—transfer and other hippocampal-dependent function should be spared. In fact, this appears to be the case in our learned irrelevance, discrimination, and acquired equivalence tasks (Myers et al., 2003a,b; Shohamy, Myers, Geghman, Sage, & Gluck, 2005).

3. Converging evidence from functional brain imaging in humans

Another method for testing predictions of the model for human learning is functional brain imaging. In particular, our model expects that the medial temporal lobes should be very active early in training when subjects are learning about stimulus–stimulus regularities and evolving a new stimulus representation, but less active later in training when other brain regions (e.g. the basal ganglia) are using these representations to perform the task. In a functional neuroimaging study (Poldrack, Clark, Pare-Blagoev, Shohamy, Creso-Moyano and Myers, 2001), we used a probabilistic category learning task developed in our laboratory that we call the 'weather prediction' task based on an early study by Gluck and Bower (1988). This task involves learning to predict the weather ('sun' or 'rain') based on the presence or absence of four different tarot cards with various geometric features. Each of these cards is partially diagnostic of either rain or sun. Although the probabilistic nature of this task precludes perfect performance, subjects do gradually improve their ability to predict the weather correctly based on these cards. As expected by the Gluck and Myers (1993) model, fMRI documented that activity in the hippocampal-region (medial temporal lobe) was highest early in training and then tapered off; in contrast, basal ganglia activity was low at first and increased during training. This was consistent with our earlier amnesic study using the same task (Knowlton, Squire, & Gluck, 1994), in which we found a deficit in amnesic patients during late-stages of training which our modeling suggests is due to a failure to acquire appropriate stimulus representations early in training (Gluck, Oliver, & Myers, 1996).

4. Biological substrates of representational processing in the hippocampal region

While the original cortico-hippocampal model (Gluck & Myers, 1993, 2001) had considerable success at accounting for behaviors of intact and hippocampal-lesioned animals and humans, it has several limitations. The first and most obvious is that it treats the entire hippocampal region as a single functional unit. We describe below how aspects of the earlier model have been mapped onto more physiological mechanisms.

5. Learning rates and septo-hippocampal cholinergic modulation

The medial septum/diagonal band complex, a structure lying in the basal forebrain, provides important cholinergic inputs to the hippocampus. Physiological studies and prior computational models by Michael Hasselmo suggest that this septo-hippocampal cholinergic input can modulate whether the hippocampus is acting primarily to store new, incoming information, or to retrieve previously-stored information (Hasselmo, 1999; Hasselmo & Schnell, 1994; Meeter, Talamini & Murre, 2004). Further, the hippocampus has outputs that travel back to the septum, suggesting that the hippocampus might be able to self-regulate its own storage/retrieval dynamics by dynamically adjusting the amount of cholinergic inputs it receives. Collaborating with Hasselmo, we implemented this idea within our computational model by assuming that the hippocampal learning rate is dependent on the amount of septo-hippocampal cholinergic input (Fig. 5(A)), and that the amount of cholinergic input is determined by a self-regulating feedback loop from hippocampus to medial septum (Myers, Ermita, Harris, Hasselmo, Solomon and Gluck, 1996; Myers, Ermita, Hasselmo, & Gluck, 1998; Rokers, Myers, & Gluck, 2000).

This model of septo-hippocampal cholinergic function can account for a range of data regarding the effects of septal lesion as well as the effects of various cholinergic drugs. For example, as shown in Fig. 5(B), the model correctly predicts that reducing levels of hippocampal acetylcholine (through septal lesions or administration of an anticholinergic drug) reduces learning, while increasing acetylcholine (through administration of a cholinergic agonist) can speed learning—but only to a certain point. Too much acetylcholine leads to a very high learning rate in the model, at which point the system becomes unstable. The same effect is seen in normal animals given a high dose of cholinergic drugs (see Myers et al., 1996). In recent studies in our lab, we have also confirmed several predictions of this model regarding the consequences of medial septal lesions on eyeblink conditioning (Allen et al., 2002b).

One interesting prediction of our septo-hippocampal model is that the memory deficits following hippocampal disruption (e.g. via medial septal lesion) might be subtly different from those following outright hippocampal removal. That is, with hippocampal lesion, areas like the

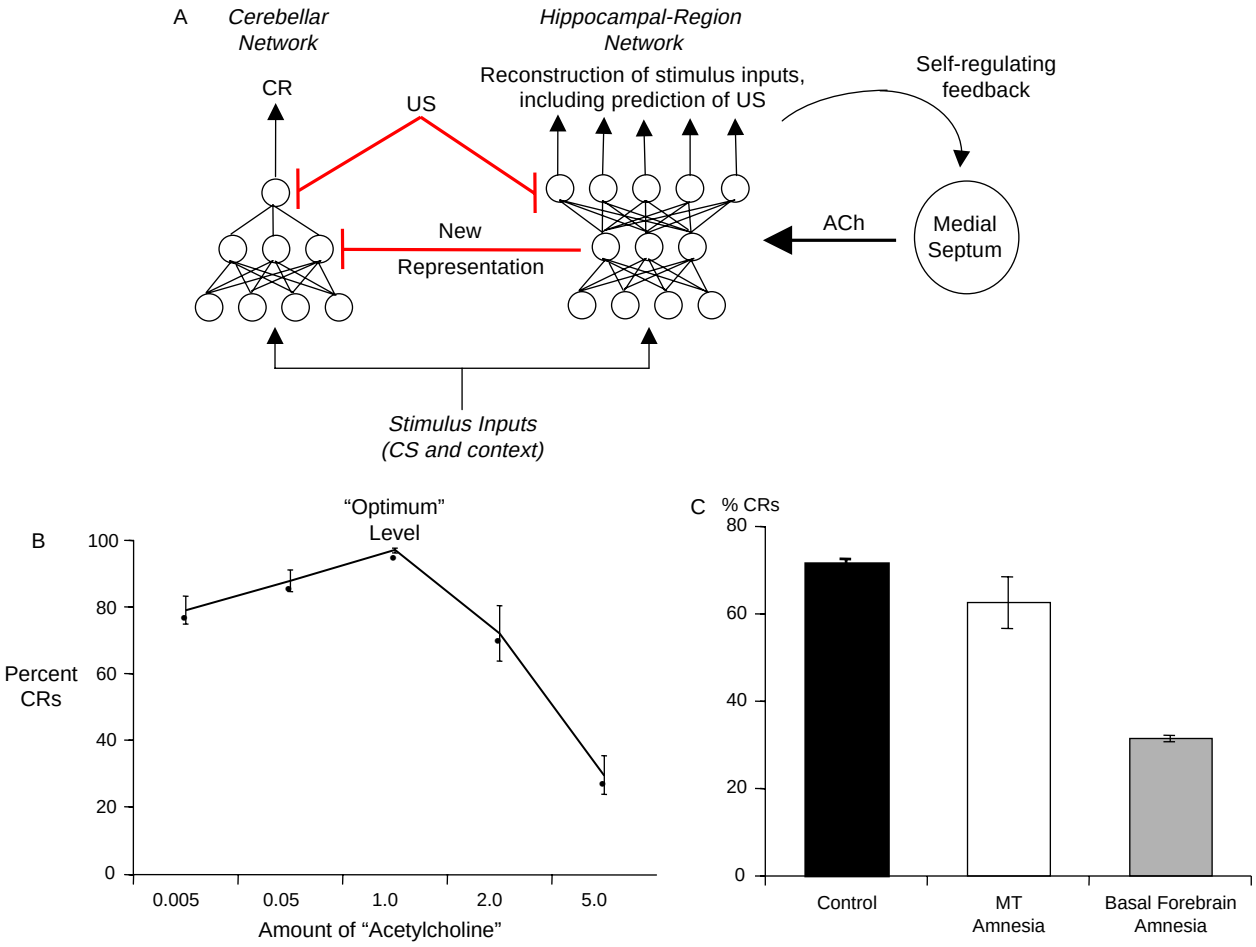

Fig. 5. (A) Adding medial septal–hippocampal interactions to the computational model. (B) The model correctly predicts that increasing septohippocampal cholinergic levels will facilitate learning—to a point, beyond which learning is actually retarded. (C) Individuals with basal forebrain damage often show amnesia superficially similar to individuals with medial temporal damage; however, unlike medial temporal amnesics, basal forebrain amnesics *are* impaired at simple delay eyeblink classical conditioning. (B adapted from Myers et al., 1996; C adapted from Myers et al., 2002).

cerebellum may be able to accomplish some simple learning on their own. However, with the hippocampus present but dysfunctional, 'noisy' hippocampal outputs will disrupt cerebellar function, slowing learning. In other words, a dysfunctional hippocampus might be worse than none at all.

We tested this prediction by comparing patients with anterograde amnesia following either medial temporal (hippocampal) or basal forebrain damage. Both patient groups show a similar abolition of new declarative learning. However, whereas medial temporal amnesia spares simple delay eyeblink conditioning in animals and humans, basal forebrain amnesia impairs it (Fig. 5(C); Myers, DeLuca, Schultheis, Schnirman, Ermita and Diamond, 2001; Myers, Bryant, DeLuca, & Gluck, 2002a,b). This finding has potential clinical relevance: currently, 'amnesia' is generally treated as a single clinical syndrome; better understanding of the ways in which non-declarative memory is differentially affected by medial temporal vs. basal forebrain damage should allow the development of rehabilitation techniques targeted at each population's unique pattern of impaired and spared abilities.

6. Redundancy compression and the entorhinal cortex

The hippocampal region is comprised of several distinct and interacting units, including the hippocampus proper (subfields CA1 and CA3), the dentate gyrus, the entorhinal cortex, and the subiculum (Fig. 6). Each of these areas has unique anatomical and physiological characteristics, and each may be expected to provide a unique information-processing function that contributes to the workings of the hippocampal region as a whole.

The entorhinal cortex is the primary path by which sensory information reaches hippocampus. Building on an earlier model of paleocortex by Richard Granger and colleagues (Ambros-Ingerson, Granger, & Lynch, 1990; Coultrip, Granger, & Lynch, 1992), we noted (Gluck & Granger, 1993) that the anatomical and physiological

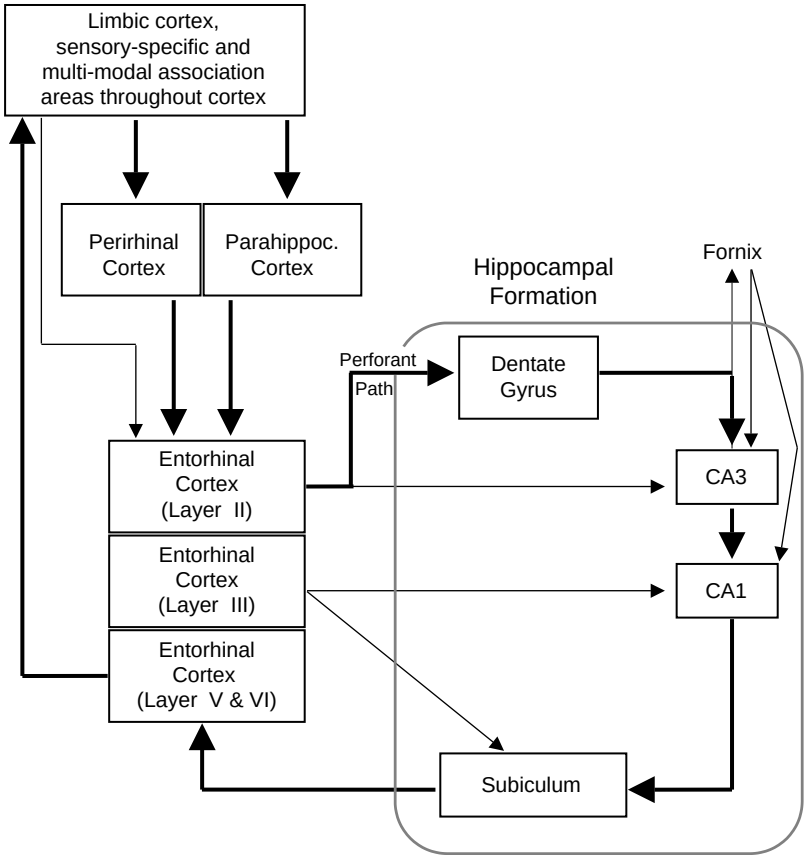

Fig. 6. Schematic of some important substructures and information pathways within the hippocampal region.

characteristics of the entorhinal cortex would be consistent with an emergent function of clustering and compressing input patterns. In collaboration with Granger, we suggested that the entorhinal cortex might, therefore be the place where redundancy compression, one of our postulated biases for hippocampal-region re-representation, could take place (Myers et al., 1995).

In Myers et al. (1995), we proposed a model of superficial entorhinal cortex as an unsupervised competitive network: a layer of units representing superficial entorhinal layer II excitatory neurons, receiving sparse multimodal inputs, and grouped into small patches whose members are reciprocally connected with an inhibitory interneuron (Fig. 7(A)). The result of excitatory–inhibitory interaction in each patch is the emergence of lateral competition, approximating 'winner-take-all' activity. The winning nodes undergo LTP-like plasticity, increasing their likelihood of winning the competition when similar inputs are presented in future. The resulting network performs unsupervised clustering based on surface similarities among stimuli, but also based on stimulus co-occurrence: two stimuli which appear together are treated as a single, compound stimulus and assigned to the same cluster. Later, learning about one stimulus will generalize to other stimuli in the same cluster. This assignment of co-occurring stimuli to the same representational cluster results in the same kind of representational compression which Gluck and Myers (1993) previously proposed to be one constraint biasing new stimulus representations in the hippocampal region. Myers et al. (1995) therefore proposed that the entorhinal cortex contained sufficient circuitry to implement stimulus–stimulus redundancy compression.

When this entorhinal cortex model is connected to our existing cerebellar model (Fig. 7(B)), the resulting model continues to show those hippocampal-region-dependent phenomena that depend on redundancy compression. For example, latent inhibition and learned irrelevance are two effects that we have explained as resulting from compression of cues and contexts during the initial exposure phase (Myers & Gluck, 1994). Both effects are disrupted by broad hippocampal-region damage in animals and in the computational model; however, the entorhinal-cerebellar model of Fig. 8(B) can demonstrate both these effects, suggesting that the effects should survive selective hippocampal lesions (sparing entorhinal cortex) in animals. We confirmed these predictions of our model empirically: selective hippocampal lesions disrupt both latent inhibition and learned irrelevance in rabbit eyeblink conditioning, while lesions that include the entorhinal cortex abolish both effects (Allen et al., 2002a,b; Shohamy et al., 2000).

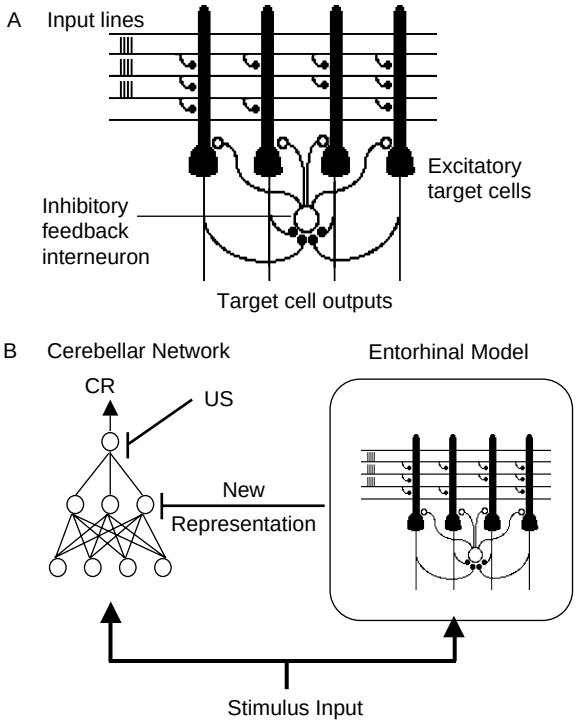

Fig. 7. (A) The entorhinal cortex model incorporates known anatomical, physiological and plasticity constraints. (B) In tandem with the cerebellar model, the resulting function approximates a selective lesion that eliminates hippocampus but spares entorhinal cortex.

Of course, the entorhinal cortex is only one component of the hippocampal region, and our entorhinal model only computes one portion (redundancy compression) of our proposed hippocampal-region function. Ongoing work, described in more detail below, seeks to add additional modules representing other hippocampal-region substructures with additional functionality, to try to capture the full range of hippocampal-region contributions to incremental learning.

7. General discussion and future directions

Our general approach has been to work 'top-down', starting with connectionist network models that instantiate information-processing theories of the computations required to explain behavior, and then showing how these functions could arise 'bottom-up' from the anatomy and physiology of specific brain regions. In this way, computational models can help to elucidate the principles of learning and memory at multiple levels of analysis, from behavioral processes through neural circuits.

There are many limitations to our work to date, some of which we have begun to address in ongoing work. We turn now to review these opportunities for future modeling. They fall into two main categories: (1) modeling which seeks to add greater *biological depth* and detail to the past work, and (2) modeling which expands the *behavioral breadth* to a wider range of learning and memory phenomena.

8. Future aims to enhance biological depth

Although we have made some progress at replacing components of our original Gluck and Myers (1993) 'top-down' model with biologically-plausible components, much still remains to be done. As described above, we have argued that certain features of our earlier theory of hippocampal-region function can be mapped to specific brain substrates: redundancy compression to the entorhinal cortex and learning rate modulation to septo-hippocampal feedback loops. Several important future aims for elaborating the biological details of our model framework for associative learning are noted below, along with some preliminary results where available.

8.1. Dentate gyrus

The entorhinal model described above assumes that the entorhinal cortex is a substrate for representational compression, one of the functions that Gluck and Myers (1993) ascribed to the hippocampal region. However, in addition to representational compression, Gluck and Myers (1993) proposed that the hippocampal region could mediate representational differentiation: increasing the difference between representations of stimuli that predict different future events (e.g. reinforcement). Others have suggested that the anatomy of the dentate gyrus would be ideally suited for this function (Hasselmo & Wyble, 1997; O'Reilly & McClelland, 1994).

Evidence for localizing differentiation in the dentate gyrus comes from studies by Mark West and others, who have used multicellular recordings from the hippocampal region in behaving animals to detect representational changes. The pattern of firing activity across a set of neurons is analogous to the activities across a set of nodes in a neural network, and can be viewed as the brain's representation of the current inputs. In one study, recordings were taken from dentate gyrus while rats learned a discrimination (respond to stimulus A+ but not to B−) (Deadwyler, West & Lynch, 1979). Early in discrimination learning, neuronal activity looked similar after presentation of either A or B. However, as the discrimination was learned, neuronal discharge in the dentate gyrus differentiated the two stimuli; specifically, neurons might respond to both stimuli, but only the rewarded stimulus (A+) elicited sustained activity. This and related findings suggest that stimulus representations in the dentate gyrus are differentiated if the stimuli are associated with different reinforcement outcomes and/or different responses (see also Bostock, Muller, & Kubie, 1991; Cahusec, Rolls, Miyashita, & Niki, 1993). Further,

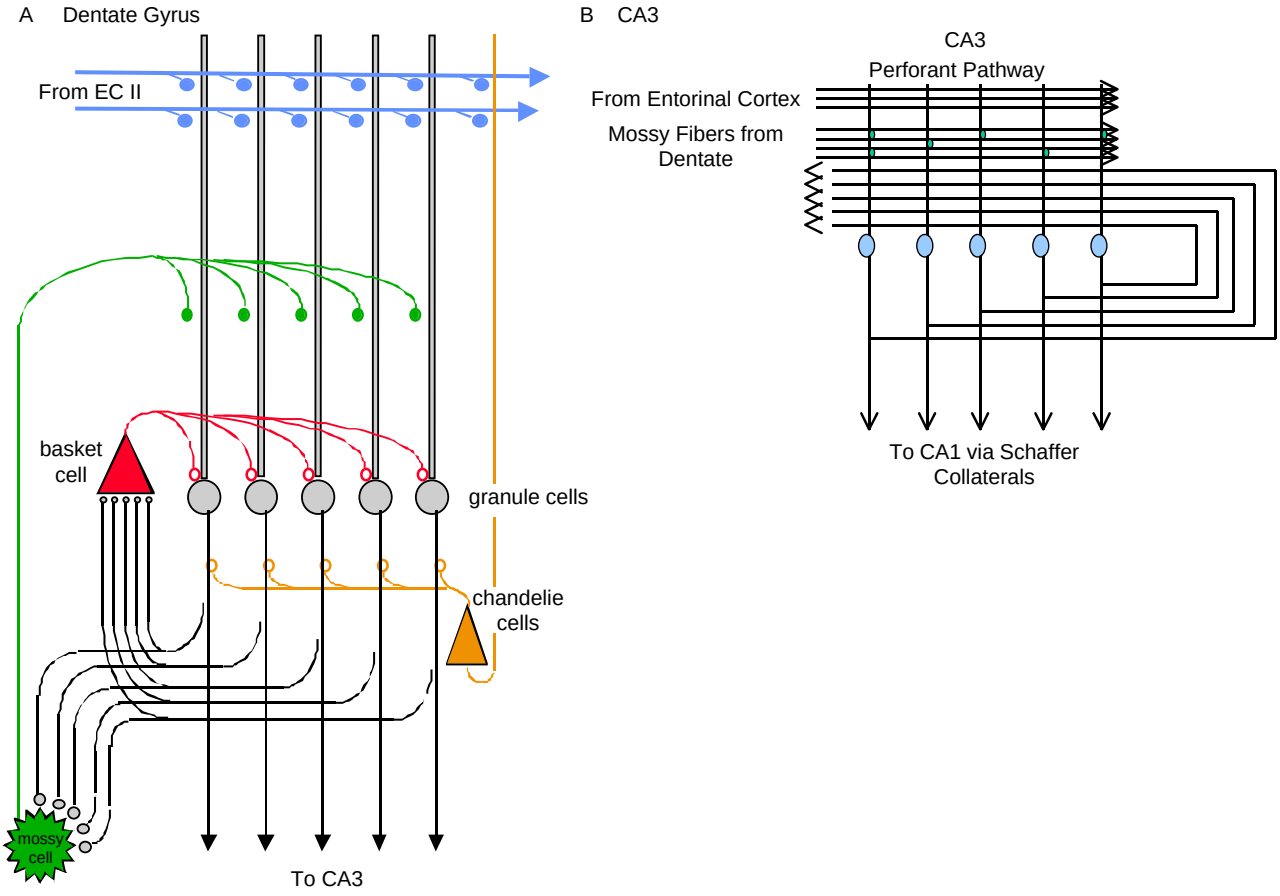

Fig. 8. (A) Schematic of dentate gyrus model. (B) Schematic of hippocampal subfield CA3 model.

this differentiation is not visible in entorhinal cortex, which provides the dentate gyrus's principal input, suggesting that it is the dentate gyrus that is performing the representational differentiation (Deadwyler et al., 1979).

In preliminary work, we have constructed a computational model that is based on known anatomical and physiological characteristics of dentate gyrus (Fig. 8(A)). The model includes 500 granule cells that receive sparse and weakly excitatory input from 100 entorhinal afferents via the perforant path; granule cell outputs (the mossy fibers) form the principal output of the dentate model. The granule cells are inhibited by local circuit interneurons, basket cells, which serve to roughly normalize overall firing activity and allow 'winner-take-all' processing among clusters of granule cells. Plasticity is implemented via a learning rule that incorporates aspects of LTP and LTD as observed in the dentate gyrus.

This preliminary model sparsifies inputs: it takes an input pattern from a relatively small number of entorhinal nodes and expands it into a representation across a much larger number of dentate granule cells. This has the effect of differentiating representations in general, as posited in the original hippocampal-region model of Gluck and Myers (1993).

8.2. Hippocampal subfields CA3 and CA1

The hippocampus proper can be broken down into several subfields, including CA3 and CA1. Many previous researchers have argued that CA3 may function as an autoassociator (Fig. 8(B)), based on its high degree of internal recurrency. This function would allow CA3 to rapidly store random patterns and then retrieve them later when given whole or partial cues (Levy, 1996; McClelland, McNaughton, & O'Reilly, 1995; Treves & Rolls, 1994). We are currently working to connect such a CA3 autoassociator into our system, where it would receive inputs that had been pre-processed by entorhinal cortex and dentate gyrus and store them.

From CA3, information travels to CA1 via the Schaffer collaterals as well as via a direct pathway from entorhinal cortex (the perforant path). Previous researchers have suggested that this dual input would allow CA1 to function as a comparator, specifically comparing the direct entorhinal input against the reconstructed pattern provided from CA3 (Hasselmo & Schnell, 1994). If the patterns are highly similar, this would mean that the information had been successfully stored in and reconstructed in CA3. Such a comparator function is consistent with the finding of so-called 'mis-match' cells in CA1, which become active when

an animal is confronted with novel or unexpected inputs. This in turn would signal the system that learning is needed to store the new inputs, possibly by signaling the septum for acetylcholine to facilitate hippocampal learning. Finally, outputs from CA3 and/or CA1 travel back out via subiculum and entorhinal cortex and eventually to the cortical areas where they arose.

In future work, we plan to include CA3 and CA1 modules, together with entorhinal, dentate, and cortico/cerebellar modules, and compare them against existing data from intact and hippocampal-region-lesioned animals, as well as with existing data regarding the effects of selective lesions, such as ibotenate lesions of hippocampus that spare entorhinal cortex.

Additional elaboration of these models would allow for examining how distribution of different receptor subtypes (e.g. muscarinic vs. nicotinic cholinergic receptors) might affect hippocampal processing and response to pharmacological agents, and investigating possible roles for other neuromodulators in the hippocampal region such as dopamine, stress hormones such as glucocorticoids, and reproductive hormones such as estrogen.

8.3. Basal ganglia and dopamine

We mentioned above that the basal ganglia play an important role in incrementally-acquired non-declarative learning, and that such incrementally-acquired learning is disrupted in Parkinson's patients, who have basal ganglia dysfunction (e.g. Myers et al., 2003a; Shohamy et al., 2005).

Neurons in the ventral tegmental area (VTA) and the substantia nigra pars compacta send dopaminergic projections to the basal ganglia, medial prefrontal and other limbic forebrain regions. Stimulation of this system is a strong reinforcer in animals, suggesting a role in reward. Dopamine was initially seen as a generic reward signal, but release from this system correlates more strongly with predictors of a reward than with reward itself (Schultz, Apicella, & Ljungberg, 1993).

One role that has been suggested for dopamine is as a reinforcing reward signal, guiding instrumental learning and the association of stimuli to reward (Daw & Touretzky, 2001; Doya, 2000; Schultz, Dayan, & Montague, 1997). As would be expected if it has such a role, dopaminergic output obeys a temporal difference rule, which means that it signals not reward as such but a change in the likelihood of reward (Schultz et al., 1993; 1997). One may wonder what in turn controls this reward signal. An intriguing possibility is that projections from the basal ganglia play a role in the control of substantia nigra/VTA dopamine release (Brown, Bullock, & Grossberg, 1999). Integration of this control in a model of the basal ganglia will be of great value to a better understanding of learning and memory.

Interaction between the basal ganglia and medial temporal lobe structures may have an important role in the generation of firing patterns correlated with instrumental behavior. The basal ganglia, and in particular the nucleus accumbens, have strong recurrent connections with the medial temporal lobe (especially field CA1 of the hippocampus and the entorhinal cortex). It has been shown that neurons in hippocampal system subfields exhibit stimulus-locked firing in instrumental conditioning (e.g. Deadwyler, West, & Lynch, 1979). It may thus be hypothesized that hippocampal inputs to the nucleus accumbens are key to the time-locked firing in this latter structure shown in instrumental behavior (Peoples et al., 1997), perhaps by directly activating a representation of the conditioned response in the nucleus accumbens. We plan to integrate these ideas into our model system, by allowing the hippocampal-system modules to interact with a basal ganglia module that is in turn modulated by dopaminergic inputs carrying reinforcement signals.

9. Future aims to enhance behavioral breadth

To date, most of our computational modeling has focused on a relatively constrained area of learning: simple conditioning and those behaviors (like human category learning) that can be understood in terms of conditioning (Gluck & Bower, 1988). However, the hippocampus and associated structures play clear and important roles in other domains, including (but not limited to) declarative learning (Squire, 1987), temporal learning (Levy, 1989) and spatial learning (O'Keefe & Nadel, 1978). Several future directions for expanding the behavioral breadth of these models are suggested below.

9.1. Temporal processing

Even within the limited domain of classical conditioning, our models do not account for all findings; for example, one important finding is that when the CS and US do not overlap (a paradigm known as trace conditioning), hippocampal lesion disrupts learning (e.g. Moyer et al., 1990). Our models do not include temporal information, and so cannot be used to examine these and other paradigms where time is a parameter. To some extent, this can be addressed by adding feedback connections to the model (Zackheim, Myers, & Gluck, 1998); our working hypothesis is that many of the 'temporal' aspects of hippocampal-region processing can be understood as emerging from the information-processing functions of the region (Gluck, Allen, Myers, & Thompson, 2001). Work by Chip Levy and colleagues suggests one approach to incorporating temporal and sequence learning into hippocampal models (August & Levy, 1999; Levy, 1989; 1996).

9.2. Episodic memory

The hippocampus is not only important in associative learning, but also in learning of autobiographical events

('episodes') (Eichenbaum, 1992; Marr, 1971; Scoville & Milner, 1957). Episodic learning seems to pose very different demands on the hippocampus than the associative learning tasks discussed above. Associative learning is sensitive to behavioral outcome, while episodic learning is thought of as unsupervised, automatic coding of whatever is present. Episodic learning is fast (often one trial), while associative learning is incremental and slow. And yet a single brain system—the hippocampal region—appears to play a critical role in both kinds of learning.

We set out to reconcile the Gluck and Myers model with a generic version of an episodic memory model (Meeter, Myers, & Gluck, 2005). Instead of using a hippocampal autoencoder, we started with a multilayer model capable of forming episodic memories, loosely based on earlier episodic memory models (Meeter, Murre, & Talamini, 2002; Talamini, Meeter, Murre, Elvevåg, & Goldberg, 2005). This simple model of episodic memory, as shown in Fig. 9, contains three layers. An input layer, modeling the neocortex, codes for stimuli and context features. A second layer stands for the parahippocampal region: the perirhinal, entorhinal and postrhinal/parahippocampal cortices. This layer has integrated representations, with some nodes coding mostly for context features and some mostly for stimuli, but all nodes also getting input of the other kind. The third layer stands for the hippocampus proper, with nodes representing dentate granule cells and/or pyramidal cells in CA3 and CA1. This hippocampal layer forms a compact code for the whole situation in which the organism finds itself, for which we use the term 'ensemble' (Murnane, Phelps, & Malmberg, 1999). Such representations form the basis of episodic memory. (Later, of course, this network could be elaborated to include more physiological detail, as outlined above, but in these initial studies, the purpose was to attempt to reconcile associative and episodic behavior in a single, simple model system.)

As in the Gluck and Myers (1993) model, the network simulating the hippocampal region interacts with other regions in the brain, in which the outputs of memory are coded. For classical conditioning, the cerebellum is most relevant, but for other tasks of incremental learning one would have to include output regions for rewards and operant behaviors (basal ganglia) and for fear responses (amygdala). The cerebellar circuit is the only output structure implemented to date. All three layers of the episodic network project to the output modules. These connections allow the output modules to attach behavioral significance to simple and complex representations of the same set of stimuli, thereby allowing stimulus configurations to have different associations than the constituent stimuli on their own.

The model also follows earlier learning theories, such as Wagner's Sometimes Opponent Process (Wagner, 1981) in assuming that responses adapt to familiar stimuli. The first time a stimulus is presented, nodes respond strongly, but this response decreases with repeated presentations of the same stimulus. This has indeed been found experimentally (Li, Miller, & Desimone, 1993; Xiang & Brown, 1998).

Since the 1970s, many researchers in the field of episodic recognition memory have argued that recognition judgments can be based on a fuzzy feeling that the item matches old memories, usually referred to as familiarity (e.g. Atkinson & Juola, 1974; Humphreys, Bain, & Pike, 1989; Mandler, 1980; Yonelinas, 2002). When neural responses in the perirhinal cortex to novel and familiar stimuli are directly compared, neural responses to a stimulus decrease

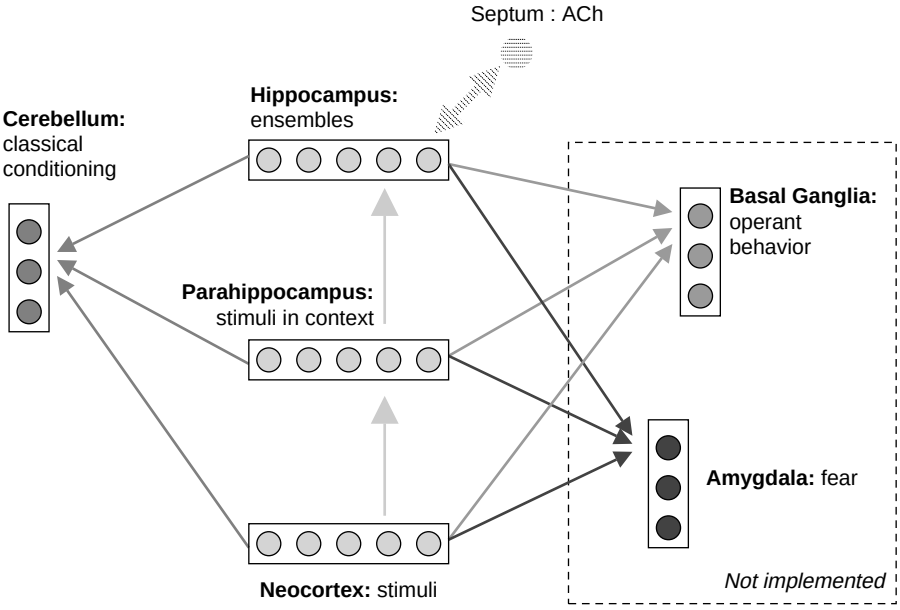

Fig. 9. Brain regions that play a part in incremental learning. See main text for explanation.

in the perirhinal cortex with increasing familiarity (Li, Miller, & Desimone, 1993; Xiang & Brown, 1998). These results are exactly what one would expect from SOP and our model: a state of high activity at the first presentations, with decreasing activity after stimulus repetitions. In Meeter et al. (2005), we agued that this familiarity effect, the decreased parahippocampal response to familiar stimuli, is what causes the effects of stimulus novelty on the speed of conditioning (other accounts can be found in Sohal & Hasselmo, 2000; Bogacz & Brown, 2002; 2003).

Other analyses presented in Meeter et al. (2005) show how this model can capture a wide range of behavioral data on both elementary associative learning (such as conditoning) as well as episodic retrieval, recognition, and familiarity. An important caveat, however, is that while this broader model integrates aspects of the earlier Gluck and Myers (1993) cortico-hippocampal model of classical conditioning, it is still a connectionist model using abstract nodes and links. Further work will be required to integrate these broader behavioral models with the more detailed and biologically-constrained models described earlier for specific brain regions of the hipppocampal region. See also work by James McClelland, Randy O'Reilly, and colleagues on relevant models of cortico-hippocampal function in episodic memory (McClelland et al., 1995; Norman & O'Reilly, 2003; O'Reilly & Norman, 2002; O'Reilly & Rudy, 2001).

9.3. Consolidation

An issue that often comes up with regard to the hippocampus is whether memories are consolidated from the hippocampus to the neocortex or not (for review, see Meeter & Murre, 2004). Marr (1971) suggested, for computational reasons, that the hippocampus was only a temporary store, with older memories being transferred from the hippocampus to the neocortex. Squire, Cohen and Nadel (1984) came to the same hypothesis to explain an old neuropsychological mystery, namely that recent memories seemed more vulnerable to hippocampal damage than more remote memories. This hypothesis was also included in the Gluck and Myers (1993) model. More recently, the view that memories are encoded has come under attack (Nadel & Moscovitch, 1997; Nadel, Samsonovitch, Ryan, & Moscovitch, 2000), and our recent modeling (Meeter et al., 2005) has shown that the consolidation hypothesis is not essential to explain associative learning phenomena.

10. Conclusion

A large body of neurobiological and behavioral data on associative learning can be accounted for by assuming that the hippocampal region—including hippocampus proper, entorhinal cortex, and other associated structures—performs an information processing function: compressing (or making more similar) the representations of inputs that co-occur or are otherwise redundant, while differentiating (or making less similar) the representations of inputs that predict different future events. This process can be instantiated in a computational model, where the hippocampal region is conceived as a functional unit performing predictive autoassociation. This leads to an important question: does such a computational function emerge naturally from the brain substrate? Here, we have reviewed our progress to date, both in developing a 'top-down' model to account for behavioral data, as well as attempting to instantiate the functionality of that top-down model via 'bottom-up' modules that are consistent with known features of the biological substrate.

The current modeling lays the groundwork for future directions that would increase the depth of detail of the biological modeling, as well as the breadth of behavioral phenomena addressed. In particular, we are working now to reconcile these kinds of incremental associative learning models with other models of the hippocampal region that account for the rapid formation of declarative memories.

References

Allen, M., Chelius, L., & Gluck, M. (2002a). Selective entorhinal and nonselective cortical–hippocampal region lesions, but not selective hippocampal lesions, disrupt learned irrelevance in rabbit eyeblink conditioning. *Cognitive, Affective, and Behavioral Neuroscience*, *2*(3), 214–226.

Allen, M. T., Padilla, Y., & Gluck, M. A. (2002b). Ibotenic acid lesions of the medial septum retard delay eyeblink conditioning in rabbits (*Oryctolagus cuniculus*). *Behavioral Neuroscience*, *116*(4), 733–738.

Ambros-Ingerson, J., Granger, R., & Lynch, G. (1990). Simulation of paleocortex performs hierarchical clustering. *Science*, *247*, 1344–1348.

Atkinson, R. C., & Juola, J. F. (1974). Search and decision processes in recognition memory. In D. H. Krantz, R. C. Atkinson, R. D. Luce, & P. Suppes (Eds.), *Contemporary developments in mathematical psychology*. San Franscisco, CA: Freeman.

August, D. A., & Levy, W. B. (1999). Temporal sequence compression by an integrate-and-fire model of hippocampal area CA3. *Journal of Computational Neuroscience*, *6*, 71–90.

Baldi, P., & Hornik, K. (1989). Neural networks and principal component analysis: Learning from examples without local minima. *Neural Networks*, *2*, 53–58.

Bogacz, R., & Brown, M. W. (2002). The restricted influence of sparseness of coding on the capacity of familiarity discrimination networks. *Network: Computations in Neural Systems*, *13*, 457–485.

Bogacz, R., & Brown, M. W. (2003). Comparison of computational models of familiarity discrimination in the perirhinal cortex. *Hippocampus*, *13*, 494–524.

Bostock, E., Muller, R., & Kubie, J. (1991). Experience-dependent modifications of hippocampal place cell firing. *Hippocampus*, *1*(2), 193–206.

Brown, J., Bullock, D., & Grossberg, S. (1999). How the basal ganglia use parallel excitatory and inhibitory learning pathways to selectively respond to unexpected rewarding cues. *The Journal of Neuroscience*, *19*, 10502–10511.

Cahusec, P., Rolls, E., Miyashita, Y., & Niki, H. (1993). Modification of the responses of hippocampal neurons in the monkey during the learning of a conditional spatial response task. *Hippocampus*, *3*(1), 29–42.

Coultrip, R., Granger, R., & Lynch, G. (1992). A cortical model of winner-take-all competition via lateral inhibition. *Neural Networks*, *5*, 47–54.

Daw, N. D., & Touretzky, D. S. (2001). Operant behavior suggests attentional gating of dopamine system inputs. *Neurocomputing*, *38–40*, 1161–1167.

Deadwyler, S., West, M., & Lynch, G. (1979). Activity of dentate granule cells during learning: Differentiation of perforant path input. *Brain Research*, *169*, 29–43.

de Leon, M., George, A., Stylopoulos, L., Smith, G., & Miller, D. (1989). Early marker for Alzheimer's disease: The atrophic hippocampus. *The Lancet*, *2*(8664), 672–673.

Doya, K. (2000). Complementary roles of basal ganglia and cerebellum in learning and motor control. *Current Opinion in Neurobiology*, *10*, 732–739.

Eichenbaum, H. (1992). The hippocampal system and declarative memory in animals. *Journal of Cognitive Neuroscience*, *4*, 217–231.

Eichenbaum, H., Mathews, P., & Cohen, N. J. (1989). Further studies of hippocampal representation during odor discrimination learning. *Behavioral Neuroscience*, *103*(6), 1207–1216.

Gabrieli, J., McGlinchey-Berroth, R., Carrillo, M., Gluck, M., Cermack, L., & Disterhoft, J. (1995). Intact delay-eyeblink classical conditioning in amnesia. *Behavioral Neuroscience*, *109*(5), 819–827.

Gluck, M. A., Allen, M. T., Myers, C. E., & Thompson, R. F. (2001). Cerebellar substrates for error-correction in motor-reflex conditioning. *Neurobiology of Learning and Memory*, *76*, 314–341.

Gluck, M., & Bower, G. (1988). From conditioning to category learning: An adaptive network model. *Journal of Experimental Psychology: General*, *117*(3), 225–244.

Gluck, M. A., & Granger, R. (1993). Computational models of the neural bases of learning and memory. *Annual Review of Neuroscience*, *16*, 667–706.

Gluck, M., & Myers, C. (1993). Hippocampal mediation of stimulus representation: A computational theory. *Hippocampus*, *3*, 491–516.

Gluck, M., & Myers, C. (1995). Representation and association in memory: A neurocomputational view of hippocampal function. *Current Directions in Psychological Science*, *4*(1), 23–29.

Gluck, M., & Myers, C. (1996). Integrating behavioral and physiological models of hippocampal function. *Hippocampus*, *6*(6), 643–653.

Gluck, M., & Myers, C. (2001). *Gateway to memory: An introduction to neural network modeling of the hippocampus in learning and memory*. Cambridge, MA: MIT Press.

Gluck, M., Myers, C., & Thompson, R. (1994). A computational model of the cerebellum and motor-reflex conditioning. In S. Zornetzer, J. Davis, C. Lau, & T. McKenna (Eds.), *An introduction to neural and electronic networks* (pp. 91–98). New York: Academic Press.

Gluck, M., Oliver, L., & Myers, C. (1996). Late-training amnesic deficits in probabilistic category learning: A neurocomputational analysis. *Learning and Memory*, *3*, 326–340.

Gormezano, I., Kehoe, E. J., & Marshall, B. S. (1983). Twenty years of classical conditioning research with the rabbit. *Progress in Psychobiology and Physiological Psychology*, *10*, 197–275.

Hasselmo, M. E. (1999). Neuromodulation: Acetylcholine and memory consolidation. *Trends in Cognitive Sciences*, *3*, 351–359.

Hasselmo, M. E., & Schnell, E. (1994). Laminar selectivity of the cholinergic suppression of synaptic transmission in rat hippocampal region CA1: Computational modeling and brain slice physiology. *Journal of Neuroscience*, *14*(6), 3898–3914.

Hasselmo, M. E., & Wyble, B. P. (1997). Free recall and recognition in a network model of the hippocampus: Simulating effects of scopolamine on human memory function. *Behavioural Brain Research*, *89*, 1–34.

Hinton, G. (1989). Connectionist learning procedures. *Artificial Intelligence*, *40*, 185–234.

Humphreys, M. S., Bain, J. D., & Pike, R. (1989). Different ways to cue a coherent system: A theory for episodic, semantic and procedural tasks. *Psychological Review*, *96*, 208–233.

Knowlton, B. J., Squire, L. R., & Gluck, M. A. (1994). Probabilistic category learning in amnesia. *Learning and Memory*, *1*, 106–120.

Levy, W. B. (1989). A computational approach to hippocampal function. In R. Hawkins, & G. Bower (Eds.), *Psychology of learning and motivation* (pp. 243–304). London: Academic Press.

Levy, W. B. (1996). A sequence predicting CA3 is a flexible associator that learns and uses context to solve hippocampal-like tasks. *Hippocampus*, *6*, 579–590.

Levy, W. B., Brassel, S. E., & Moore, S. D. (1983). Partial quantification of the associative synaptic learning rule of the dentate gyrus. *Neuroscience*, *8*(4), 799–808.

Li, L., Miller, E. K., & Desimone, R. (1993). The representation of stimulus familiarity in anterior inferior temporal cortex. *Journal of Neurophysiology*, *13*, 1918–1929.

Lubow, R. (1973). Latent inhibition. *Psychological Bulletin*, *79*, 398–407.

Mandler, G. (1980). Recognizing: The judgment of previous occurrence. *Psychological Review*, *87*, 252–271.

Marr, D. (1971). Simple memory: A theory for archicortex. *Philosophical Transactions of the Royal Society B*, *262*, 23–81.

McClelland, J. L., McNaughton, B. L., & O'Reilly, R. C. (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. *Psychology Review*, *102*, 419–457.

Meeter, M., & Murre, J. M. J. (2004). Consolidation of long-term memory: Evidence and alternatives. *Psychological Bulletin*, *130*, 843–857.

Meeter, M., Murre, J. M. J., & Talamini, L. M. (2002). A computational approach to memory deficits in schizophrenia. *Neurocomputing*, *44*, 929–936.

Meeter, M., Myers, C. E., & Gluck, M. A. (2005). Integrating incremental learning and episodic memory models of the hippocampal region. *Psychological Review, 112*, 560–585.

Meeter, M., Talamini, L. M., & Murre, J. M. J. (2004). Mode shifting between storage and recall based on novelty detection in oscillating hippocampal circuits. *Hippocampus*, *14*, 722–741.

Moyer, J., Deyo, R., & Disterhoft, J. (1990). Hippocampectomy disrupts trace eye-blink conditioning in rabbits. *Behavioral Neuroscience*, *104*(2), 243–252.

Murnane, K., Phelps, M. P., & Malmberg, K. (1999). Context-dependent recognition memory: The ICE theory. *Journal of Experimental Psychology: General*, *128*, 403–415.

Myers, C., Bryant, D., DeLuca, J., & Gluck, M. (2002a). Dissociating basal forebrain and medial temporal amnesic syndromes: Insights from classical conditioning. *Integrative Physiological and Behavioral Science*, *37*(2), 85–102.

Myers, C., DeLuca, J., Schultheis, M., Schnirman, G., Ermita, B., & Diamond, B. (2001). Impaired delay eyeblink classical conditioning in individuals with anterograde amnesia resulting from anterior communicating artery aneurysm rupture. *Behavioral Neuroscience*, *115*(3), 560–570.

Myers, C., Ermita, B., Harris, K., Hasselmo, M., Solomon, P., & Gluck, M. (1996). A computational model of the effects of septohippocampal disruption on classical eyeblink conditioning. *Neurobiology of Learning and Memory*, *66*, 51–66.

Myers, C., Ermita, B., Hasselmo, M., & Gluck, M. (1998). Further implications of a computational model of septohippocampal cholinergic modulation in eyeblink conditioning. *Psychobiology*, *26*(1), 1–20.

Myers, C., & Gluck, M. (1994). Context, conditioning and hippocampal re-representation. *Behavioral Neuroscience*, *108*(5), 835–847.

Myers, C., & Gluck, M. (1996). Cortico-hippocampal representations in simultaneous odor discrimination learning: A computational interpretation of Eichenbaum, Mathews & Cohen (1989). *Behavioral Neuroscience*, *110*(4), 685–706.

Myers, C., Gluck, M., & Granger, R. (1995). Dissociation of hippocampal and entorhinal function in associative learning: A computational approach. *Psychobiology*, *23*(2), 116–138.

Myers, C., Kluger, A., Golomb, J., Ferris, S., de Leon, M., Schnirman, G., et al. (2002b). Hippocampal atrophy disrupts transfer generalization in non-demented elderly. *Journal of Geriatric Psychiatry and Neurology*, *15*(2), 82–90.

Myers, C., McGlinchey-Berroth, R., Warren, S., Monti, L., Brawn, C. M., & Gluck, M. A. (2000). Latent learning in medial temporal amnesia: Evidence for disrupted representational but preserved attentional processes. *Neuropsychology*, *14*(1), 3–15.

Myers, C., Shohamy, D., Gluck, M., Grossman, S., Kluger, A., Ferris, S., et al. (2003a). Dissociating hippocampal vs. basal ganglia contributions to learning and transfer. *Journal of Cognitive Neuroscience*, *15*(2), 185–193.

Myers, C., Shohamy, D., Gluck, M., Grossman, S., Onlaor, S., & Kapur, N. (2003b). Dissociating medial temporal and basal ganglia memory systems with a latent learning task. *Neuropsychologia*, *41*, 1919–1928.

Nadel, L., & Moscovitch, M. (1997). Memory consolidation, retrograde amnesia and the hippocampal complex. *Current Opinion in Neurobiology*, *7*, 217–227.

Nadel, L., Samsonovitch, A., Ryan, L., & Moscovitch, M. (2000). Multiple trace theory of human memory: Computational, neuroimaging and neuropsychological results. *Hippocampus*, *10*, 352–368.

Norman, K. A., & O'Reilly, R. C. (2003). Modeling hippocampal and neocortical contributions to recognition memory: A complementary learning systems approach. *Psychological Review*, *110*, 611–646.

O'Keefe, J., & Nadel, L. (1978). *The hippocampus as a cognitive map*. Oxford: Clarendon University Press.

O'Reilly, R. C., & McClelland, J. L. (1994). Hippocampal conjunctive encoding, storage, and recall: Avoiding a trade-off. *Hippocampus*, *4*, 661–682.

O'Reilly, R. C., & Norman, K. A. (2002). Hippocampal and ceocortical contributions to memory: Advances in the complementary learning systems framework. *Trends in Cognitive Sciences*, *6*, 505–510.

O'Reilly, R. C., & Rudy, J. W. (2001). Conjunctive representations in learning and memory: Principles of hippocampal and cortical function. *Psychological Review*, *108*, 311–345.

Peoples, L. L., Uzwiak, A. J., Fred, G., & West, M. O. (1997). Operant behavior during sessions of intravenous cocaine infusion is necessary and sufficient for phasic firing of single nucleus accumbens neurons. *Brain Research*, *757*, 280–284.

Poldrack, R. A., Clark, J., Pare-Blagoev, E. J., Shohamy, D., Creso-Moyano, J., Myers, C. E., et al. (2001). Interactive memory systems in the brain. *Nature*, *414*, 546–550.

Rescorla, R., & Wagner, A. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and non-reinforcement. In A. Black, & W. Prokasy, *Classical Conditioning II: Current Research and Theory* (pp. 64–99). New York: Appleton-Century-Crofts.

Rokers, B., Myers, C., & Gluck, M. (2000). A dynamic model of learning in the septo-hippocampal system. *Neurocomputing*, *32–33*, 501–507.

Rumelhart, D., Hinton, G., & Williams, R. (1986). Learning internal representations by error propogation. In D. Rumelhart, & R. McClelland (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition* (pp. 318–362). Cambridge, MA: MIT Press.

Schmaltz, L., & Theios, J. (1972). Acquisition and extinction of a classically conditioned response in hippocampectomized rabbits (*Oryctolagus cuniculus*). *Journal of Comparative and Physiological Psychology*, *79*, 328–333.

Schultz, W., Apicella, P., & Ljungberg, T. (1993). Responses of monkey dopamine neurons to rewards and conditioned stimuli during successive steps of learning a delayed response task. *Journal of Neuroscience*, *13*, 900–913.

Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science*, *275*, 1593–1599.

Scoville, W. B., & Milner, B. (1957). Loss of recent memories after bilateral hippocampal lesions. *Journal of Neurology, Neurosurgery and Psychiatry*, *20*, 11–21.

Shohamy, D., Allen, M., & Gluck, M. (2000). Dissociating entorhinal and hippocampal involvement in latent inhibition. *Behavioral Neuroscience*, *114*(5), 867–874.

Shohamy, D., Myers, C. E., Geghman, K. D., Sage, J. & Gluck, M. A. (2005, in press). L-Dopa impairs learning, but not generalization, in Parkinson's disease. *Neuropsychologia*.

Sohal, V. S., & Hasselmo, M. E. (2000). A model for experience-dependent changes in the responses of enfero-temporal neurons. *Network: Computations in Neural Systems*, *11*, 169–190.

Solomon, P., & Moore, J. (1975). Latent inhibition and stimulus generalization of the classically conditioned nictitating membrane response in rabbits (*Oryctolagus cuniculus*) following dorsal hippocampal ablation. *Journal of Comparative and Physiological Psychology*, *89*, 1192–1203.

Squire, L. R. (1987). *Memory and brain*. New York: Oxford University Press.

Squire, L. R., Cohen, N. J., & Nadel, L. (1984). The medial temporal region and memory consolidation: A new hypothesis. In H. Weingarter, & E. Parker (Eds.), *Memory consolidation* (pp. 185–210). Hillsdale, NJ: Lawrence Erlbaum.

Talamini, L. M., Meeter, M., Murre, J. M. J., Elvevag, B., & Goldberg, T. E. (2005). Redyced parahippocampal connectivity produces schizophrenia-like deficits in simulated neural circuits with reduced parahippocampal connectivity. *Archives of General Psychiatry*, 62, 485–493.

Thompson, R. (1986). The neurobiology of learning and memory. *Science*, *233*, 941–947.

Treves, A., & Rolls, E. T. (1994). A computational analysis of the role of the hippocampus in memory. *Hippocampus*, *4*, 373–391.

Wagner, A. (1981). SOP: A model of automatic memory processing in animal behavior. In N. Spear, & R. R. Miller (Eds.), *Information processing in animals: Memory mechanisms* (pp. 5–47). Hillsdale, NJ: Erlbaum.

Widrow, B., & Hoff, M. (1960). Adaptive switching circuits. *Institute of Radio Engineers, Western Electronic Show and Convention Record*, *4*, 96–104.

Xiang, J. Z., & Brown, M. W. (1998). Differential neuronal encoding of novelty, familiarity, and recency in regions of the anterior temporal lobe. *Neuropharmacology*, *37*, 657–676.

Yonelinas, A. P. (2002). The nature of recollection and familiarity: A review of 30 years of research. *Journal of Memory and Language*, *46*, 441–517.

Zackheim, J., Myers, C., & Gluck, M. (1998). A temporally sensitive recurrent network model of occasion setting. In N. Schmajuk, & P. Holland (Eds.), *Occasion setting: Associative learning and cognition in animals* (pp. 319–342). Washington, DC: American Psychological Association.