

available at www.sciencedirect.comwww.elsevier.com/locate/brainres

**BRAIN
RESEARCH**

Research Report

A neural network model of hippocampal–striatal–prefrontal interactions in contextual conditioning

Matthew Turnock, Suzanna Becker*

Department of Psychology, Neuroscience and Behaviour, McMaster University, Building 34, Room 312, 1280 Main Street West, Hamilton, Canada ON L8S 4K1

ARTICLE INFO
Article history:

Accepted 20 June 2007

Available online 26 July 2007

Keywords:

Hippocampus

Context

Nucleus accumbens

Prefrontal cortex

Conditioned behavior

ABSTRACT

The hippocampus is thought to be critical for encoding contextually bound memories and setting the context for ongoing behavior. However, the mechanisms by which the hippocampal–cortical system controls behavior are poorly understood. We propose a computational model in which the hippocampus exerts contextual control over motivated behavior by gating prefrontal cortex inputs to the nucleus accumbens. The model integrates the episodic memory functions of the hippocampus, the prefrontal role in representing the motivational stimuli and cognitive control, and the role of striatal regions in conditioned learning within a single theoretical framework. Simulation results are consistent with the hypothesis that hippocampal–prefrontal interactions may act as the neural substrate that allows contextual cues to override conditioned responses at the level of the nucleus accumbens. Prefrontal and hippocampal input overrides the predominant CS–US association if the context is inconsistent, and promotes flexible selection of previously learned associations and behaviors. Simulated transection of the fornix, effectively eliminating hippocampal and prefrontal influence over the nucleus accumbens, abolishes normal contextual modulation of behavior. The model is consistent with a wide range of empirical data.

© 2007 Elsevier B.V. All rights reserved.

1. Introduction

Context, which may be defined as the complex set of relatively persistent background cues associated with an event, exerts a potent influence over nearly all aspects of learning, perception, memory, and behavior. Everyday examples of contextually modulated behavior include our failure to recognize the familiar face of the postman in an unfamiliar context like riding on a bus, and the reformed smoker's sudden desire for a cigarette in a bar-room context consistent with past smoking

habits. Such phenomena are also well documented empirically. For example, conditioned stimuli whose associations have been extinguished in a novel context can elicit strong immediate reinstatement of conditioned responses when presented in their learned context (Bouton, 2004). This phenomenon is not only governed by the presence or absence of external environmental cues, but also by the internal environment of the animal. For example, Eich (1975) found that students who studied while under the influence of marijuana performed better on their exams while in the

* Corresponding author. Fax: +1 905 529 6225.

E-mail address: becker@mcmaster.ca (S. Becker).URL: <http://www.science.mcmaster.ca/Psychology/sb.html> (S. Becker).

same physical state (under the influence of marijuana) than those who wrote their exams drug-free.

The usage of the term context varies throughout the psychological literature. Cohen and Servan-Schreiber (1992) proposed that a function of the PFC is the active representation of context needed to modulate behavioral processing based on task requirements. In the PFC literature, these representations have been referred to as “cue”, “working memory”, and “goal representations”. In the HPC literature, however, context is used to refer to external environmental cues that are encoded and stored by the HPC in long-term memory. In order to reconcile the two uses of the term context, it is important at this point to distinguish between internal context, subserved by the PFC, and external context, subserved by the HPC (J.D. Cohen, personal communication, April 23, 2007). Throughout this paper, context refers to external environmental cues, while the term ‘motivated behavior’ refers to goal-directed behavior, based on PFC representations of internal context, elicited in response to appetitive or aversive stimuli.

What are the neural mechanisms by which context exerts such a potent influence over motivated behavior? We propose that such contextual control arises from interactions between three major brain regions: (1) the hippocampal complex (HPC), involved in contextual coding and contextual gating; (2) the prefrontal cortex (PFC), involved in selection of appropriate goal-directed behavior; and (3) the nucleus accumbens (NA), involved in the generation of motivated behavior. We consider the separate roles of each of these brain regions briefly in turn, before describing our computational model of the integrated circuit combining all three regions.

The hippocampus is strongly implicated in contextual memory functions including episodic memory (see e.g. Squire, 1992; Tulving and Markowitsch, 1998) and spatial memory and navigation (e.g. Maguire et al., 1997, 1998; Devan et al., 1996; Diamond et al., 1976; Feigenbaum and Rolls, 1991; O’Keefe, 1979; Muller, 1987; Burgess et al., 2002). Given its critical role in the encoding of context, it is not surprising that the hippocampus is also critical for contextually controlled behavior. For example, in aversive conditioning tasks, Phillips and LeDoux (1994) found that HPC-lesioned rats were severely impaired in their ability to learn the context specificity of a conditioned stimulus. Whereas intact rats would freeze in response to the CS only in the conditioned context, HPC-lesioned rats’ freezing responses would generalize to other non-conditioned contexts. Similarly, Selden et al. (1991) reported normal avoidance responses in HPC-lesioned rats when in the presence of a CS that predicts an aversive US, but an impaired ability to select a safe environment based on contextual cues alone. On the other hand, HPC lesions have no effect on other measures of fear conditioning such as the fear-potentiated startle (Gewirtz et al., 2000), nor any effect on the rats’ freezing response to context-only or CS-only states. Rudy and O’Reilly (1999) suggest that hippocampally encoded contextual representations are conjunctive in nature. By comparing the effects of pre-exposure to a conditioned context in intact rats with the effects of pre-exposure to the component features of the context, they only observed conditioned fear responses to the entire context.

One interpretation of these results is that the HPC may be involved in forming stimuli–context and context–context

associations, as opposed to encoding each “contextual” stimulus separately (Winocur, 1997). In the absence of HPC input, contextual cues may be encoded individually by the limbic system or cortex (Rudy et al., 2002). The exact neurobiology of HPC encoding of context as it relates to conditioned behavior is largely unknown. The goal of this paper is to investigate a potential mechanism by which contextual representations of the current environment may directly modulate the activation of conditioned behaviors, via HPC–PFC–NA interactions.

The PFC is strongly implicated in many aspects of goal-oriented and controlled behavior. While the rat prefrontal cortex is proportionately smaller and less developed than that of primates, it is nonetheless a heterogeneous region subserving a wide range of executive functions including working memory (Broersen et al., 1995), representation of the incentive value of stimuli (Schoenbaum et al., 1998, 1999), strategic learning (Winocur and Moscovitch, 1990), attentional set-shifting (Birrell and Brown, 2000), and the ability to inhibit prepotent associations and monitor response conflict (De Wit et al., 2006). Single unit recordings in primates (Badre and Wagner, 2004) and an fMRI study conducted in humans (Matsumoto et al., 2003) also support the PFC’s role in selecting goal-directed behavior. Here we focus particularly on the role of the medial or “pre-limbic” regions of the rat PFC in the representation of incentive value and the control of motivated behavior based on appropriate goal selection.

The NA, a part of the ventral striatum, is strongly implicated in motivated behavior (e.g. Fletcher and Korth, 1999; Schultz, 1998; Ikemoto and Panksepp, 1999), and has been characterized as the limbic-motor interface (Mogenson et al., 1980). It receives input from the amygdala, ventral tegmental area, HPC, and PFC and projects to the ventral pallidum, which in turn projects to the thalamus and PFC. One of the major targets of PFC behavioral control is the ventral striatum; fronto-striatal disconnection disrupts executive control over response selection (Dunnett et al., 2005). Additionally, the rewarding effect of drugs of abuse including cocaine and amphetamines is believed to be mediated by the ventral striatum by elevating dopamine levels and inhibiting GABA neurons (Fletcher and Korth, 1999). Furthermore, rats will actively seek direct electrical stimulation of the NA (Prado-Alcalá and Wise, 1984), while firing rates of NA neurons in awake primates are unrelated to motor activity (reviewed in Pennartz et al., 1994), supporting involvement of the NA in the reward pathway. In addition, NA units in rats trained to self-administer cocaine were active, but unaffected by whether or not the rat actually obtained the drug (Carelli and Ijames, 2000). These results suggest that NA neurons may be more generally related to appetitive and goal directed behavior. Moreover, the majority of NA neuronal responses are not only highly dependent on reward and behavioral task but also on context (Rolls and Williams, 1987).

Pennartz et al. (1994) propose that the conditional nature of NA neuron activation can only be generated if the neuron receives two or more coincident inputs simultaneously.

Consistent with this idea, intracellular recordings by O’Donnell and Grace (1995) revealed three distinct categories of NA neuronal responses: silent; spontaneous, constant, and with a low firing rate; and biphasic (up/down) firing states

mediated by HPC activation that may facilitate activation in groups. Fornix transection, eliminating communication between the hippocampus and the NA, prevented biphasic NA neurons from entering this depolarized active state, which suggests that hippocampal input is necessary for NA neurons to enter an “up” state. PFC input was unable to activate NA neurons unless they were in an “up” state. This indicates a potential dependency on the HPC for PFC activation of the NA (O’Donnell et al., 1999).

In developing a framework for understanding the pathophysiology of schizophrenia, Grace (2000) suggests that information flow from the PFC to the NA is gated by the HPC, and that disruption in this gating mechanism is partly responsible for the etiology of schizophrenia. The heightened sensitivity to novel and aversive stimuli, and prevalence of negative symptoms such as apathy, low motivation, and reduced emotion in schizophrenics might be the result of a lack of inhibition or overriding by the PFC of amygdala and limbic input in the NA, due to the HPC pathology commonly found in schizophrenics. This HPC-mediated modulation of PFC function is implemented computationally in our model to test hypotheses regarding how the HPC uses contextual cues to modulate conditioned behavior at the level of the NA.

The main structures considered by our model are those most directly relevant to appetitive and aversive conditioning: the HPC, amygdala, NA, and PFC. The full range of structures activated during conditioning is far too extensive to model, so we have limited the initial model to these critical structures. The model draws upon several influential ideas in the literature: (1) an appetitive dopaminergic system in conjunction with an aversive opponent system proposed by Daw et al.

(2002), both trained by temporal difference (TD)-learning (Sutton and Barto, 1987), for learning the conditioning tasks as in previous models (Daw and Touretzky, 2000; Suri, 2002; Smith et al., 2004); (2) a simplified version of the hippocampal model designed by O’Reilly and Rudy (2000), which simulates a wide range of empirical data including contextual fear conditioning and learning of conjunctive codes. Our model adopts the fast learning and sparsely coded, conjunctive representations of O’Reilly and Rudy’s (2000) HPC model to encode incoming stimuli and their corresponding context.

In our model, the CS and US send inputs to the amygdala, which projects to the NA to activate its motivational units. The input layer to the HPC consists of amygdala output combined with the CS and context. The HPC output is combined with the CS and US and projects to the PFC, which in turn projects to the NA. Two novel contributions of our model are: (1) the ability of the PFC to override amygdala-driven motivational states at the level of the NA based on the contextual input of the HPC and (2) simulating the mechanism by which HPC contextual representation gates the flow of information from the PFC to the NA.

The model was tested under a number of conditioning paradigms: simple conditioning, extinction, and latent inhibition. In the latent inhibition paradigm, the model was presented with the CS in a context that was either consistent or inconsistent with the context in which the CS was presented without a US. In the final simulation, the model is presented with contextual conditioning in multiple reversing contexts and USs, both with and without a simulated fornix transection that reproduces the one performed by O’Donnell and Grace (1995). Consequently, we show that hippocampal gating of

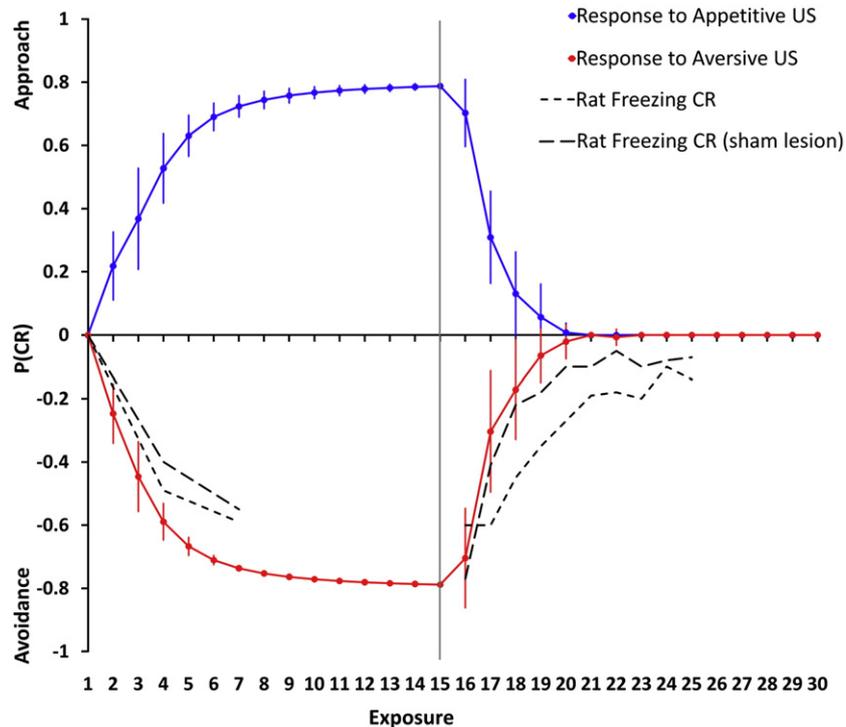


Fig. 1 – Response to extinction. Probability of a conditioned response (\pm SEM) by the model in response to each exposure to the CS and either appetitive (blue) or aversive (red) US. No US is presented past exposure 15. Rat data adapted from Fig. 2 of [Anglada-Figueroa and Quirk \(2005\)](#) for both control rats experiencing sham lesions before conditioning and after conditioning but before extinction.

prefrontal inputs to the nucleus accumbens is a potential mechanism by which context influences motivated behavior.

2. Results

2.1. Simulation 1: simple conditioning and extinction

To establish the validity of the model, we first simulated several simple conditioning tasks. The network was taught a single CS–US association, using either an aversive or appetitive stimulus. Learning curves for both appetitive and aversive conditioning are shown in Fig. 1 (top and bottom curves respectively, trials 1 to 15). Given that the context remains constant across conditioning trials, the response of the model is driven mainly by the emotional input of the amygdala, which learns the CS–US association according to the TD-learning algorithm. For the aversive conditioning case, the learning curve of the model appears to follow the same general trend as rabbit eye blink data obtained by Solomon et al. (1983), as both curves reached a CR probability close to 1 rapidly in response to a CS that predicts an aversive US.

The model was then tested under conditions that produce extinction of a learned CS–US association. Performance curves for the model during the extinction trials are plotted in Fig. 1 (trials 15 onwards), along with experimental data from a fear conditioning and extinction experiment (rat freezing CRs adapted from Fig. 2, Anglada-Figueroa and Quirk, 2005).

The time course of the learning does not correspond exactly with the experimental data but provides a good qualitative fit; the model's probability of a CR to the CS asymptotically approaches 1 as it learns to associate the CS with the appropriate US, but quickly diminishes to zero in the absence of the US.

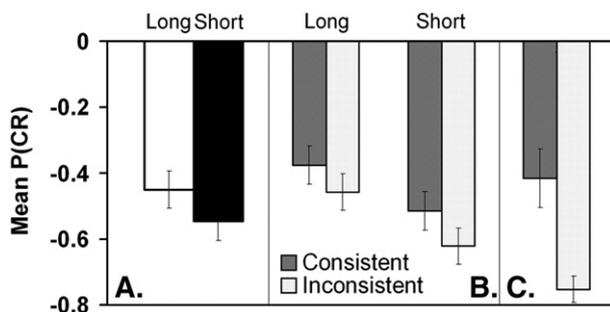


Fig. 2 – Response to latent inhibition. Mean probability of producing an avoidance response (\pm SEM) during the first 15 exposures to the CS–US pairing (A) after exposure to the CS in the absence of a US for 7 trials (short duration, black) and 15 trials (long duration, white), or (B) after exposure to the CS in the absence of a US in a context consistent (dark grey) or inconsistent (light grey) with the pre-exposure context. (C) Rat data adapted from Fig. 1 of Yap and Richardson (2005) for rats with exposure to the CS in a context consistent (dark grey) and inconsistent (light grey) with the pre-exposure context.

2.2. Simulation 2: latent inhibition

The first simulation examining contextual effects involved the latent inhibition paradigm. The original latent inhibition paradigm involves pre-exposure to an unpaired CS prior to exposure to CS–US pairings; the pre-exposure has the effect of retarding acquisition of the CS–US association. In the contextual version of the paradigm, the unpaired and paired trials either occur in the same or different contexts. In our simulations of latent inhibition, in the contextually consistent trials, CS pre-exposure and CS–US pairing both occurred in context A, while in the inconsistent trials, CS pre-exposure occurred in context A, and CS–US pairing occurred in context B. The pre-exposure either occurred for 7 trials (short duration condition) or 15 trials (long pre-exposure condition). A Student's Paired t-test of the difference in mean probability of a CR, averaged over the first 15 exposures to an aversive US, in the short versus long duration pre-exposure conditions revealed a significant effect of the length of latent inhibition [$t(29)=5.96, P=8.88 \times 10^{-7}$] (Fig. 2A). More importantly, the effect of context on latent inhibition was highly significant for both short [$t(14)=9.74, P=6.44 \times 10^{-8}$] and long duration pre-exposures [$t(14)=9.77, P=6.26 \times 10^{-8}$] (Fig. 2B), where pre-exposure to a different context reduced the effect of latent inhibition. This pattern of results is consistent with data obtained by Yap and Richardson (2005) (Fig. 2C) who examined the effects of context and latent inhibition on the freezing response of rats to a CS (odor) that predicted an aversive US (shock).

Consistent with the literature on latent inhibition, the network's probability of making a CR appears to be negatively correlated with the amount of prior exposure to the CS; more pre-exposure leads to slower formation of the CS–US association. On the other hand, there is significantly less latent inhibition when the network is pre-exposed to the CS in a context inconsistent with that in which the CS–US association was learned. In our model, this context-specificity in the CS–US association is a result of the HPC-gating mechanism. While some degree of latent inhibition is still evident when the pre-exposure occurs in a different context than CS–US pairings, its effect appears to be reduced significantly.

The disparity of the magnitude of the difference between means for the simulated data and animal data is likely a result of the arbitrarily chosen learning rate, and the small number of trials selected for the simulation. The averages of two functions that differ in their rate of change drift apart as more trials are added. A simulation run with a smaller learning rate over a greater number of trials might produce results that reflect the animal data more closely; however, the differences are statistically significant and thus provide at least a good qualitative fit to the animal data. Future studies will investigate whether the optimization of the learning rate could produce better quantitative data fits.

2.3. Simulation 3: contextual conditioning with multiple reversals and fornix transection

The final simulations directly explore the impact of the HPC-gating mechanism more extensively by presenting the model with abrupt reversals in the appetitive/aversive valence of the US in concert with contextual changes. In effect, the same CS is

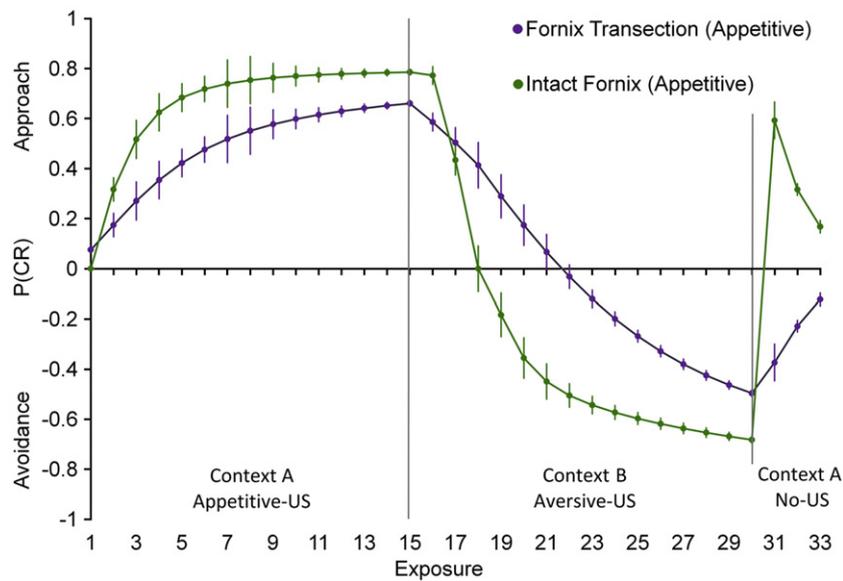


Fig. 3 – Response to modulated CS–US pairing and contextual changes with an intact and transected fornix. Probability of a conditioned response (\pm SEM) by the model in response to each exposure to the CS and either appetitive or aversive US, for both an intact fornix (green) and transected fornix (purple). Changes in context and US occur at exposures 15 and 30.

presented in each context (A or B) in conjunction with context-specific appetitive or aversive stimuli.

CRs during the first 15 exposures resemble the classical conditioning simulation, as the network asymptotically learns to produce correctly an approach response in association with the appetitive US. After the first reversal in US valence, for the next 15 exposures, the network adapts quickly to the change in US, and by the 22nd exposure the model responds correctly with an avoidance CR almost 50% of the time. At the 30th exposure, the network was placed back into the original context, but instead of gradually unlearning the CS–US association in the absence of a US, it responds immediately with an approach response. Although the probability of a CR is less than that at exposure 15, the motivational state is consistent with the CS–appetitive US association learned last time the network was presented with context A.

To simulate a fornix transection (O'Donnell and Grace, 1995), we removed the connections (weights) between the PFC and NA, effectively eliminating contextual modulation of PFC input via the HPC. The exact same parameters as the previous simulation were employed, and the model's responses were analyzed. Fig. 3 plots the responses to various CSs, USs, and contexts for the both the transected fornix and intact fornix for comparison.

Although the fornix-transected model is slower to respond to the change in context and US valence, the motivational state selections are similar for both simulations in response to the US and context changes up until the final 3 exposures. During the final 3 exposures, with an intact fornix, the model responded with an approach response consistent with the behavior exhibited the previous time it was presented with context A; however, when the fornix is transected, an avoidance response was produced despite the change in context from B to A. Another noticeable trend is the model's slow learning curve in the absence of PFC input to the NA. This

effect on the change in probability of producing an approach response is highly significant for both the initial learning [$t(14)=7.89$, $P=8.02 \times 10^{-7}$], and the adaptive learning between exposures 15 and 30 [$t(14)=6.18$, $P=1.18 \times 10^{-5}$]. With the fornix transection, the model did not begin to produce the appropriate response until after seven exposures with the new US at 3.14% probability, whereas with an intact fornix it responded 22.3% of the time by the fourth exposure.

3. Discussion

The model's performance on both classical conditioning and extinction tasks attests to its validity. In both cases, the network produced appropriate responses to the given CS–US pair, consistent with observed animal behaviors and with previous models of classical conditioning based on TD-learning. In these two simulations, the driving force behind its performance lay entirely in the amygdala component's representation of emotional output to the NA, because the context in which conditioning took place remained consistent throughout conditioning and extinction trials. These results are expected given that in our model, learning in the amygdala relies entirely on the TD-learning algorithm, which has been employed previously to simulate conditioned behavior, as discussed earlier.

More importantly, the ability of the PFC to activate NA neurons appears to be critical for the network's ability to associate a context with a CS–US association. This allows the network to select motivational states consistent with the context, as opposed to more recently learned CS–US associations. Disrupting the HPC's ability to allow PFC excitation of the NA, via transection of the fornix, eliminates this phenomenon entirely. Unexpectedly, the speed and extent of learning were compromised after removing the fornix. This suggests that the

PFC may not only be responsible for planning and executing goal-directed behavior, but also for providing feedback to the limbic and HPC systems to facilitate learning. This may suggest one aspect of the neural mechanism responsible for PFC involvement in learning (Dalley et al., 2004), and why forebrain size, cognitive ability, and social complexity in various species may be correlated (Burish et al., 2004; Gagliardo et al., 1996). However, further exploration of this line of research is beyond the scope of this study.

The network integrates two well developed paradigms, TD-Learning and contextual processing by the HPC, to investigate the extensively studied effect that context has on behavior (e.g. Bouton, 1993, 2004; Bouton and Bolles, 1979; Channell and Hall, 1983; Gisquet-Verrier et al., 1999; Hall and Channell, 1986; Hall et al., 1996; Lovibond et al., 1984; Rudy et al., 2004; Winocur, 1997; Yap and Richardson, 2005). These models have been shown to reflect current theories of conditioned behavior (e.g. Daw and Touretzky, 2000; Schultz, 1998; Suri, 2002; Suri et al., 2001; Smith et al., 2004) and HPC memory and function (e.g. Becker, 2005; Frank et al., 2003; O'Reilly and Rudy, 2001). The critical mechanism of HPC-gating in the model is motivated by evidence from intracellular recordings by O'Donnell and Grace (1995) that revealed that stimulation of HPC efferents causes NA neurons to transition into an "up state" that allows them to respond to PFC input. The close relationship between the HPC, amygdala, and NA, as well as the role of the PFC in directing goal-oriented behavior (Carmichael and Price, 1995; Dalley et al., 2004; Goto and Grace, 2005a,b) prompted us to consider their findings as a potential mechanism by which context can affect behavior. By simulating conditioning phenomenon, such as latent inhibition, extinction, and conditioned avoidance, and correlating the network's output with experimentally obtained data, evidence is provided that supports this theoretical mechanism.

Our model is greatly simplified and therefore has a number of limitations. We mention four of these here. First, our model assumed that both aversive and appetitive stimuli are processed in the same amygdala-NA circuits, whereas there is some evidence to suggest two discrete and opposing, yet complementary systems (see Daw et al., 2002). Second, it is important to note that there are multiple reciprocal connections among the structures considered here, the HPC, NA, amygdala, and PFC (Carmichael and Price, 1995), and these structures connect with many other regions including motor cortex (Mogenson et al., 1980), and numerous other systems (Mark et al., 1995), whereas we have only considered a small subset of the connections among these structures. Third, we used a highly abstracted, unrealistic method of simulating the HPC-gating of PFC-NA pathway; the biophysical mechanisms of these state transitions are not well understood but modelers have begun to address these issues (see e.g. Holzman and Tsodyks, 2006). Lastly, the temporal properties of the simulated exposures are largely ignored in our model. Though the number of exposures and learning rate is relatively unimportant and can be selected to fit the time scale of any data set, it has been shown that the inter-stimulus interval (ISI), the time between each exposure to the CS and US, has a critical effect on behavior. The HPC has been implicated in this effect (Solomon et al., 1986), and a TD-learning approach to

this phenomenon has also been explored computationally by Smith et al. (2004).

Despite these limiting assumptions, the model is able to account for several different learning phenomena: classical conditioning, extinction, latent inhibition, US valence reversals, and the context-specificity of all of these phenomena. The model provides the first computational support for a theory that completely relieves the HPC of direct behavioral modulation, and instead places it in the control of a higher executive structure, a notion far more consistent with current beliefs about the function of the HPC and PFC. The model can accommodate multiple CSs, and multiple contexts, including adding internal or external contextual cues from other cortical structures to the EC layer of the HPC. It can also accommodate different US strengths by implementing a range of values between 0 and 1 instead of binary values to contrast, for example, the value of a single food pellet versus a mound of food pellets.

Other hippocampal-cortical models have been developed, in addition to our own, to investigate hippocampal function and classical conditioning, using rather different mechanisms. For example, Schmajuk and DiCarlo (1992) use cortical structures and connections to model far more complex circuits for processing stimuli for HPC input, error signals, and network output. O'Reilly and Rudy (2001) model the HPC as part of a system that forms conjunctive representations of stimulus elements. In their model, both the HPC and cortex share the responsibility of learning: the HPC rapidly learns individual experiences, while the cortex captures general trends over multiple exposures through slow learning. Simulated HPC lesions revealed the importance of using an HPC-formed conjunctive representation of a context and an associated fear response. They show that, in a delayed shock simulation, pre-exposure to the context can provoke a greater fear response than in an immediate shock simulation (an effect that is not observed in the lesioned HPC trials). In our model, this would be explained by the inability of the HPC to use the limited contextual input from the immediate shock condition to form an appropriate contextual-based gating connection to the NA. This would prevent the PFC from exciting the NA and facilitating a fear response. The delayed shock condition would allow the HPC to use the abundance of contextual input to form appropriate NA gating connections that would allow the PFC to provoke a greater fear response when the model is exposed to the context prior to the onset of the CS. This contrasts with their approach by implicating the PFC, a higher-level cortical structure, in directly manipulating behavior, while leaving the HPC to encode the contextual cues and stimulus associations.

Other modelers, in particular, O'Reilly and colleagues, have also addressed dynamic gating mechanisms of the prefrontal cortex. For example, O'Reilly and Frank (2006) modeled the dynamic gating processes involved in controlling the updating of working memory, while Rougier and O'Reilly (2002) modeled the gating actions of the PFC in dynamic task-switching situations.

Although computational models have begun to shed light on the mechanisms underlying conditioned learning, further research will be required to delineate the precise roles of the various brain structures involved. The importance context plays in learning and memory cannot be denied, and further

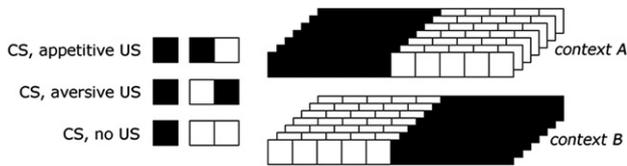


Fig. 4 – Experimentally selected unit activation patterns. Each unit is either on (black) or off (white). (A) Three CS–US association patterns: CS-appetitive US, CS-aversive US, CS-no US. (B) The two contextual input patterns, each with 35 active (black) units. Context A has the first half active; context B has the second half active.

research that investigates the extent of its influence will shed more light on these fundamental elements of behavior.

4. Experimental procedures

The same model was used in all simulations reported here. We first describe the details of the model: the input representation, training patterns, architecture, and operation of each module (PFC, HPC, NA, and amygdala); further details can be found in Appendix A. We then describe the methods for our three simulations: (1) simple conditioning and extinction, (2) latent inhibition, and (3) conditioning with multiple reversals and fornix lesion.

4.1. Input representation and training patterns

Fig. 4 depicts the structure of the input patterns that were presented to the network. Inputs were separated into three groups: a CS unit, two US units, and 70 contextual units. The CS unit indicates presence (on) or absence (off) of the conditioned stimulus, e.g. an auditory stimulus (tone), or visual stimulus (light) (Fig. 4A). Each US unit's activity reflects whether an appetitive or aversive stimulus is present (Fig. 4A); absence of a US is reflected by a lack of any active units in the US pattern. Each input unit projects to 10 EC units with identical levels of activity to facilitate pattern separation and completion by the HPC. Shifts in context are accomplished by assigning a unique activation pattern to the 70 unit context layer. In the simulations reported here, only two contexts are ever used, *context A* or *context B* (Fig. 4B). Novel contexts can be accommodated within the model by creating new patterns; however, these must also be reflected by an increase in the number of PFC units. By repeatedly exposing the network to these patterns with appropriate experimenter-selected input unit activations, many traditional experimental paradigms including conditioning, extinction, and latent inhibition can be simulated along with contextually modulated versions of those experiments.

4.2. Neural network model architecture

Our model incorporated three interacting modules in order to simulate contextually gated behavior: the striatum, the HPC, and the PFC. There are a number of distinct modeling frameworks that have been developed to examine the

functions of each of these areas separately; however, we require a harmonious integration of these modules. To accomplish this, we draw on the established TD-learning algorithm to simulate basic conditioned learning. To process contextual information, we build on a multi-layered HPC model first proposed by Marr (1971), and more recently explored in computational simulations (e.g. Becker, 2005; O'Reilly and Rudy, 2001). Finally, we introduce a rudimentary mechanism that simulates critical prefrontal cortical functionality. Fig. 5 illustrates the combined model integrating the separate structures and the gating mechanism. The main principles are presented in the next 3 sub-sections, and the full details of the model can be found in Appendix A.

4.3. TD-learning and amygdala–nucleus accumbens interactions

The TD model of learning has roots in the earliest computational models of reinforcement learning (Sutton and Barto, 1987). The TD-error signal, which is calculated as the difference between actual and predicted reward, has a striking resemblance to the phasic firing patterns of dopamine (DA) neurons

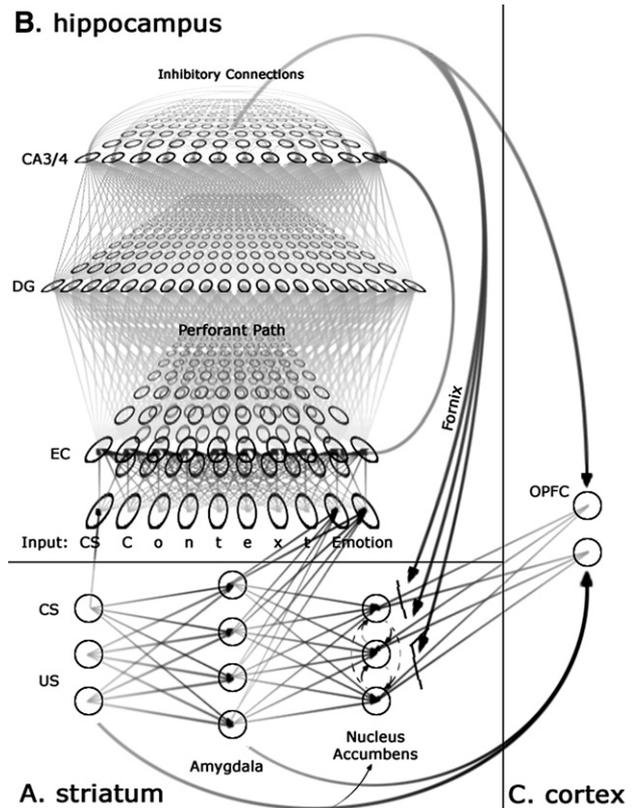


Fig. 5 – Schematic overview of the neural network. The model showing three separate interacting structures: the striatum (amygdala and nucleus accumbens), the hippocampus, and prefrontal cortex. (A) Amygdala layer conveys emotional input to the nucleus accumbens. (B) Three sparsely activated layers of HPC with full connectivity perform pattern separation and completion. (C) Prefrontal units activate in response to US input, amygdala input, and HPC input and output to the appropriate unit of the NA based on the gating properties of the HPC.

in response to rewarding and conditioned stimuli (Schultz, 1998; Schultz et al., 1997). Although aversive and appetitive USs may be handled by separate neural circuits, Daw et al. (2002) incorporate opponent aversive and appetitive processes within a single TD framework by simply representing aversive USs with a negative reward. Similarly, in our model, positive values of reward indicate appetitive USs, negative values indicate aversive USs, and zero reward indicates an absence of US.

Fig. 5A outlines a schematic view of the network's striatum component. In the model, the amygdala receives input representations of individual CSs and USs, associates them together, and outputs the emotional valence of the CS based on past associations with appetitive or aversive USs. Four units in the amygdala layer code emotional valence in a graded fashion, with activation toward either end of the layer signaling a stronger positive or negative emotional valence. Weights between the CS and amygdala are updated via TD-learning. Amygdala output is then sent to the NA via weights that are updated using a Hebbian learning rule based on incoming US signals and amygdala input. In turn, the NA acts as an action biasing layer, combining amygdala input, US input, and lateral inhibition to determine whether to produce an approach or avoidance response, or no response at all. This choice of NA representation is based on evidence that injections of a GABA agonist into various sites in the NA shell elicit either increases in approach (i.e. eating) behaviors or in defensive behaviors depending on the location of the injection (Reynolds and Berridge, 2001), suggesting that the motivational valence coded by the NA is segregated in a graded fashion, a characteristic that has been captured in the model. A more complete model would include a diverse action selection layer, incorporating the basal ganglia and other motor structures, that would receive input from the NA and learn to select the best response based on operant conditioning procedures.

4.4. Hippocampal model and context discrimination

The computational structure of the model is anatomically and physiologically similar to the HPC formation (Fig. 5B) according to Marr (1971). We assume that the overall computational role of the HPC is to perform rapid encoding (either one-shot learning or learning after only a very few exposures) and accurate cued recall. To achieve this, as discussed in the preceding paragraph, the hippocampal circuitry is ideally designed to reach a compromise between pattern separation and pattern completion. Pattern separation is achieved by using a small number of highly selective units that generate sparse activation of the separate layers, which closely reflects activity levels found in the rat HPC (Boss et al., 1985, 1987; Squire et al., 1989). This is achieved by applying a *k*-winners-take-all (*k*WTA) activation function to all layers (O'Reilly and McClelland, 1994). The HPC receives inputs¹ from a wide range

¹ The US does not serve as input to the HPC. The model associates the US with the context and CS indirectly via its emotional valence through the amygdala. The absence of direct US input to the HPC is not critical for the model performance, as the main driving force behind the stimulus–contextual associations that result in modulation of behavior that comes from the amygdala.

of sources within the cortex via the fully connected perforant path from the 100 unit, 25% activation EC network layer to the DG and CA3 layers. In the rat, the DG layer's sparse level of activity (about 0.5% of neurons active; Boss et al., 1985) provides the putative mechanism by which pattern separation is achieved, and is modeled here by a large 500 unit sparsely coded layer with 1% activation. The CA3 layer receives mossy fiber pathway input from the DG and perforant pathway input from the EC. Pattern completion is achieved within the 150 unit CA3 layer with 10% activation by incorporating widespread lateral (recurrent) connections, such that partial activation of a pattern can invoke activation of missing units (Marr, 1971). Because we are simulating relatively simple conditioned responses involving sub-cortical structures, it was unnecessary to include the CA1 layer and backward connections to the EC. The EC back-projection allows the HPC to influence many regions in the neocortex, and these connections would be important to include in simulations of more complex tasks.

4.5. Prefrontal cortex modulation at the nucleus accumbens

Exactly how the HPC uses contextual cues to modulate behavior is unclear. As mentioned in the Introduction, O'Donnell and Grace (1995) found that biphasic NA neurons could only be driven into the "up" state by stimulation of HPC afferents (and not PFC, amygdala, or thalamic stimulation), and could only be driven to fire by PFC inputs when in the up state. The strong relationship between the NA, dopamine-modulated limbic systems, and goal-directed behavior (Mogenson et al., 1980; Goto and Grace, 2005a,b) provides a theoretical framework for a potential mechanism by which contextual information processed by the HPC can influence conditioned behavior. Our implementation of this HPC modulation of PFC influence over the NA is the major novel contribution of our model, and allows us to simulate tasks not previously dealt with by either reinforcement learning or hippocampal models alone.

We employ a highly simplified neural circuit to represent the PFC; for more elaborate and realistic models of PFC function, see e.g. O'Reilly et al. (2002) and Rougier and O'Reilly (2002). The output from the PFC is computed as a function of contextual and emotional input received from the HPC and amygdala. The PFC is represented by two units that respond individually to the two simulated contexts and emotional input (Fig. 5C), with output connections to each of the NA units (representing approach and avoidance) that strengthen or weaken extremely rapidly based on a Hebbian plasticity rule, capturing co-occurring PFC contextual representations and NA motivational drive representations. Since weight updates between the NA and PFC are contingent upon HPC input to the PFC, the connections between the PFC and NA act as the target of the HPC-gating mechanism.² Input from the PFC to the appropriate motivational unit in the NA is much stronger than the emotional input to the NA,

² Note that in our model there are no direct HPC connections to the NA. HPC activity affects the NA only through its modulatory action on the PFC. If the PFC layer was removed, the model would perform similarly to a model with a fornix transection.

and through the lateral inhibitory connections within the simulated NA, a motivational unit that would otherwise have been excited by amygdala input may be inhibited by the stronger PFC input. Without any PFC input, or limited input, the amygdala remains as the only driving force behind NA activity, and any unit activations are selected based purely on emotional valence, without any contextual influence.

4.6. Performance measure

The performance of the model was measured as the probability of producing a conditioned response, $P(\text{CR})$, for each exposure to the CS. Positive values indicate approach responses. Probabilities are negated to indicate avoidance responses. The time course of learning is largely irrelevant because the learning rate is an arbitrarily chosen number, and can easily be modified to fit the time-scale of any data set (refer to Appendix A for model details).

4.7. Methods for simulation 1: simple conditioning and extinction

In simulation 1, the model's input consisted of a CS, an aversive or appetitive US, and a single context for 15 exposures, after which extinction trials began. During extinction, the model was exposed to the CS and context in the absence of a US for 15 trials. Each set of exposures was run 10 times to simulate results obtained from 10 subjects.

4.8. Methods for simulation 2: latent inhibition

The effects of latent inhibition on the model were examined by pre-exposing the network to the CS for various time durations and contexts. Short and long durations of latent inhibition were simulated by employing 7 and 15 pre-exposures of the CS, prior to conditioning trials. Furthermore, the model experienced the CS in a context that was either consistent or inconsistent with the context presented during subsequent conditioning trials. During conditioning trials, the CS was paired with either an appetitive or aversive US for the remaining exposures. Each set of exposures was run 10 times to simulate results obtained from 10 subjects.

4.9. Methods for simulation 3: contextual conditioning with multiple reversals and fornix transection

The final simulations explored the specific effect of the gating mechanism by exposing the model to a single CS in varying contexts with varying USs. For the first 15 trials, the model was exposed to the CS with an appetitive or aversive US in context A, then for the next 15 trials it was exposed to the same CS but with a US of opposite valence in context B, and for the final 3 trials, the model was exposed to the CS in the original context A without any US to avoid influencing motivational selection. The simulation was run with the network as presented earlier, intact, and with the HPC contribution to the connections from the PFC to the NA removed in order to simulate a fornix transection. Each set of

exposures was run 10 times to simulate results obtained from 10 subjects.

Appendix A

A.1. Neuron architecture

Neurons are interconnected via weighted connections. Except where otherwise noted (see below), a neuron's output is calculated using a binary threshold function: if the weighted summed input exceeds the threshold, θ , its output is 1, otherwise its output is 0. All weights are initialized to random values between -1 and 1 . The weight values indicate the strength of the synaptic connections, and are updated according to the learning rule appropriate to that unit in the network, as described below.

A.2. Learning algorithms

Weights are updated either using Hebbian learning or TD-learning. All weights between layers in HPC, and connections from the CS, amygdala, PFC to the NA are updated using Hebbian learning, while TD-learning occurs only between the CS and amygdala units. Both learning algorithms will be discussed in the relevant subsections.

A.3. Hippocampus structure

The hippocampus is represented by three layers: the entorhinal cortex (EC), dentate gyrus (DG), and CA3. The EC is a 100 unit layer with 25% of units active, and full connectivity to the DG and CA3 layers. The DG is a 500 unit layer with 1% of units active, and full connectivity to the CA3 layer. The CA3 is a 150 unit layer with 10% of units active, full recurrent connectivity, and full connectivity to the PFC. Sparse activation is achieved through the use of a k -winners-take-all function (k WTA), which limits the number of active units by applying a global inhibitory value (Eq. (1)) to all units in a layer, and increasing inhibition until all but k of the units' activations are below threshold:

$$\Delta x_i = x_k + q(x_{k+1} - x_k) - \theta \quad (1)$$

where x_{k+1} is the output value of the last unit to be activated, x_k is the output value of the first inhibited unit, x_i is the inhibitory adjustment to the activity of all units, q is a parameter that ensures the k th unit is always above threshold (usually set to 0.25), and θ is the threshold value. All weights between the hippocampal layers are updated by applying Hebbian learning. Hebbian learning occurs when an input unit and the receiving unit's outputs, x_i and x_j have both reached threshold:

$$\Delta W_{ij} = \alpha * x_i x_j (1 - W_{ij}) \quad (2)$$

where α is the learning rate, and W_{ij} is the value of the weight on the connection between the i th and j th units.

A.4. Ventral striatum and amygdala

The amygdala consists of a 4 unit layer representing emotional valences of $\gamma = -0.75, -0.25, +0.25,$ and $+0.75$. Thus, the job of predicting the reward is distributed across

the four units. Weights between the CS and amygdala units are updated via a modified TD-learning rule, where the TD-error is calculated as the difference between the actual “reward” value associated with the US (positive for reward, negative for punishment) and the expected emotional reward value (sum of the outputs of each amygdala unit weighted by its target emotional valence, γ):

$$TD = r(US) - \Sigma(\gamma_j)(x_j)$$

$$\Delta W_{ij} = \alpha * x_i * x_j (1 - W_{ij})$$

$$= TD * \gamma_j * x_i \quad (3)$$

The NA is a 3 unit layer, with each unit representing one of three motivational states: approach, neutral, and avoidance. NA units are probabilistically activated, with the probability of activating a particular NA unit calculated using the softmax function (Bridle, 1990). Weights are updated using a simplified version of Q-learning (Watkins, 1989), where the value of the US acts as a reward (+1) or punishment (-1), the particular NA unit probabilistically chosen to be active is treated as the “action”, and that unit’s weights undergo positive or negative Hebbian learning depending on whether the “action” is consistent with the value of the US. PFC input increases the probability of activation of one of the units, making it more likely to be activated over another unit that might have been activated based on amygdala and US input alone. If there is no US, only the neutral unit is updated, with a reward value of +1 for that unit. This reflects the reduction of CS saliency over time (Ungless, 2004) and helps facilitate extinction and latent inhibition.

A.5. Prefrontal cortex and HPC-gating action

The prefrontal cortex (PFC) layer contains 2 units that activate according to a WTA function similar to Eq. (1) (but with only 1 winner rather than k WTA). PFC activation is determined by HPC, amygdala, and US input. The weights between the PFC and NA are updated via Hebbian learning (Eq. (2)), so that combinations of context and US are associated with the appropriate motivational state for that context. The gating action is established by: (1) eliminating PFC output to the NA when there is no HPC contextual input, leaving the activation of the NA dependant solely on amygdala input, and (2) normalizing the incoming weights to the NA so that the PFC inputs are always much stronger and override the amygdala inputs when present. The simulated transection of the fornix was achieved by removing HPC input to the PFC, thus eliminating context-dependent NA input from the PFC.

REFERENCES

- Anglada-Figueroa, D., Quirk, G.J., 2005. Lesions of the basal amygdala block expression of conditioned fear but not extinction. *J. Neurosci.* 25, 9680–9685.
- Badre, D., Wagner, A.D., 2004. Selection, integration, and conflict monitoring: assessing the nature and generality of prefrontal cognitive control mechanisms. *Neuron* 41, 473–487.
- Becker, S., 2005. A computational principle for hippocampal learning and neurogenesis. *Hippocampus* 15, 722–738.
- Birrell, J.M., Brown, V.J., 2000. Medial frontal cortex mediates perceptual attentional set shifting in the rat. *J. Neurosci.* 20, 4320–4324.
- Boss, B.D., Peterson, G.M., Cowan, W.M., 1985. On the numbers of neurons in the dentate gyrus of the rat. *Brain Res.* 338, 144–150.
- Boss, B.D., Turlejski, K., Stanfield, B.B., Cowen, W.M., 1987. On the numbers of neurons in fields CA1 and CA3 of the hippocampus of Sprague–Dawley and Wistar rats. *Brain Res.* 406, 280–287.
- Bouton, M.E., 1993. Context, time, and memory retrieval in the interference paradigms of pavlovian learning. *Psychol. Bull.* 114, 80–99.
- Bouton, M.E., 2004. Context and behavioral processes in extinction. *Learn. Mem.* 11, 485–494.
- Bouton, M.E., Bolles, R.C., 1979. Role of conditioned contextual stimuli in reinstatement of extinguished fear. *J. Exp. Psychol., Anim. Behav. Processes* 5, 368–378.
- Bridle, J.S., 1990. Training stochastic model recognition algorithms as networks can lead to maximum mutual information estimation of parameters. In: Touretzky, D.S. (Ed.), *Advances in Neural Information Processing Systems*, vol. 2. Morgan Kaufmann, San Mateo, pp. 211–217.
- Broersen, L.M., Heinsbroek, R.P.W., Debruin, J.P.C., Uylings, H.B.M., Olivier, B., 1995. The role of the medial prefrontal cortex of rats in short-term-memory functioning—further support for involvement of cholinergic, rather than dopaminergic mechanisms. *Brain Res.* 674, 221–229.
- Burgess, N., Maguire, E.A., O’Keefe, J., 2002. The human hippocampus and spatial and episodic memory. *Neuron* 35, 625–641.
- Burish, M.J., Kueh, H.Y., Wang, S.S.H., 2004. Brain architecture and social complexity in modern and ancient birds. *Brain Behav. Evol.* 63, 107–124.
- Carelli, R.M., Ijames, S.G., 2000. Nucleus accumbens cell firing during maintenance, extinction, and reinstatement of cocaine self-administration behavior in rats. *Brain Res.* 866, 44–54.
- Carmichael, S.T., Price, J.L., 1995. Limbic connections of the orbital and medial prefrontal cortex in macaque monkeys. *J. Comp. Neurol.* 363, 615–641.
- Channell, S., Hall, G., 1983. Contextual effects in latent inhibition with an appetitive conditioning procedure. *Anim. Learn. Behav.* 11, 67–74.
- Cohen, J., Servan-Schreiber, D., 1992. Context, cortex, and dopamine: a connectionist approach to behavior and biology in schizophrenia. *Psychol. Rev.* 99, 45–77.
- Dalley, J.W., Cardinal, R.N., Robbins, T.W., 2004. Prefrontal executive and cognitive functions in rodents: neural and neurochemical substrates. *Neurosci. Biobehav. Rev.* 28, 771.
- Daw, N.D., Touretzky, D.S., 2000. Behavioral considerations suggest an average reward TD model of the dopamine system. *Neurocomputing* 32, 679–684.
- Daw, N.D., Kakade, S., Dayan, P., 2002. Opponent interactions between serotonin and dopamine. *Neural Netw.* 15, 603–616.
- Devan, B.D., Goad, E.H., Petri, H.L., 1996. Dissociation of hippocampal and striatal contributions to spatial navigation in the water maze. *Neurobiol. Learn. Mem.* 66, 305–323.
- De Wit, S., Kosaki, Y., Balleine, B.W., Dickinson, A., 2006. Dorsomedial prefrontal cortex resolves response conflict in rats. *J. Neurosci.* 26, 5224–5229.
- Diamond, M.C., Ingham, C.A., Johnson, R.E., Bennett, E.L., Rosenzweig, M.R., 1976. Effects of environment on morphology of rat cerebral-cortex and hippocampus. *J. Neurobiol.* 7, 75–85.
- Dunnett, S.B., Meldrum, A., Muir, J.L., 2005. Frontal-striatal disconnection disrupts cognitive performance of the frontal-type in the rat. *Neuroscience* 135, 1055–1065.

- Eich, J.E., 1975. State-dependent accessibility of retrieval cues in retention of a categorized list. *J. Verbal Learn. Verbal Behav.* 14, 408–417.
- Feigenbaum, J.D., Rolls, E.T., 1991. Allocentric and egocentric spatial information-processing in the hippocampal-formation of the behaving primate. *Psychobiology* 19, 21–40.
- Fletcher, P.J., Korth, K.M., 1999. Activation of 5-HT_{1B} receptors in the nucleus accumbens reduces amphetamine induced enhancement of responding for conditioned reward. *Psychopharmacology* 142, 165–174.
- Frank, M.J., Rudy, J.W., O'Reilly, R.C., 2003. Transitivity, Flexibility, Conjunctive Representations and the Hippocampus: II. A Computational Analysis. *Hippocampus* 13, 341–354.
- Gagliardo, A., Bonadonna, F., Divac, I., 1996. Behavioural effects of ablations of the presumed 'prefrontal cortex' or the corticoid in pigeons. *Behav. Brain Res.* 78, 155–162.
- Gewirtz, J.C., McNish, K.A., Davis, M., 2000. Is the hippocampus necessary for contextual fear conditioning? *Behav. Brain Res.* 110, 83–95.
- Gisquet-Verrier, P., Dutriex, G., Richer, P., Doyere, V., 1999. Effects of lesions to the hippocampus on contextual fear: evidence for a disruption of freezing and avoidance behaviour but not context conditioning. *Behav. Neurosci.* 113, 507–522.
- Goto, Y., Grace, A.A., 2005a. Dopaminergic modulation of limbic and cortical drive of nucleus accumbens in goal-directed behavior. *Nat. Neurosci.* 8, 805–812.
- Goto, Y., Grace, A.A., 2005b. Dopamine-dependent interactions between limbic and prefrontal cortical plasticity in the nucleus accumbens: disruption by cocaine sensitization. *Neuron* 47, 255–266.
- Grace, A.A., 2000. Gating of information flow within the limbic system and the pathophysiology of schizophrenia. *Brains Res. Rev.* 31, 330–341.
- Hall, G., Channell, S., 1986. Context specificity of latent inhibition in taste-aversion learning. *Q. J. Exp. Psychol., B Comp. Physiol. Psychol.* 38, 121–139.
- Hall, G., Purves, D., Bonardi, C., 1996. Contextual control of conditioned responding in rats with dorsal hippocampal lesions. *Behav. Neurosci.* 110, 933–945.
- Holcman, D., Tsodyks, M., 2006. The emergence of up and down states in cortical networks. *PLOS Comput. Biol.* 2, 174–181.
- Ikemoto, S., Panksepp, J., 1999. The role of nucleus accumbens dopamine in motivated behavior: a unifying interpretation with special reference to reward-seeking. *Brains Res. Rev.* 31, 6–41.
- Lovibond, P.F., Preston, G.C., Mackintosh, N.J., 1984. Context specificity of conditioning, extinction, and latent inhibition. *J. Exp. Psychol., Anim. Behav. Processes* 10, 360–375.
- Maguire, E.A., Frackowiak, R.S.J., Frith, C.D., 1997. Recalling routes around London: activation of the right hippocampus in taxi drivers. *J. Neurosci.* 17, 7103–7110.
- Maguire, E.A., Burgess, N., Donnett, J.G., Frackowiak, R.S.J., Frith, C.D., O'Keefe, J., 1998. Knowing where and getting there: a human navigation network. *Science* 280, 921–924.
- Mark, L.P., Daniels, D.L., Naidich, T.P., Hendrix, L.E., 1995. Limbic connections. *Am. J. Neuroradiol.* 16, 1303–1306.
- Marr, D., 1971. Simple memory: a theory for archicortex. *Philos. trans. R. Soc. Lond., B* 262, 23–81.
- Matsumoto, K., Suzuki, W., Tanaka, K., 2003. Neuronal correlates of goal-based motor selection in the prefrontal cortex. *Science* 301, 229–232.
- Mogenson, G.J., Jones, D.L., Yim, C.Y., 1980. From motivation to action: functional interface between limbic system and the motor system. *Prog. Neurobiol.* 14, 69–97.
- Muller, R.U., 1987. The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. *J. Neurosci.* 7, 1951–1968.
- O'Donnell, P., Grace, A.A., 1995. Synaptic-interactions among excitatory afferents to nucleus-accumbens neurons—hippocampal gating of prefrontal cortical input. *J. Neurosci.* 15, 3622–3639.
- O'Donnell, P., Greene, J., Pabello, N., Lewis, B.L., Grace, A.A., 1999. Modulation of cell firing in the nucleus accumbens. *Ann. N.Y. Acad. Sci.* 877, 157–175.
- O'Keefe, J., 1979. Review of the hippocampal place cells. *Prog. Neurobiol.* 13, 419–439.
- O'Reilly, R.C., Frank, M.J., 2006. Making working memory work: a computational model of learning in the prefrontal cortex and basal ganglia. *Neural Comput.* 18, 283–328.
- O'Reilly, R.C., McClelland, J.L., 1994. Hippocampal conjunctive encoding, storage, and recall: avoiding a tradeoff. *Hippocampus* 4, 661–682.
- O'Reilly, R.C., Rudy, J.W., 2000. Computational principles of learning in the neocortex and hippocampus. *Hippocampus* 10, 389–397.
- O'Reilly, R.C., Rudy, J.W., 2001. Conjunctive representations in learning and memory: principles of cortical and hippocampal function. *Psychol. Rev.* 108, 311–345.
- O'Reilly, R.C., Noelle, D.C., Braver, T.S., Cohen, J.D., 2002. Prefrontal cortex and dynamic categorization tasks: representational organization and neuromodulatory control. *Cereb. Cortex* 12, 246–257.
- Pennartz, C.M.A., Groenewegen, H.J., Lopes da Silva, F.H., 1994. The nucleus-accumbens as a complex of functionally distinct neuronal ensembles—an integration of behavioral, electrophysiological and anatomical data. *Prog. Neurobiol.* 42, 719–761.
- Phillips, R.G., LeDoux, J.E., 1994. Lesions of the dorsal hippocampal-formation interfere with background but not foreground contextual fear conditioning. *Learn. Memory* 1, 34–44.
- Prado-Alcalá, R., Wise, R.A., 1984. Brain-stimulation reward and dopamine terminal fields: 1. Caudate-putamen, nucleus accumbens and amygdala. *Brain Res.* 297, 265–273.
- Reynolds, S.M., Berridge, K.C., 2001. Fear and feeding in the nucleus accumbens shell: rostrocaudal segregation of GABA-elicited defensive behaviour versus eating behaviour. *J. Neurosci.* 21, 3261–3270.
- Rolls, E.T., Williams, G.V., 1987. Neuronal activity in the ventral striatum of the primate. In: Carpenter, M.B., Jayaraman, A. (Eds.), *The basal ganglia II- structure and function - current concepts*. Plenum, New York, pp. 349–356.
- Rougier, N.P., O'Reilly, R.C., 2002. Learning representations in a gated prefrontal cortex model of dynamic task switching. *Cogn. Sci.* 26, 503–520.
- Rudy, J.W., O'Reilly, R.C., 1999. Contextual fear conditioning, conjunctive representations, pattern completion, and the hippocampus. *Behav. Neurosci.* 113, 867–880.
- Rudy, J.W., Barrientos, R.M., O'Reilly, R.C., 2002. Hippocampal formation supports conditioning to memory of a context. *Behav. Neurosci.* 116, 530–538.
- Rudy, J.W., Huff, N.C., Matus-Amat, P., 2004. Understanding contextual fear conditioning: insights from a two-process model. *Neurosci. Biobehav. Rev.* 28, 675–685.
- Schmajuk, N.A., DiCarlo, J.J., 1992. Stimulus configuration, classical-conditioning, and hippocampal function. *Psychol. Rev.* 99, 268–305.
- Schoenbaum, G., Chiba, A.A., Gallagher, M., 1998. Orbitofrontal cortex and basolateral amygdala encode expected outcomes during learning. *Nat. Neurosci.* 1, 155–159.
- Schoenbaum, G., Chiba, A.A., Gallagher, M., 1999. Neural encoding in orbitofrontal cortex and basolateral amygdala during olfactory discrimination learning. *J. Neurosci.* 19, 1876–1884.
- Schultz, W., 1998. Predictive reward signal of dopamine neurons. *J. Neurophysiol.* 80, 1–27.

- Schultz, W., Dayan, P., Montague, P.R., 1997. A neural substrate of prediction and reward. *Science* 275, 1593–1599.
- Selden, N.R.W., Everitt, B.J., Jarrard, L.E., Robbins, T.W., 1991. Complementary roles for the amygdala and hippocampus in aversive-conditioning to explicit and contextual cues. *Neuroscience* 42, 335–350.
- Smith, A., Li, M., Becker, S., Kapur, S., 2004. A model of antipsychotic action in conditioned avoidance: a computational approach. *Neuropsychopharmacology* 29, 1040–1049.
- Solomon, P.R., Solomon, S.D., Vander Schaaf, E., Perry, H.E., 1983. Altered activity in the hippocampus is more detrimental to classical-conditioning than removing the structure. *Science* 220, 329–331.
- Solomon, P.R., Vander Schaaf, E.R., Thompson, R.F., Weisz, D.J., 1986. Hippocampus and trace conditioning of the rabbit's classically conditioned nictitating membrane response. *Behav. Neurosci.* 100, 729–744.
- Squire, L.R., 1992. Memory and the hippocampus—a synthesis from findings with rats, monkeys, and humans. *Psychol. Rev.* 99, 195–231.
- Squire, L.R., Shimamura, A.P., Amaral, D.G., 1989. Memory and the hippocampus. In: Byrne, J.H., Berry, W.O. (Eds.), *Neural Models of Plasticity: Experimental and Theoretical Approaches*. Academic Press, San Diego, pp. 208–239.
- Suri, R., 2002. TD models of reward predictive responses in dopamine neurons. *Neural Netw.* 15, 523–533.
- Suri, R.E., Bargas, J., Arbib, M.A., 2001. Modeling functions of striatal dopamine modulation in learning and planning. *Neuroscience* 103, 65–85.
- Sutton, R.S., Barto, A.G., 1987. A temporal-difference model of classical conditioning. *Proceedings of the Ninth Conference of the Cognitive Science Society*, pp. 355–378.
- Tulving, E., Markowitsch, H.J., 1998. Episodic and declarative memory: role of the hippocampus. *Hippocampus* 8, 198–204.
- Ungless, M.A., 2004. Dopamine: the salient issue. *Trends Neurosci.* 27, 702–706.
- Watkins, C., 1989. Learning from delayed rewards. PhD thesis, University of Cambridge, Psychology Department.
- Winocur, G., 1997. Hippocampal lesions alter conditioning to conditional and contextual stimuli. *Behav. Brain Res.* 88, 219–299.
- Winocur, G., Moscovitch, M., 1990. Hippocampal and prefrontal cortex contributions to learning and memory: analysis of lesion and aging effects on maze learning in rats. *Behav. Neurosci.* 104, 544–551.
- Yap, C.S.L., Richardson, R., 2005. Latent inhibition in the developing rat: an examination of context-specific effects. *Dev. Psychobiol.* 47, 55–65.