Review

Basal ganglia and dopamine contributions to probabilistic category learning

D. Shohamy*, C.E. Myers, J. Kalanithi, M.A. Gluck

Department of Psychology, 386 Schermerhorn Hall, Columbia University, New York, NY 10027, USA

Abstract

Studies of the medial temporal lobe and basal ganglia memory systems have recently been extended towards understanding the neural systems contributing to category learning. The basal ganglia, in particular, have been linked to probabilistic category learning in humans. A separate parallel literature in systems neuroscience has emerged, indicating a role for the basal ganglia and related dopamine inputs in reward prediction and feedback processing. Here, we review behavioral, neuropsychological, functional neuroimaging, and computational studies of basal ganglia and dopamine contributions to learning in humans. Collectively, these studies implicate the basal ganglia in incremental, feedback-based learning that involves integrating information across multiple experiences. The medial temporal lobes, by contrast, contribute to rapid encoding of relations between stimuli and support flexible generalization of learning to novel contexts and stimuli. By breaking down our understanding of the cognitive and neural mechanisms contributing to different aspects of learning, recent studies are providing insight into how, and when, these different processes support learning, how they may interact with each other, and the consequence of different forms of learning for the representation of knowledge.

© 2007 Elsevier Ltd. All rights reserved.

Keywords: Memory; Hippocampus; Reinforcement learning; Categorization

Contents

1. Introduction	2
2. The basal ganglia and learning	2
2.1. Probabilistic category learning and the basal ganglia	2
2.2. Memory systems and the basal ganglia: limitations and open questions	3
3. The basal ganglia and outcome prediction	4
3.1. How do people learn? Strategies in probabilistic category learning	4
3.1.1. Computational modeling of the basal ganglia in probabilistic learning	5
3.2. When do people learn? Temporal dynamics of the basal ganglia during learning	7
3.3. What do people learn? Basal ganglia and representation of learned knowledge	9
4. Feedback and reward modulate learning	11
4.1. The role of feedback in incremental learning	11
4.2. Dopamine modulation of feedback-based learning	13
5. Summary	15
Acknowledgments	15
References	15

*Corresponding author.

E-mail address: shohamy@psych.columbia.edu (D. Shohamy).

1. Introduction

In everyday life, decisions and actions are often guided by the ability to classify events and objects into distinct categories. In many cases, categorization is based on a specific memory deriving from a single past experience. Other times, categorization may not be based on a specific memory, but instead may follow a “gut feeling” that is based on a gradual accumulation of experiences over time. Decades of research into the neural bases of learning and memory suggest that these different forms of memory are supported by distinct cognitive and neural systems. The hippocampus and medial temporal lobes (MTL) support explicit memories for events or episodes, often referred to as ‘declarative’ memory (Cohen and Eichenbaum, 1993; Eichenbaum and Cohen, 2001; Schacter and Wagner, 1999; Squire, 1987, 1992). The basal ganglia are thought to support a distinct and independent system that contributes to gradual learning of stimulus–response associations over many trials—a form of non-declarative memory often referred to as ‘procedural’ or ‘habit’ learning (Gabrieli, 1998; Knowlton et al., 1996; Robbins, 1996; White, 1997).

In recent years, studies of the MTL and basal ganglia memory systems have been extended to understand the neural systems contributing to category learning. The basal ganglia, in particular, have been linked to probabilistic category learning in humans. A separate parallel literature in systems neuroscience has emerged, indicating a role for the basal ganglia and related dopamine inputs in reward prediction and feedback processing (Fiorillo et al., 2003; Hollerman and Schultz, 1998; Schultz, 1998; Schultz et al., 1997).

In this paper, we review behavioral, neuropsychological, functional neuroimaging, and computational studies of basal ganglia contributions to learning in humans. We first review early neuropsychological studies that provided the initial link between the basal ganglia and probabilistic category learning, implicating the basal ganglia in non-declarative habit learning. We then turn to recent neuropsychological and neuroimaging studies that break away from the declarative/non-declarative distinction to understand more specifically how the basal ganglia contribute to different aspects of category learning, specifically *how*, *when*, and *what* people learn during categorization. Finally, we review how dopamine and feedback modulate learning, drawing on recent pharmacological studies in healthy individuals and those with disrupted basal ganglia function.

2. The basal ganglia and learning

The basal ganglia are a group of highly interconnected subcortical nuclei. The main input structure of the basal ganglia is the striatum (caudate and putamen), which receives widespread projections from cortex and serves as the primary source of basal ganglia input (Alexander et al., 1986). The striatum also receives input from dopamine

projections from the Substantia nigra compacta (SNc), which modulate cortico-striatal plasticity (Albin et al., 1989; Calabresi et al., 1992; Cepeda et al., 1993; Wickens et al., 1996). Output from the basal ganglia projects back, via thalamus, to many of the same areas from which they receive input (Alexander et al., 1986). Thus, overall, the basal ganglia can be viewed as an interface between cortex and thalamus, integrating cortical information and mapping it onto behavior (Alexander et al., 1986).

Preliminary data suggesting the basal ganglia contribute not only to motor function, but are important for learning, came from animal and patient studies demonstrating a dissociation in the pattern of memory impairments following damage to the MTL and damage to the basal ganglia. Basal ganglia damage was found to impair performance on a variety of incremental, stimulus–response learning tasks (Downes et al., 1989; Kesner et al., 1993; Knowlton et al., 1996; McDonald and White, 1993; Owen et al., 1993a; Packard, 1999; Packard et al., 1989; Packard and McGaugh, 1996; Saint-Cyr et al., 1988; Shohamy et al., 2005, 2006, 2004a,b; Swanson et al., 2000), but spared performance on tasks that involve declarative memory (Knowlton et al., 1996). The opposite pattern was observed in individuals with damage to the MTL: striking declarative memory deficits, but spared incremental learning of stimulus–response associations (Gabrieli, 1998; Knowlton et al., 1996). These findings are among those supporting the idea that there are different forms of memory that are subserved by different systems, an idea that has been prominent in cognitive and neural sciences for decades (Eichenbaum and Cohen, 2001; Gabrieli, 1998).

2.1. Probabilistic category learning and the basal ganglia

In humans, particularly strong evidence for basal ganglia contributions to learning comes from neuropsychological and neuroimaging studies of *probabilistic category learning* (Gluck and Bower, 1988; Knowlton et al., 1994, 1996; Poldrack et al., 1999, 2001; Shohamy et al., 2004a,b). One widely explored paradigm is known as the “weather prediction” task, developed by Gluck and colleagues at Rutgers University based on an early task by Gluck and Bower (1988). In this category learning task, subjects view one or more cards with different geometric shapes on each trial, are asked to predict a category outcome (“rain” or “sunshine”), and receive feedback on their decision. There are four cards, and the actual weather outcome is differentially associated with each card with a particular probability. For example, the triangle card might usually (but not always) predict rain, while the circle card might usually (but not always) predict sun. Sample stimuli and probabilities for each of the cards are shown in Fig. 1.

In a seminal paper, Knowlton and colleagues demonstrated that this sort of probabilistic classification depends on the basal ganglia (Knowlton et al., 1996), and not the MTL. Knowlton and colleagues reasoned that because of

Fig. 1. A popular probabilistic classification task is often referred to as the “weather prediction” task. (A) Each of 4 visual cues—cards with shapes—is independently and probabilistically associated with either “rain” or “sun”. (B) On each trial, a combination of one to three cards is shown. Subjects respond based on their prediction of the weather for that trial, and receive response-contingent feedback.

the probabilistic nature of the associations, declarative memory for any single trial event cannot support learning, so learning must depend on non-declarative, associative learning processes over many trials. Knowlton and colleagues examined learning among two patient groups: a group of amnesic patients with severe memory impairments (due to either MTL or diencephalic damage) and a group of patients with disrupted basal ganglia function due to moderate to advanced Parkinson’s disease (Knowlton et al., 1996, 1994). The Parkinson’s patients were impaired at learning the task, an impairment that was particularly pronounced in patients in an advanced stage of the disease (Knowlton et al., 1996). The amnesic patients performed as well as healthy controls early in the task (over the first 50 trials), but were impaired later in learning, as training progressed (Knowlton et al., 1996). After training on the task, both patient groups were tested for declarative memory of the experiment. Here, the two patient groups showed the reverse pattern: the Parkinson’s patients were able to recall details of the stimuli and task events, while amnesic patients were able to recall few if any details. This double dissociation suggested a dissociation between basal ganglia and limbic system contributions to different forms of memory: the basal ganglia were necessary for incremental, stimulus-response learning, while the early spared performance among the amnesic patients suggested that the limbic system is not, at least in the earliest phases.

The Knowlton et al. (1996) study of probabilistic classification impacted the field in two important ways. First, in demonstrating a double dissociation between the Parkinson’s patients and the amnesic patients, the study has been central in supporting the popular notion that different forms of memory are supported by distinct and

independent neural systems. Second, by implicating the basal ganglia in category learning in humans, this study added to a growing literature suggesting the same learning processes may underlie both simple stimulus-response learning, as well as “higher cognitive” processes such as categorization (e.g. Gluck and Bower, 1988; Rumelhart and McClelland, 1986). Thus, the study ultimately led to the widely held view that the basal ganglia are critically involved in category learning, at least in some cases. Recent neuroimaging, computational and behavioral studies have elaborated on this initial finding, further exploring basal ganglia contributions to probabilistic category learning (Delgado et al., 2005; Foerde et al., 2006; Frank et al., 2004; Poldrack et al., 2001, 1999; Seger and Cincotta, 2005; Shohamy et al., 2004a, b), as well as to other forms of category learning (Ashby et al., 2003; Nomura et al., 2006; Reber et al., 2003b).

2.2. Memory systems and the basal ganglia: limitations and open questions

The memory systems view of the basal ganglia and non-declarative learning has provided a very useful framework for understanding how the basal ganglia contribute to learning and memory, and for understanding preliminary evidence linking the basal ganglia with category learning. Nonetheless, many open questions remain regarding the specific contributions of the basal ganglia to learning.

First, converging evidence from recent patient and functional neuroimaging (fMRI) studies suggest that there is no one-to-one mapping between non-declarative learning and the basal ganglia. Patients with basal ganglia damage are sometimes spared on non-declarative learning tasks

(Bondi and Kaszniak, 1991; Harrington et al., 1990; Heindel et al., 1989; Reber and Squire, 1999; Smith, 2001; Witt et al., 2002), and are sometimes impaired on declarative memory tasks (Bondi and Kaszniak, 1991; Breen, 1993; Owen et al., 1993a; Pillon et al., 1996; Whittington et al., 2000). Indeed, fMRI often reveals *both* MTL and basal ganglia activation during *both* declarative and non-declarative tasks (Aizenstein et al., 2004; Degonda et al., 2005; Poldrack et al., 2001; Rose et al., 2002; Schendan et al., 2003). This suggests that in the healthy brain, multiple cognitive processes and multiple neural systems may contribute to learning, and a given task can most likely be learned in more than one way. If so, there may be inherent limitations in an approach that defines a task as either declarative or non-declarative. Rather, it may be useful instead to examine specific cognitive strategies and processes that support learning and how these dynamically change over time.

A second limitation is that the double dissociations underlying the dual-systems approach do not provide insight to the neural mechanisms involved in learning. During the 1990s, an extensive literature on the physiology and chemistry of the basal ganglia emerged, suggesting a specific role for the basal ganglia in learning to predict rewards. How do these data relate to the role of the basal ganglia in non-declarative learning? Recent studies have built on the important initial findings regarding basal ganglia contributions to probabilistic category learning to understand (a) how the basal ganglia contribute more generally to learning to predict outcomes in incremental learning contexts, and (b) how to relate the role of the basal ganglia in probabilistic learning to findings regarding the neurophysiological and neurochemical properties of the basal ganglia and its dopaminergic afferents.

3. The basal ganglia and outcome prediction

Recent studies have sought to understand the basic cognitive and neural processes that underlie incremental, feedback-based learning, breaking away from assumptions regarding the declarative or non-declarative nature of a task. These studies have examined how individuals learn in probabilistic feedback-based settings. In the next section, we review studies aimed to elucidate basal ganglia contributions to specific *cognitive strategies*, the *temporal profile* of different forms of learning and their neural substrates, and the *nature of representations* formed during learning.

3.1. How do people learn? Strategies in probabilistic category learning

As reviewed above, the weather prediction task involves learning to predict a category outcome based on the combined presentation of 4 individual cues, which are associated independently and probabilistically with each of 2 category outcomes. Because the cue-outcome associa-

tions are probabilistic, it has been assumed that subjects learn these associations incrementally (and therefore presumably non-declaratively), much as if there were four independent conditioning processes going on in parallel, with subjects' choice on each trial reflecting the accumulated associations among all the present cues (Gluck and Bower, 1988). In fact, this is how the task has been scored, with performance on each trial considered to be correct if a subject's choice reflects the optimal choice for that cue combination, even if the actual weather outcome on that trial was different (e.g., Knowlton et al., 1994, 1996; Poldrack et al., 1999, 2001; Shohamy et al., 2004a, b). This would indeed be the optimal choice strategy, and would allow an ideal learner to score 100% 'optimal correct' responses. However, in the weather prediction task, healthy controls rarely approach optimal levels of performance. This suggests that subjects may be learning the task using sub-optimal strategies. Such strategies may depend on rapid, non-incremental learning processes (and as such may be attributed to declarative memory).

The weather prediction task is particularly amenable to sub-optimal learning strategies, given its specific structure and its complexity. For example, in the weather prediction task used by Knowlton et al. (1994, 1996) and others, two of the cues are highly predictive of the weather, with each associated with one outcome approximately 75% of the time. The other two cards are less predictive, associated with one or the other outcome with about 57% probability. Thus, a subject who focuses attention on just one of the highly predictive cards, and then responds 'sun' or 'rain' based only on the presence or absence of this one card, could achieve 75% 'optimal correct' responses, which is similar to the level of correct responding that most healthy subjects actually achieve (e.g. Gluck et al., 2002; Knowlton et al., 1996; Poldrack et al., 2001; Shohamy et al., 2004b). Thus, such a "one-cue" strategy could conceivably account for the behavior of subjects in probabilistic classification tasks. Further, such a "one-cue" strategy need not be learned over many trials, but could be adopted after a single trial in which the subject experienced a particular cue paired with a particular outcome. In other words, although the amnesic and control groups in the Knowlton et al. (1996) study showed similar percent optimal responding, it is difficult to know whether the two groups were actually using the same strategies or whether qualitatively different strategies might underlie learning in the two groups. Similarly, although the Parkinson's patients performed worse than controls and amnesic patients, the study could not determine if the Parkinson's patients were using the same strategies as the other groups, but doing so less effectively, or if they were using qualitatively different (and less effective) strategies than the other groups.

To address this question, we have used mathematical models to investigate whether subjects' behavior derived from the optimal strategy, or from non-optimal simpler strategies that can putatively be learned based on a single trial (such as a simple "one-cue" rule based on one of the

highly predictive cues). To assess the strategy that each subject used during learning, we generated model response profiles based on how “ideal” participants would respond on each trial if they had been following the optimal strategy, or a simpler non-optimal strategy. We then compared subjects’ actual trial-by-trial choice with the predicted choice for each model, and calculated the degree to which the “ideal” mathematical model fit each participant’s data.

We first examined strategies and learning in healthy young adults (Gluck et al., 2002). The findings suggested that previous assumptions about the dependence of probabilistic learning on incremental, non-declarative processes were wrong: simple non-optimal rules—which could be learned based on a single episode—accounted for much of the behavior of healthy subjects. In fact, a sub-optimal strategy provided a better fit than the optimal strategy for 90% of the subjects. Interestingly, however, the findings also indicated changes in strategies over time, with a shift towards the optimal, incremental strategy later in learning. This suggested that early on, healthy subjects’ choices derive from a sub-optimal strategy; but, with learning, choices gradually come to be driven by the optimal strategy.

This approach to classifying strategies allowed us to ask a related central question: do the basal ganglia contribute equally to different forms of learning, or are they particularly necessary for the incremental processes underlying optimal performance in probabilistic learning? To answer this question, we tested individuals with mild Parkinson’s disease on the weather prediction task, and compared learning and strategies with age-matched healthy controls. To examine how strategies change over time, we extended training to 3 times the training in previous studies (Gluck et al., 2002; Knowlton et al., 1994, 1996), for a total of 600 trials over 3 separate days.

Performance and strategies for the Parkinson’s patients are shown in Fig. 2 (Shohamy et al., 2004b). Overall, the Parkinson’s patients made fewer optimal responses than did controls, consistent with previous findings (Knowlton et al., 1996). However, we found that this impairment was particularly pronounced later in training, rather than early. Next, we examined strategies throughout training. Early in learning (day 1), there was no difference between the patients and the age-matched controls. Both groups’ choices derived primarily from a sub-optimal strategy. However, while the healthy controls gradually shifted over time to the optimal strategy, the Parkinson’s patients did not. Instead, the patients’ choices continued to derive from a sub-optimal strategy throughout the course of the 600 training trials. Thus, this study found that individuals with basal ganglia damage relied on simple rule-based strategies, even more so than did healthy controls (Shohamy et al., 2004b).

We also found that amnesic patients with selective, bilateral hippocampal damage (confirmed via MRI) were impaired at this task, even at the earliest stages, and this

Fig. 2. Individuals with basal ganglia disruption due to mild-to-moderate Parkinson’s disease and healthy controls were tested on probabilistic classification learning using the weather prediction task. (A) Performance (% correct) among the groups over the course of 600 training trials demonstrates that Parkinson’s patients are impaired, especially later in learning. (B) Strategy analyses reveal further differences between the groups, with Parkinson’s patients’ choices deriving from sub-optimal strategies over the course of learning (Shohamy et al., 2004a).

was related to a failure to consistently engage in any strategy (see Hopkins et al., 2004; Meeter et al., 2006, and Meeter et al. in this issue for a more detailed account of these findings, and of how they relate to the earlier Knowlton et al., 1996 findings). These findings suggest that both the MTL and the basal ganglia are necessary for probabilistic category learning. The MTL contributes *early* in learning, consistent with its hypothesized role in rapid encoding of relations between stimuli (Cohen and Eichenbaum, 1993; Eichenbaum and Cohen, 2001). The basal ganglia, by contrast, contribute to the optimal, incrementally learned stimulus–response associations that support *later* learning.

3.1.1. Computational modeling of the basal ganglia in probabilistic learning

The strategy analyses described above indicate that healthy individuals initially approach a difficult probabilistic categorization task by using simple, easily verbalizable strategies and then gradually shift towards more complex optimal strategies over the course of many training trials. How and why do the healthy individuals shift from sub-optimal to optimal strategies? One possibility is that there may be a “shifting mechanism”, that decides when and

how to shift, and that this mechanism is selectively impaired in the Parkinson's patients. Indeed, studies have suggested that Parkinson's patients are impaired at shifting between stimuli and rules (Cools et al., 2001b; Downes et al., 1989; Owen et al., 1993b), a deficit that is typically attributed to dysfunction in frontal cortical areas. Alternatively, there may be a simpler and more parsimonious explanation for this apparent drift from simple to complex strategies in the normal, but not Parkinson's, subjects. As discussed below, simple associative learning mechanisms may result in gradual changes in the extent to which subjects' choices reflect sub-optimal vs. optimal strategies.

A common principle for learning found in biological theories of the basal ganglia (Daw et al., 2005; Frank, 2005; Schultz et al., 1997), as well as in psychological models of classical conditioning (Rescorla and Wagner, 1972) and cognitive models of category learning (Gluck and Bower, 1988; Gluck et al., 1996) is error-correction learning, whereby associative links between stimuli and outcomes are adjusted on a trial-by-trial basis to minimize future expected errors in prediction of the outcome. These error-correcting learning models all share the property that the weights that change fastest early in learning are those which will produce the most rapid decreases in the probability of future errors. The classic example of such a model is the Rescorla–Wagner rule, which accounts for a wide body of data in the classical conditioning domain (Rescorla and Wagner, 1972). Error-correction models have also been applied to category learning data. For example, Gluck and Bower showed that in the initial phases of training, error-correction models rely primarily on single-cue solutions because they are the solutions that provide for the quickest reduction in expected future errors (Gluck and Bower, 1988). Late in training, when and if these single-cue solutions prove insufficient for reducing all the possible error, more complex configural solutions emerge to reduce the error even further. Gluck and Bower demonstrated that the shift from the simple to complex solutions did not require any explicit hypothesis-testing mechanism, but, rather, was a natural emergent property of the error-correction principle in the model.

Could the shifts in strategy by human subjects performing the weather prediction task also be understood as emerging from a single error-correction learning process without an explicit strategy-shifting mechanism? If so, what would this then mean for interpreting the lack of shifting to complex strategies seen in individuals with Parkinson's disease?

Gluck et al. (1993) built upon earlier behavioral models of animal and human learning, to show how a psychobiological model of cortico-hippocampal function in animal conditioning (Gluck and Myers, 1993) could be applied to category learning in normal and amnesic subjects. In the intact model, stimulus–stimulus regularities experienced during early phases of learning contribute to the develop-

ment—via hippocampal mediation—of enriched stimulus representations that allow for higher levels of accuracy and faster learning later in training (Gluck et al., 1996). By this view, the hippocampus contributes to early encoding of stimuli that may both support early simple strategies, as well as facilitate feedback-based stimulus-response learning later on. To evaluate this psychobiological model of category learning we sought to simulate the effects of Parkinson's disease in the same model based on the data from the probabilistic category learning described above (Shohamy et al., 2004b).

Converging data demonstrate that the dopamine system provides a reward-related error-correcting learning signal that is computationally similar to the error-correction signal described above. The dopamine neurons providing this signal are precisely those neurons which are severely depleted even in the earliest stages of Parkinson's disease (Dauer and Przedborski, 2003). One potential implication of this depleted dopamine reward signal in Parkinson's patients would be that the reward feedback on each learning trial is proportionally less in Parkinson's patients than in healthy controls. Within an error-correction-based learning model, this is equivalent to reducing the learning rate, which lessens the amount of learning that will take place on each trial. Although there is no explicit basal ganglia module in Gluck and Myers' original cortico-hippocampal model of conditioning and the extension to category learning (Gluck et al., 1996), the abstract cortical module can be viewed as representing any of the long-term memory regions in the brain, other than the hippocampus, in which learning occurs and is stored, including the cerebral cortex, cerebellum, and the corticostriatal loops which include the basal ganglia. Thus, under this assumption, we explored a working hypothesis that probabilistic category learning in Parkinson's patients might be modeled within Gluck and Myers' (1993) cortico-hippocampal model by reducing the learning rate parameter which controls the trial by trial changes in associative weights in the long-term memory storage region of the model, while leaving intact all other parameters in the model (see Frank et al., 2004, for a related interpretation of reduced dopamine in Parkinson's patients that has an analogous net effect of slowing the rate of learning).

Fig. 3 displays the simulated performance and strategies in the intact and “Parkinson's disease” model. Fig. 3B shows that the intact model appears to move from sub-optimal to optimal strategies across blocks of learning, just as healthy controls do (compare Fig. 2B). These simulations demonstrate that single-system learning models based on error correction show a natural emergent shift from simple, sub-optimal single-cue strategies to complex, optimal, multi-cue strategies during category learning, as shown in Fig. 3B, echoing the earlier results of Gluck and Bower (Gluck and Bower, 1988).

These simulations also show that this model captures the overall pattern found in the Parkinson's patients: the

Fig. 3. Simulated performance and strategies in a computational model of probabilistic category learning based on the Gluck and Myers cortico-hippocampal model. Parkinson's disease is simulated by assuming lower learning rates in learning modules outside of the hippocampus. (A) Overall % correct performance in the simulations show that the Parkinson's model is slower to learn relative to the intact model. (B) The models also capture the overall pattern in the strategy fits from the human data, with the Parkinson's model better fit by sub-optimal strategies. (C) Strategy fits in the Parkinson's model demonstrate that with much extensive training, eventually the models are able to produce optimal responding.

model is slower to learn relative to the intact model; furthermore, it does not adopt the optimal, complex strategy later in training, in contrast to the intact model. The modeling data therefore suggest that no special shifting mechanism need be invoked to explain the failure of patients to use the optimal strategy; their data can be accounted for simply by a reduced learning rate in the cortical module.

Given that both the human data and the model show a trend towards improved learning and a higher proportion of use of the optimal strategy with training, one question is whether the Parkinson's patients would reach performance levels comparable to control subjects with further extended training. In order to examine this question, the model was run on a large number of trials (3000) to examine its performance at asymptote. As shown in Fig. 3C, over 15 sessions, choices in the impaired model (with lower cortical learning rate) eventually reflect the optimal strategy. These data suggest that the optimal strategy reflects a continuous error-correcting learning process based on incremental learning of stimulus–response associations over time. The model suggests that Parkinson's patients are slower to learn than controls, but that with enough training, they will learn to optimally predict probabilistic outcomes. In other words, the deficit in Parkinson's patients may not be qualitative (loss of a specific learning system) but may be quantitative (a generalized slowing in feedback learning, which leads to impairments in shifting from a simple to a complex learning strategy).

3.2. When do people learn? Temporal dynamics of the basal ganglia during learning

The studies reviewed above suggest that there are multiple forms of learning that contribute to probabilistic categorization, and that may emerge at different times during the learning process. Early on, behavior derives from sub-optimal choices about single cues, and these choices are independent of the basal ganglia. Rather, they may reflect memory operations dependent on interactions between the MTL and the prefrontal cortex (PFC). Over time, and later in learning, behavior among healthy controls gradually shifts to reflect optimal choices representing incremental learning of stimulus–response associations across many stimulus cues, behavior that depends on the basal ganglia. To the extent that these optimal choices involving multiple stimulus cues require sensitivity to stimulus–stimulus relations among the different cues, the basal ganglia may be additionally dependent on the hippocampus and other MTL structures and their critical role in the development of appropriate stimulus representations that emerged during early training trials.

This hypothesis is also consistent with data from recent pharmacological studies of probabilistic selection (Frank et al., 2004, 2006). Frank et al. developed a paradigm where subjects were required on each trial to select between two stimuli, each associated with a positive outcome with differing probabilities (e.g. Frank et al., 2004, 2006). These studies revealed that learning the stimulus–outcome

associations depends on the basal ganglia and its dopaminergic afferents (Frank et al., 2004; see further discussion below). Furthermore, Frank and colleagues demonstrated that the MTL also contributes to learning—but only early on: pharmacological disruption of the MTL impairs performance during the first phases of the probabilistic selection task, while later performance is spared.

Similar results have been obtained with functional magnetic resonance imaging (fMRI), which allows an examination of dynamic changes in activity in different brain regions. Poldrack and colleagues investigated activity in the basal ganglia and MTL during probabilistic classification, using the weather prediction task (Poldrack et al., 1999, 2001). Over the course of learning, basal ganglia activity started low, and increased as learning progressed. Interestingly, the opposite pattern was found in the MTL, where activity was high early on, but *decreased* with learning (Poldrack et al., 2001). Activity changes in the MTL and the basal ganglia were negatively correlated, suggesting that these memory systems may interact during learning.

The results from the pharmacological and fMRI data are consistent with the strategy analyses, patient, and behavioral data obtained from the weather prediction task. These findings all suggest an early role for MTL and a later role for basal ganglia. Overall, these data suggest that early MTL-based learning may be necessary for the development of appropriate stimulus representations (and hence appropriate “rules”) which allow subsequent development of optimal strategies based on these representations by the basal ganglia-based optimal strategies, much as was suggested by the earlier computational models of Gluck and Myers (1993) and Gluck et al. (1996).

Alternatively, MTL and basal ganglia may contribute in parallel to learning, with each system governing behavior under different circumstances. This hypothesis is consistent with recent computational studies of the role of the basal ganglia (specifically, the caudate) in reinforcement learning. In particular, Daw et al. (2005) have proposed that the caudate stores the gradual accumulation of knowledge regarding stimulus–outcome associations, integrated over many trials, while PFC—which is interconnected with MTL and critical for episodic memory (e.g. Wagner et al., 1998, 1999)—supports rapidly formed, goal-directed representations of stimulus–outcome contingencies. This model assumes that these caudate- and PFC-based learning processes take place in parallel, with each system governing behavior depending on the specific circumstances (Daw et al., 2005). Although in their model Daw and colleagues do not specify the relative learning rate of each system, this view implies that under probabilistic conditions, the basal ganglia will support optimal performance later in learning, while PFC (and possibly MTL) may guide behavior early on, based on rapidly formed memories supporting sub-optimal strategies. A similar computational approach has been put forth by Frank and Claus (2006), who propose that the basal ganglia slowly integrates the probability of

reward, while PFC maintains information about recent learning trials in working memory (Frank and Claus, 2006).

In superficial contrast to this view, fMRI studies using simpler probabilistic tasks have demonstrated learning-related changes in the basal ganglia that start early and *decrease* later in learning (Delgado et al., 2005; Seger and Cincotta, 2005). For example, Delgado and colleagues developed a “gambling” task, where subjects made categorical decisions whether the numerical value of cards would be higher or lower than 5 (Delgado et al., 2005). A single card was presented on each trial, and a shape on the card predicted whether the card was probabilistically (70%), deterministically (100%), or randomly (50%) associated with one outcome. Delgado and colleagues found caudate activations that appeared early in the experiment and that increased with learning. Once the associations had been well learned, caudate activity decreased, suggesting that caudate activity was related to learning of cue–outcome associations, but not to the ability to act based on previously learned associations. Similar results were obtained with other recent fMRI studies, suggesting that the caudate may be processing the properties of feedback in a reinforcement learning context to improve choice behavior (Tricomi et al., 2004). These studies suggest that *after* learning, behavior may eventually come to be guided by MTL-based declarative strategies (Haruno et al., 2004), or by PFC (Delgado et al., 2005). Similar results have been obtained with electrophysiological studies from monkeys engaged in reversal of extensively trained stimulus–response associations, with learning-related changes in caudate appearing early and PFC supporting later performance (Pasupathy and Miller, 2005). Finally, some computational models have also suggested that early changes in the basal ganglia may be required for later, long term storage of learned stimulus–response associations in PFC (Beiser and Houk, 1998; Frank, 2005; O’Reilly and Frank, 2006).

These findings of early learning-related activity in basal ganglia appear to contrast with fMRI and behavioral data from the weather prediction task. However, the many differences between the paradigms make it difficult to directly compare them. In particular, the weather prediction task is more complex, involves the presentation of multiple stimuli on each trial, and results in a slower learning curve with healthy controls only reaching optimal performance after several hundred trials. Notably, the fMRI studies of the weather prediction task were run only for 150 trials, although optimal performance in this task is achieved after several hundred trials, as shown in Fig. 2. This suggests that extended learning in the weather prediction task might have revealed a later decline in basal ganglia activity after learning had reached asymptote. Furthermore, those studies that showed early changes in basal ganglia (e.g. Delgado et al., 2005; Pasupathy and Miller, 2005) did not examine changes in MTL. Thus it is possible that the basal ganglia activity was preceded by

transient activity in MTL, just as observed in the weather prediction and probabilistic selection tasks. If so, the simpler tasks might invoke exactly the same qualitative pattern of brain activity as more complex tasks, like the weather prediction task—but in the former learning is simply much faster overall than in the latter.

Recent studies further indicate that the temporal profile of brain activation may vary depending on the particular subregion of the basal ganglia investigated. Seger and colleagues (Seger and Cincotta, 2005) examined basal ganglia contributions to a probabilistic classification task where the association between a single cue and a category outcome (“rain” or “sun”) was either probabilistic (70%), deterministic (100%) or random (50%). The study revealed activity in the body and tail of the caudate and in putamen that increased over the course of learning, while activity in the head of the caudate (and the ventral striatum) was related to feedback processing, and decreased over the course of learning. Similar results have been obtained with other behavioral paradigms, as well (Cincotta and Seger, 2007; Haruno and Kawato, 2006; Lehericy et al., 2005; Williams and Eskandar, 2006). Interestingly, the region investigated in the Delgado et al. (2005) study was indeed in the head of the caudate, while the Poldrack et al. (2001) paper focused on a region in the body of the caudate.

In summary, fMRI, electrophysiological and computational studies collectively indicate a role for the basal ganglia in incremental stimulus–response learning. These studies further demonstrate that multiple neural systems may contribute to category learning, either in parallel, or in a competitive interaction. Specifically, data from various probabilistic classification tasks emphasize a role for MTL activity early in learning, while the basal ganglia appear to contribute later in learning as behavior gradually shifts to optimal, integrative strategies. Other studies suggest that basal ganglia activity—especially in the head of the caudate—drives learning of stimulus–response associations early on and this activation decreases once associations become well learned, with behavior perhaps shifting to PFC-guided mechanisms. Future studies are necessary to fully examine the dynamics of MTL and basal ganglia during learning and how these neural changes relate to changes in memory.

3.3. What do people learn? Basal ganglia and representation of learned knowledge

Converging data reviewed above suggest that *both* the MTL and the basal ganglia contribute to probabilistic learning. Therefore, one question is: what are the implications of using one system vs. the other, in terms of the subsequent representation of knowledge? In a parallel line of research, we have focused on this issue, asking when—and how—do people acquire flexible mnemonic representations that allow transfer of knowledge about a category to new instances? How do the MTL and the basal ganglia contribute to such flexible transfer and generalization?

To address these questions, we have been using two-phase learning and transfer tasks to assess representational changes during learning. In these studies, subjects first engage in incremental stimulus–response learning, then are probed to transfer, generalize, or reverse what they have learned to novel contexts, stimuli, or feedback (Myers et al., 2003; Shohamy et al., 2006).

In one such study, subjects engaged in a concurrent discrimination task. On each trial, subjects viewed a pair of objects and were required to choose one object; the chosen object was then raised to show the presence or absence of a smiley face that signaled reward (Fig. 4A). Multiple different pairs of objects were trained concurrently. The cue–outcome association was deterministic, so that the same object in each pair always predicted the smiley face. This task draws on a rich literature of concurrent discrimination in animals (Eichenbaum et al., 1989; Eichenbaum and Cohen, 2001). But, it can also be thought of as a categorization task, with subjects learning to categorize objects (colored shapes) as predicting one of two outcomes (“smiley face” or “no smiley face”).

After subjects learn the associations, they are tested with a surprise transfer phase with new pairs of objects (Fig. 4B). What allows subjects to transfer what they have learned is that the trained and new objects share a common feature. Specifically, during training, each pair differs in either shape, or in color (but not both). During transfer, the relevant feature stays the same, while the irrelevant feature changes. Thus, the initial discrimination shown in Fig. 4A can be learned in two different ways: subjects can learn based on the relationship *between the two stimuli* (e.g. “red beats yellow”). Alternatively, subjects can learn based on the *specific stimulus–response relationship*, regardless of the other non-rewarded stimulus (e.g. “the red hexagon is hiding the smiley face”). Each of these approaches could support optimal responding during the learning phase. However, subjects’ response to the transfer can tell us something about the representational changes that supported learning. A subject who encoded the stimulus–stimulus relationships during learning should generalize perfectly, because the relationship between the stimuli hasn’t changed (e.g., in Fig. 4B, red still beats yellow). By contrast, learning a specific stimulus–response association will not support transfer, because the specific stimulus has changed (the red hexagon of Fig. 4A is no longer a stimulus in Fig. 4B). Notably, these two approaches map well onto the characteristics attributed to the MTL and basal ganglia memory systems: the MTL is thought to support the formation of representations based on stimulus–stimulus relations, and to allow flexible transfer (Cohen and Eichenbaum, 1993; Eichenbaum and Cohen, 2001; Gluck and Myers, 1993). The basal ganglia, by contrast, are thought to support gradual learning of stimulus–response associations, and to result in relatively inflexible representations.

To examine basal ganglia and MTL contributions to learning and transfer, we tested patients with damage to the

Fig. 4. A two-phase concurrent discrimination task: (A) subjects first engage in a series of discrimination trials where they choose the rewarded of two objects on each trial (Learning phase); (B) subsequently, subjects are presented with a series of new pairs where the previously relevant dimension (here, color) remains the same, but the previously irrelevant dimension (here, shape) has changed (transfer phase). (C and D) The data demonstrate a double-dissociation between basal ganglia and MTL contributions to learning: Parkinson's patients (BG) are slow to learn the associations, but are able to transfer; individuals with MTL atrophy (MTL) learn the associations well, but are impaired at transfer (Shohamy et al., 2006; Myers et al., 2003).

basal ganglia (Parkinson's disease), and elderly individuals with mild hippocampal atrophy assessed with structural neuroimaging (Myers et al., 2002; Shohamy et al., 2006). Healthy controls learn the initial discrimination pairs quickly (Fig. 4C), and also transfer well, making very few errors on the new pairs (Fig. 4D). Individuals with hippocampal damage learn as quickly as controls, but their learning is based on pathological, hyper-specific representations, impairing their ability to transfer what they have learned. Basal ganglia damage leads to the opposite pattern: slow learning, but successful transfer, indicating the formation of flexible representations (Myers et al., 2002, 2003; Shohamy et al., 2006). These findings support the idea that the hippocampus and the basal ganglia both contribute, in different ways, to incremental learning. The hippocampus forms flexible representations that can be used in new settings. The basal ganglia form specific inflexible representations that do not generalize well. Healthy people, who are likely to have a more balanced access to both hippocampal and basal ganglia learning systems, optimally access the appropriate representation when making decisions under these circumstances. Thus, in novel contexts, healthy people are able to flexibly draw on past experience to inform decisions.

A recent fMRI study demonstrated similar findings with healthy controls engaged in the weather prediction task (Foerde et al., 2006). Foerde and colleagues manipulated the attentional load during learning, with subjects learning

a set of associations under single-task (full attention) or dual task (split attention) conditions. The study revealed that overall learning under single vs. dual-task conditions elicited relatively more MTL activity and less basal ganglia activity (despite similar levels of performance under both conditions). Foerde and colleagues also administered a post-test questionnaire to assess subjects' ability to flexibly express what they had learned. MTL activation during learning was correlated with performance on the flexibility test, but only for associations learned under single-task conditions. By contrast, basal ganglia activity correlated with learning, but not with flexible transfer, only under dual-task conditions. These findings suggest that associative stimulus-response learning can be supported by both hippocampal and basal ganglia activity, with important qualitative differences in the representation of learned knowledge depending on the neural system engaged during learning.

Finally, recent studies have demonstrated that in some cases what may appear as flexible transfer of knowledge may in fact be supported by reinforcement based stimulus-response learning mechanisms in the basal ganglia, independent of the MTL (Frank et al., 2006). One paradigm that is considered a good index of flexible transfer is transitive inference—the ability to “infer” from learned associations (e.g. A beats B; B beats C) about the relation between stimulus pairs that were never before experienced (A beats C). Several studies have indeed

demonstrated that the MTL is necessary for such inferences (Dusek and Eichenbaum, 1997; Heckers et al., 2004; Preston et al., 2004), consistent with a role for the MTL in flexible transfer. However, Frank and colleagues have hypothesized that this task can be learned in more than one way, and that—at least in some cases—transitive inference can be driven by incremental, implicit, reinforcement-based stimulus–response learning alone (Frank et al., 2006). In support of this hypothesis, pharmacological disruption of declarative memory processes (presumably via disruption of MTL processes) *facilitated*, rather than impaired, transitive inference (Frank et al., 2006).

To summarize, converging data demonstrate that multiple distinct cognitive and neural processes contribute to how people learn to predict outcomes. This emphasizes the importance of breaking away from a priori assumptions regarding the nature of a task (e.g. declarative vs. non-declarative), as well as the value of using model-based approaches to understand the cognitive components contributing to learning. Taken together, these data support the hypothesis that *both* the MTL and basal ganglia contribute to incremental learning with distinct temporal profiles, that there may exist a competitive interaction between them, and that the involvement of each system has important implications for the nature of the representations formed during learning.

4. Feedback and reward modulate learning

A central goal of cognitive neuroscience is to relate cognitive processes to the neural characteristics of underlying brain structures. Significant advances have been made in recent years into the functional neurophysiological, neurochemical, and neurocomputational characteristics of the basal ganglia and its dopaminergic projections (e.g. Bayer and Glimcher, 2005; Beiser and Houk, 1998; Daw and Doya, 2006; Daw et al., 2005; Schultz, 2000; Schultz et al., 1997). Collectively, these studies suggest that dopamine neurons in the basal ganglia are critical for learning to predict rewarding outcomes.

This idea is based on a series of seminal studies demonstrating that midbrain dopamine neurons in animals implement a reward-related “prediction error” (Fiorillo et al., 2003; Hollerman and Schultz, 1998; Schultz, 1998; Schultz et al., 1997). Three key findings link dopamine and reward prediction. First, dopamine neurons produce a strong phasic response when an animal receives an unexpected reward (e.g. juice). Second, if this reward is consistently predicted by a cue (e.g. a tone), then the dopamine response is elicited by the cue, and not the reward—suggesting that dopamine helps signal the prediction of an upcoming reward. Third, if a reward is expected, but is not received, there is a dip in the response of the dopamine signal—presumably indicating a negative error in the reward prediction. Similar findings have now been demonstrated in humans, using functional imaging and a

variety of rewards (e.g. Aron et al., 2004; Delgado et al., 2000; Kirsch et al., 2003; Knutson et al., 2001; McClure et al., 2004; O’Doherty, 2004; Poldrack et al., 2001).

These data suggest that the same neural circuitry implicated in incremental learning is also involved in reward prediction. How do these neuronal data relate to incremental learning in humans? Recent studies have begun to bridge the neural and behavioral perspectives to provide a more complete picture of the neurocognitive mechanisms underlying our ability to predict category outcomes based on past experience.

4.1. The role of feedback in incremental learning

Collectively, studies of the midbrain dopamine system emphasize a role for dopaminergic projections to the striatum in *modifying behavioral responses to environmentally salient stimuli based on response–contingent feedback* (Hollerman and Schultz, 1998; Schultz, 1998, 2000; Schultz et al., 1997). These findings suggest, therefore, that the basal ganglia support learning that relies on trial-by-trial feedback, but not learning by ‘observation’, without feedback. Initial support for this hypothesis came from an fMRI study of probabilistic classification learning, using the weather prediction task. In this study, we found increased activity in the basal ganglia when learning is feedback-based, but not when learning is driven by observation, despite similar levels of performance in both cases (Poldrack et al., 2001). The same effect was also found in midbrain dopaminergic regions. Subsequent studies further specified that midbrain dopamine regions respond selectively to the stimulus and the feedback during probabilistic category learning, and that the degree of activation in these regions is related to the degree of uncertainty for a given trial (Aron et al., 2004), consistent with electrophysiological data from animals engaged in stimulus–response learning (Fiorillo et al., 2003).

We then sought to obtain more direct evidence that the basal ganglia are necessary for feedback-based learning. Because neuroimaging cannot establish the necessity of particular regions for task performance, it is critical to establish that patients with damage to basal ganglia function are specifically impaired at feedback-based learning. To that end, we tested Parkinson’s patients and age-matched controls on a probabilistic classification learning task, similar to the weather prediction task (e.g. Knowlton et al., 1996; Poldrack et al., 2001; Shohamy et al., 2004b). In this study, instead of predicting the weather based on shapes, subjects viewed pictures of Mr. Potatohead dolls and predicted the flavor of ice cream that each doll would choose (chocolate or vanilla). Features on the Mr. Potatohead doll (moustache, bowtie, hat, or glasses) were probabilistically and independently associated with each ice cream flavor (analogous to the cards with shapes in the weather prediction task). Other task features and probabilities were identical to the weather prediction task (e.g. Shohamy et al., 2004b).

Fig. 5. In a probabilistic classification task, subjects predict a category outcome (vanilla or chocolate ice cream) based on features on a Mr. Potatohead doll (moustache, hat, bowtie or glasses): (A) training was either feedback-based, as in prior studies, or observational, involving no feedback; (B) controls showed similar levels of learning under both conditions; individuals with basal ganglia disruption due to Parkinson's disease (PD) were impaired at feedback-based learning, but not at observational learning (Shohamy et al., 2004b).

Subjects were tested on two versions—a 'feedback' version, and an 'observational' version (Fig. 5A). In the feedback version, subjects saw a figure, guessed the outcome, and were provided with trial-by-trial feedback based on their response to each trial, as in prior studies (e.g. Poldrack et al., 2001; Shohamy et al., 2004b). In the observational version, subjects were shown the figure together with the correct outcome on each trial, with no behavioral response required and no feedback presented. In both versions, subjects in each condition were exposed to the same stimulus–outcome information across the course of an experiment. Results, as shown in Fig. 5B, indicated that basal ganglia damage (in patients with Parkinson's disease) leads to impaired feedback-based learning, but intact observational learning of the same task (Shohamy et al., 2004a, 2006). These findings suggest a link between the role of the basal ganglia in human

learning and data from animals regarding midbrain dopamine involvement in feedback processing.

Although converging evidence implicates the basal ganglia in feedback- and reward-based learning, patient data reveal that not *all* feedback-based learning depends on the basal ganglia. For example, Parkinson's patients are spared at learning associations between a single cue and an outcome (Shohamy et al., 2005), or learning of a concurrent discrimination task with few stimuli (Swainson et al., 2006)—even when such learning involves trial-by-trial feedback. Similarly, our data with the weather prediction task indicate that Parkinson's patients are not impaired at learning the sub-optimal single cue strategy, which presumably also involves feedback-based learning of associations between a single cue and an outcome. Thus, the patient data suggest that the feedback-based learning impairments are particularly pronounced when learning depends on integrating information across multiple experiences, emphasizing the necessity of the basal ganglia for these aspects of behavior. Functional imaging and electrophysiological data, by contrast, suggest basal ganglia and midbrain dopamine activity are normally *involved* even in such low-demand feedback-based tasks, indicating a role for these regions in feedback-based learning more generally.

Feedback-based learning has also been examined in other forms of category learning. For example, Ashby and colleagues have been investigating the cognitive and neural systems involved in perceptual category learning, which they sort into two types of tasks: "rule-based" tasks where a simple one-dimensional rule defines category membership, and "information-integration" tasks, in which categories are defined based on a complex multi-dimensional rule (Ashby and Ell, 2001; Ashby et al., 2003). Although both types of tasks involve trial-by-trial feedback, Ashby and colleagues have demonstrated that feedback plays a more important role in driving information-integration tasks relative to rule-based tasks. Interestingly, however, they propose that rule-based tasks, for which feedback is less critical, are dependent upon the basal ganglia. By contrast, information-integration tasks—which are more driven by feedback-based learning—can be learned even with basal ganglia damage, at least in some cases (although this depends on the complexity of the task (Ashby et al., 2002; Filoteo et al., 2005; Maddox et al., 2003, 2004; Shohamy et al., 2005)).

One possible explanation of these apparent discrepancies may be that, in many cases, multiple different approaches and systems may support learning. We have proposed that the basal ganglia system represents gradual learning of feedback-based stimulus–response associations. To the extent that optimal performance depends on such representations, individuals with disrupted basal ganglia function will show impaired performance. However, to the extent that alternative representations may support optimal or near optimal performance, damage to the basal ganglia may not lead to overt behavioral impairments.

As discussed earlier, differential involvement of each system may not be indicated by the ability of subjects to perform at similar levels, but rather by the nature of the representation learned. Future studies are necessary to examine more systematically when Parkinson's patients are spared vs. impaired on feedback-based learning, and how such learning may be supported by alternate systems, such as the MTL or the PFC.

4.2. Dopamine modulation of feedback-based learning

The neuronal mechanisms of dopamine suggest a putative role for dopamine in specific aspects of cognition. Animal data and computational models also suggest that it is not the *absolute* level of dopamine, but rather *relative* levels and *timing* of dopamine release, that are critical for feedback-based learning. This suggests that global enhancement of absolute dopamine levels—such as occurs with many forms of dopaminergic medications—might *impair* feedback-based learning, because increased global dopamine masks the timing and relativity of stimulus-specific signals from dopaminergic neurons. This hypothesis further predicts that this impairment is selective to incremental, feedback-based learning.

Several recent studies support this hypothesis, demonstrating that medication that enhances global dopamine levels in Parkinson's patients can impair some kinds of learning (e.g. Cools et al., 2001a; Frank et al., 2004; Shohamy et al., 2006). For example, in one recent study, we found that patients tested on their normal dopaminergic medication were impaired at feedback-based learning, but not at other forms of learning, nor the ability to transfer what was learned (Shohamy et al., 2006). In this study, patients with mild to moderate Parkinson's disease were tested on the concurrent discrimination task described above (Fig. 4). One group of patients was tested "on" medication: within 3 h of taking their normal dopaminergic medication (L-dopa; a dopamine precursor), which causes an increase in dopamine levels in the striatum. A second group of patients were tested "off" medication, meaning that patients had refrained from taking their dopaminergic medication for approximately 16 h, and thus had low levels of dopamine in the striatum; any dopamine remaining in the brains of these patients was presumably due to physiological release from surviving dopamine neurons in the brain. Thus, if it is the timing and relative levels of dopamine that are critical for feedback-based learning, patients tested "on" medication should be impaired relative to patients tested "off". Indeed, as shown in Fig. 6A, patients tested "on" medication were impaired at learning, while those tested "off" could learn as well as healthy controls. To determine the degree to which the L-dopa related impairment was due to the demands for feedback-based learning, we developed an alternate version of the task in which the feedback demands were reduced, by showing the patients the correct answer the first time each pair was presented for training. Thus, patients were

Fig. 6. The role of dopamine in modulating feedback-based learning was examined by testing patients with Parkinson's disease PD either 'on' or 'off' dopaminergic medication: (A) using the feedback-learning task of Fig. 4A, patients tested 'on' medication—but not 'off'—were impaired at learning relative to healthy controls (CON). This impairment was selective to a concurrently learned condition, where error and feedback processing demands were high; (B) when tested on a low-feedback version, medicated patients were not impaired relative to matched controls (Shohamy et al., 2006).

no longer required to learn by trial-and-error based on feedback; they could learn merely by observation. Fig. 6B shows that, under these conditions, patients tested "on" medication were able to learn the task as well as healthy controls (Shohamy et al., 2006).

Others have similarly proposed, and demonstrated, that the effect of dopaminergic medication on cognition depends on the specific task demands. For example, Cools and colleagues have proposed that systemic L-dopa may result in dopamine "overdose" in those parts of the brain where dopamine is not depleted by disease; such overdose could account for the differential effects of L-dopa on various tasks, with L-dopa alleviating deficits in dopamine depleted neural circuits, but enhancing (or causing) impairments in non-depleted circuits (Cools et al., 2001a, 2006b). Specifically, the degeneration of nigrostriatal projection in Parkinson's disease typically occurs mainly in dorsal striatum early in the course of the disease, and extends to include ventral striatum as the disease progresses. Thus, Cools and colleagues have proposed that early in the disease, enhancing dopamine levels, via dopaminergic medication, may have a positive effect on tasks that depend on the depleted dorsal striatum, but may have a negative over-dosing effect on tasks that depend on

the relatively intact ventral striatum (Cools et al., 2001a, 2006b). In support of this hypothesis, Cools and colleagues found that L-dopa impaired probabilistic reversal learning (associated with ventral striatum) but enhanced task-switching performance (associated more with dorsal striatum). It is interesting to note, however, that in the Cools et al. studies, the two tasks differ not only in the neural circuitry they are presumed to rely on, but also in the kinds of learning processes they involve. In particular, while the probabilistic reversal (which was impaired with L-dopa) involves feedback-based learning that relies on temporally specific, stimulus-specific information, the task-switching ability (which was remediated with L-dopa) does not. Similar results have been obtained with functional imaging and pharmacological manipulations in both healthy and patient populations (Cools et al., 2003, 2006a, b). This suggests two complementary levels at which dopamine modulation can impact cognitive function: (1) at the synaptic level, by modulating stimulus-specific, temporally specific phasic dopamine signals and (2) at the circuit level, by modulating overall levels of dopamine in particular subregions of cortico-striatal circuits.

Another account of the effects of dopamine on feedback-based learning has been advanced recently by Frank and colleagues (e.g. Frank, 2005; Frank et al., 2004). Frank and colleagues have demonstrated that dopamine differentially impacts learning based on whether learning is driven more by positive vs. negative feedback, confirming predictions from computational modeling. Specifically, they proposed (i) that depletion of dopamine due to Parkinson's disease will impair reward-related responses that are necessary for learning based on positive feedback, but will enhance learning based on negative feedback, and (ii) that enhanced dopamine with medication will facilitate learning from positive feedback, but will impair learning from negative feedback. To test this hypothesis, they developed a probabilistic selection task where subjects learned a series of probabilistic forced-choice selections between two alternative stimuli. In each pair, one stimulus was usually rewarded and one was usually not. After learning, subjects were presented with choices between stimuli that had not been paired together during learning. These new pairs could be approached in two ways: either by selecting the stimulus which had previously been associated with reward most often, or, by avoiding the stimulus which had previously been associated with reward least often. This design allowed analysis of the extent to which each individual learned based on positive vs. negative feedback. Frank and colleagues found that Parkinson's patients tested "off" dopaminergic medication were particularly impaired at learning from positive outcomes, compared to negative outcomes, while dopaminergic medication reversed this effect: patients tested "on" medication were particularly impaired at learning based on negative outcomes compared to positive outcomes. Thus, these findings bridge between the physiological data and human learning to demonstrate that the contribution of dopamine to

feedback-based learning depends on the *valence* of the feedback.

Finally, intriguing new data suggest that midbrain dopamine may also contribute to non-feedback, episodic learning supported by the MTL (Adcock et al., 2006; Wittmann et al., 2005). In a recent fMRI study, Adcock and colleagues presented subjects with a series of pictures (each presented once), then tested their memory of the pictures the next day. During presentation, subjects were told their subsequent memory for each picture would be worth either a high or a low monetary reward; the potential monetary value for remembering it was shown prior to each picture's appearance. Adcock and colleagues found that, while subjects waited for the high-value pictures to appear, fMRI activity in midbrain dopamine regions and in the hippocampus became more tightly correlated. This increase in midbrain-to-hippocampus coupling predicted that a forthcoming picture would be remembered, so that overall, stimuli associated with a high reward were better remembered. Memory enhancements have also been demonstrated for cue stimuli which predicted reward (vs. no reward) for correct performance on an upcoming semantic decision (Wittmann et al., 2005). The effects of midbrain dopamine on episodic memory may be mediated by circuitry linking ventral midbrain regions (including the ventral striatum and the ventral tegmental area) with the MTL. Interestingly, novelty may play an important role in gating the interaction between these regions (Lisman and Grace, 2005). These data raise important questions regarding the relationship between feedback-based incremental learning supported by the basal ganglia, and rapidly formed memories supported by the MTL.

In summary, recent computational, pharmacological, and patient studies link the basal ganglia memory system directly to feedback and reward. These studies indicate an important role for the basal ganglia in feedback-based incremental learning and in reward-related learning. These studies also indicate an important role for optimal levels and timed release of dopamine for learning: pharmacological manipulations that increase global dopamine levels can result in either beneficial or detrimental effects, depending on the task. These effects may be related to the effects of systemic dopamine enhancement on the timing and stimulus-specificity of dopamine firing: i.e. receiving the 'wrong' signal at the 'wrong' time. Additionally, enhanced dopamine levels may lead to overdose effects in particular brain circuits, which may be intact or damaged in patients at different stages of Parkinson's disease. Finally, specific circuits and mechanisms may support learning for different feedback valences, resulting in differential effects of dopamine manipulations on learning about negative vs. positive outcomes.

These findings have important implications for how different memory systems contribute to category learning. Specifically, they begin to provide a mechanistic explanation of how and when the basal ganglia and dopamine contribute to category learning, and suggest that category

learning that does not depend critically on gradual trial-by-error learning may not depend on the basal ganglia, but on other systems, such as the MTL and the PFC. Importantly, many forms of category learning indeed do not involve trial-by-error correction processes. Converging evidence suggests that these forms of category learning depend on neural mechanisms not subserved by the basal ganglia (Bozoki et al., 2006; Knowlton and Squire, 1993; Reber et al., 1998, 2003a; Reed et al., 1999).

5. Summary

Converging evidence indicates an important role for the basal ganglia and midbrain dopamine system in learning, particularly in probabilistic category learning. The studies reviewed here emphasize that the basal ganglia are critical for specific aspects of learning, namely, for gradual, incremental, feedback-based learning of associations. Other cognitive strategies, which turn out to be quite important for probabilistic category learning especially early on, do not depend on the basal ganglia. Functional imaging data further suggest that the basal ganglia are specifically necessary for *learning* of associations, but may be less critical for mediating *performance* once associations have been well learned (instead, these later phases of performance may be driven by representations in PFC and/or the MTL). Thus, although prior studies had emphasized a selective role for the basal ganglia in supporting probabilistic learning, recent data suggest a more complex picture, with multiple neural systems contributing to probabilistic learning in different ways, and with different temporal profiles. Important open questions remain regarding the nature of the relationship between basal ganglia based learning and other neural systems. Preliminary evidence suggests that the basal ganglia and the MTL may compete during probabilistic category learning, given negative interactions between them during learning. Finally, the basal ganglia appear to support the formation of relatively inflexible stimulus-response associations that do not generalize to new stimuli and contexts.

In summary, there are many different ways in which healthy people can learn categories. Even within a given paradigm, multiple cognitive and neural systems may contribute in parallel to learning. By breaking down our understanding of the specific cognitive and neural mechanisms contributing to different aspects of learning, recent studies are providing insight into how, and when, these different processes support learning, how they may interact with each other, and the consequences for different forms of learning on the resulting representation of knowledge.

Acknowledgments

We are grateful to Alison Adcock, Roshan Cools, Nathaniel Daw, Mauricio Delgado, Russ Poldrack, and Anthony Wagner for many fruitful discussions and to Russ Poldrack, Lucien Cote, and Jacob Sage for collaborating

with us on some of the research discussed here. We would also like to thank the reviewers for their insightful comments on an earlier draft of the manuscript. Some of the research discussed here was supported by grants from the National Institute of Mental Health (NRSA MH072135 to DS, R01 MH065406 to CEM, and R01 NS047434-02 to MAG), and from the National Science Foundation (BCS-0223910 to MAG).

References

- Adcock, R.A., Thangavel, A., Whitfield-Gabrieli, S., Knutson, B., Gabrieli, J.D., 2006. Reward-motivated learning: mesolimbic activation precedes memory formation. *Neuron* 50 (3), 507–517.
- Aizenstein, H.J., Stenger, V.A., Cochran, J., Clark, K., Johnson, M., Nebes, R.D., et al., 2004. Regional brain activation during concurrent implicit and explicit sequence learning. *Cerebral Cortex* 14 (2), 199–208.
- Albin, R.L., Young, A.B., Penney, J.B., 1989. The functional anatomy of basal ganglia disorders. *Trends in Neuroscience* 12 (10), 366–375.
- Alexander, G.E., DeLong, M.R., Strick, P.L., 1986. Parallel organization of functionally segregated circuits linking basal ganglia and cortex. *Annual Reviews in Neuroscience* 9, 357–381.
- Aron, A.R., Shohamy, D., Clark, J., Myers, C., Gluck, M.A., Poldrack, R.A., 2004. Human midbrain sensitivity to cognitive feedback and uncertainty during classification learning. *Journal of Neurophysiology* 92 (2), 1144–1152.
- Ashby, F.G., Ell, S.W., 2001. The neurobiology of human category learning. *Trends in Cognitive Sciences* 5 (5), 204–210.
- Ashby, F.G., Maddox, W.T., Bohil, C.J., 2002. Observational versus feedback training in rule-based and information-integration category learning. *Memory Cognition* 30 (5), 666–677.
- Ashby, F.G., Noble, S., Filoteo, J.V., Waldron, E.M., Ell, S.W., 2003. Category learning deficits in parkinson's disease. *Neuropsychology* 17 (1), 115–124.
- Bayer, H.M., Glimcher, P.W., 2005. Midbrain dopamine neurons encode a quantitative reward prediction error signal. *Neuron* 47 (1), 129–141.
- Beiser, D.G., Houk, J.C., 1998. Model of cortical-basal ganglionic processing: encoding the serial order of sensory events. *Journal of Neurophysiology* 79 (6), 3168–3188.
- Bondi, M.W., Kaszniak, A.W., 1991. Implicit and explicit memory in alzheimer's disease and parkinson's disease. *Journal of Clinical Experiments in Neuropsychology* 13 (2), 339–358.
- Bozoki, A., Grossman, M., Smith, E.E., 2006. Can patients with Alzheimer's disease learn a category implicitly? *Neuropsychologia* 44 (5), 816–827.
- Breen, E.K., 1993. Recall and recognition memory in Parkinson's disease. *Cortex* 29 (1), 91–102.
- Calabresi, P., Maj, R., Mercuri, N.B., Bernardi, G., 1992. Coactivation of d1 and d2 dopamine receptors is required for long-term synaptic depression in the striatum. *Neuroscience Letters* 142 (1), 95–99.
- Cepeda, C., Buchwald, N.A., Levine, M.S., 1993. Neuromodulatory actions of dopamine in the neostriatum are dependent upon the excitatory amino acid receptor subtypes activated. *Proceedings of the National Academy of Sciences of the United States of America* 90 (20), 9576–9580.
- Cincotta, C.M., Seger, C.A., 2007. Dissociation between striatal regions while learning to categorize via feedback and via observation. *Journal of Cognitive Neuroscience* 19 (2), 249–265.
- Cohen, N.J., Eichenbaum, H., 1993. *Memory, Amnesia, and the Hippocampal System*. MIT Press, Cambridge, MA.
- Cools, R., Barker, R.A., Sahakian, B.J., Robbins, T.W., 2001a. Enhanced or impaired cognitive function in Parkinson's disease as a function of dopaminergic medication and task demands. *Cerebral Cortex* 11 (12), 1136–1143.

- Cools, R., Barker, R.A., Sahakian, B.J., Robbins, T.W., 2001b. Mechanisms of cognitive set flexibility in Parkinson's disease. *Brain* 124 (Pt. 12), 2503–2512.
- Cools, R., Barker, R.A., Sahakian, B.J., Robbins, T.W., 2003. L-dopa medication remediates cognitive inflexibility, but increases impulsivity in patients with Parkinson's disease. *Neuropsychologia* 41 (11), 1431–1441.
- Cools, R., Altamirano, L., D'Esposito, M., 2006a. Reversal learning in Parkinson's disease depends on medication status and outcome valence. *Neuropsychologia* 44 (10), 1663–1673.
- Cools, R., Lewis, S.J., Clark, L., Barker, R.A., Robbins, T.W., 2006b. L-dopa disrupts activity in the nucleus accumbens during reversal learning in Parkinson's disease. *Neuropsychopharmacology*.
- Dauer, W., Przdeborski, S., 2003. Parkinson's disease: mechanisms and models. *Neuron* 39 (6), 889–909.
- Daw, N.D., Doya, K., 2006. The computational neurobiology of learning and reward. *Current Opinion in Neurobiology* 16 (2), 199–204.
- Daw, N.D., Niv, Y., Dayan, P., 2005. Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nature Neuroscience* 8 (12), 1704–1711.
- Degonda, N., Mondadori, C.R., Bosshardt, S., Schmidt, C.F., Boesiger, P., Nitsch, R.M., et al., 2005. Implicit associative learning engages the hippocampus and interacts with explicit associative learning. *Neuron* 46 (3), 505–520.
- Delgado, M.R., Nystrom, L.E., Fissell, C., Noll, D.C., Fiez, J.A., 2000. Tracking the hemodynamic responses to reward and punishment in the striatum. *Journal of Neurophysiology* 84 (6), 3072–3077.
- Delgado, M.R., Miller, M.M., Inati, S., Phelps, E.A., 2005. An fMRI study of reward-related probability learning. *Neuroimage* 24 (3), 862–873.
- Downes, J.J., Roberts, A.C., Sahakian, B.J., Evenden, J.L., Morris, R.G., Robbins, T.W., 1989. Impaired extra-dimensional shift performance in medicated and unmedicated parkinson's disease: evidence for a specific attentional dysfunction. *Neuropsychologia* 27 (11–12), 1329–1343.
- Dusek, J.A., Eichenbaum, H., 1997. The hippocampus and memory for orderly stimulus relations. *Proceedings of the National Academy of Sciences of the United States of America* 94 (13), 7109–7114.
- Eichenbaum, H.E., Cohen, N.J., 2001. *From Conditioning to Conscious Recollection: Memory Systems of the Brain*. Oxford University Press, New York.
- Eichenbaum, H., Mathews, P., Cohen, N.J., 1989. Further studies of hippocampal representation during odor discrimination learning. *Behavioral Neuroscience* 103 (6), 1207–1216.
- Filoteo, J.V., Maddox, W.T., Salmon, D.P., Song, D.D., 2005. Information-integration category learning in patients with striatal dysfunction. *Neuropsychology* 19 (2), 212–222.
- Fiorillo, C.D., Tobler, P.N., Schultz, W., 2003. Discrete coding of reward probability and uncertainty by dopamine neurons. *Science* 299 (5614), 1898–1902.
- Foerde, K., Knowlton, B.J., Poldrack, R.A., 2006. Modulation of competing memory systems by distraction. *Proceedings of the National Academy of Sciences of the United States of America* 103 (31), 11778–11783.
- Frank, M.J., 2005. Dynamic dopamine modulation in the basal ganglia: a neurocomputational account of cognitive deficits in medicated and nonmedicated parkinsonism. *Journal of Cognitive Neuroscience* 17 (1), 51–72.
- Frank, M.J., Claus, E.D., 2006. Anatomy of a decision: striato-orbitofrontal interactions in reinforcement learning, decision making, and reversal. *Psychology Reviews* 113 (2), 300–326.
- Frank, M.J., Seeberger, L.C., O'Reilly, R.C., 2004. By carrot or by stick: cognitive reinforcement learning in Parkinsonism. *Science* 306 (5703), 1940–1943.
- Frank, M.J., O'Reilly, R.C., Curran, T., 2006. When memory fails, intuition reigns: Midazolam enhances implicit inference in humans. *Psychology and Science* 17 (8), 700–707.
- Gabrieli, J.D., 1998. Cognitive neuroscience of human memory. *Annual Reviews in Psychology* 49, 87–115.
- Gluck, M.A., Bower, G.H., 1988. From conditioning to category learning: an adaptive network model. *Journal of Experimental Psychology—General* 117 (3), 227–247.
- Gluck, M.A., Myers, C.E., 1993. Hippocampal mediation of stimulus representation: a computational theory. *Hippocampus* 3 (4), 491–516.
- Gluck, M.A., Oliver, L.M., Myers, C.E., 1996. Late-training amnesic deficits in probabilistic category learning: a neurocomputational analysis. *Learning & Memory* 3 (4), 326–340.
- Gluck, M.A., Shohamy, D., Myers, C., 2002. How do people solve the “weather prediction” task? Individual variability in strategies for probabilistic category learning. *Learning & Memory* 9 (6), 408–418.
- Harrington, D.L., Haaland, K.Y., Yeo, R.A., Marder, E., 1990. Procedural memory in Parkinson's disease: impaired motor but not visuoperceptual learning. *Journal of Clinical and Experimental Neuropsychology* 12 (2), 323–339.
- Haruno, M., Kawato, M., 2006. Different neural correlates of reward expectation and reward expectation error in the Putamen and caudate nucleus during stimulus–action–reward association learning. *Journal of Neurophysiology* 95 (2), 948–959.
- Haruno, M., Kuroda, T., Doya, K., Toyama, K., Kimura, M., Samejima, K., et al., 2004. A neural correlate of reward-based behavioral learning in caudate nucleus: a functional magnetic resonance imaging study of a stochastic decision task. *Journal of Neuroscience* 24 (7), 1660–1665.
- Heckers, S., Zalesak, M., Weiss, A.P., Ditman, T., Titone, D., 2004. Hippocampal activation during transitive inference in humans. *Hippocampus* 14 (2), 153–162.
- Heindel, W.C., Salmon, D.P., Shults, C.W., Walicke, P.A., Butters, N., 1989. Neuropsychological evidence for multiple implicit memory systems: a comparison of Alzheimer's, Huntington's, and Parkinson's disease patients. *Journal of Neuroscience* 9 (2), 582–587.
- Hollerman, J.R., Schultz, W., 1998. Dopamine neurons report an error in the temporal prediction of reward during learning. *Nature Neuroscience* 1 (4), 304–309.
- Hopkins, R.O., Myers, C.E., Shohamy, D., Grossman, S., Gluck, M., 2004. Impaired probabilistic category learning in hypoxic subjects with hippocampal damage. *Neuropsychologia* 42 (4), 524–535.
- Kesner, R.P., Bolland, B.L., Dakis, M., 1993. Memory for spatial locations, motor responses, and objects: triple dissociation among the hippocampus, caudate nucleus, and extrastriate visual cortex. *Experimental Brain Research* 93 (3), 462–470.
- Kirsch, P., Schienle, A., Stark, R., Sammer, G., Blecker, C., Walter, B., et al., 2003. Anticipation of reward in a nonaversive differential conditioning paradigm and the brain reward system: an event-related fMRI study. *Neuroimage* 20 (2), 1086–1095.
- Knowlton, B.J., Squire, L.R., 1993. The learning of categories: parallel brain systems for item memory and category knowledge. *Science* 262 (5140), 1747–1749.
- Knowlton, B.J., Squire, L.R., Gluck, M.A., 1994. Probabilistic classification learning in amnesia. *Learning & Memory* 1 (2), 106–120.
- Knowlton, B.J., Mangels, J.A., Squire, L.R., 1996. A neostriatal habit learning system in humans. *Science* 273 (5280), 1399–1402.
- Knutson, B., Adams, C.M., Fong, G.W., Hommer, D., 2001. Anticipation of increasing monetary reward recruits nucleus accumbens. *Journal of Neuroscience* 21 (16), RC159.
- Lehericy, S., Benali, H., Van de Moortele, P.F., Pelegrini-Issac, M., Waechter, T., Ugurbil, K., et al., 2005. Distinct basal ganglia territories are engaged in early and advanced motor sequence learning. *Proceedings of the National Academy of Sciences of the United States of America* 102 (35), 12566–12571.
- Lisman, J.E., Grace, A.A., 2005. The hippocampal-vta loop: controlling the entry of information into long-term memory. *Neuron* 46 (5), 703–713.
- Maddox, W.T., Ashby, F.G., Bohil, C.J., 2003. Delayed feedback effects on rule-based and information-integration category learning. *Journal of Experimental Psychology—Learning Memory and Cognition* 29 (4), 650–662.

- Maddox, W.T., Ashby, F.G., Ing, A.D., Pickering, A.D., 2004. Disrupting feedback processing interferes with rule-based but not information-integration category learning. *Memory and Cognition* 32 (4), 582–591.
- McDonald, R.J., White, N.M., 1993. A triple dissociation of memory systems: hippocampus, amygdala, and dorsal striatum. *Behavioral Neuroscience* 107 (1), 3–22.
- McClure, S.M., Li, J., Tomlin, D., Cypert, K.S., Montague, L.M., Montague, P.R., 2004. Neural correlates of behavioral preference for culturally familiar drinks. *Neuron* 44 (2), 379–387.
- Meeter, M., Myers, C.E., Shohamy, D., Hopkins, R.O., Gluck, M.A., 2006. Strategies in probabilistic categorization: results from a new way of analyzing performance. *Learning & Memory* 13 (2), 230–239.
- Myers, C.E., Kluger, A., Golomb, J., Ferris, S., de Leon, M.J., Schnirman, G., et al., 2002. Hippocampal atrophy disrupts transfer generalization in nondemented elderly. *Journal of Geriatric Psychiatry and Neurology* 15 (2), 82–90.
- Myers, C.E., Shohamy, D., Gluck, M.A., Grossman, S., Kluger, A., Ferris, S., et al., 2003. Dissociating hippocampal versus basal ganglia contributions to learning and transfer. *Journal of Cognitive Neuroscience* 15 (2), 185–193.
- Nomura, E.M., Maddox, W.T., Filoteo, J.V., Ing, A.D., Gitelman, D.R., Parrish, T.B., et al., 2006. Neural correlates of rule-based and information-integration visual category learning. *Cerebral Cortex*.
- O'Doherty, J.P., 2004. Reward representations and reward-related learning in the human brain: insights from neuroimaging. *Current Opinion in Neurology* 14 (6), 769–776.
- 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 Computing* 18 (2), 283–328.
- Owen, A.M., Beksinska, M., James, M., Leigh, P.N., Summers, B.A., Marsden, C.D., et al., 1993a. Visuospatial memory deficits at different stages of Parkinson's disease. *Neuropsychologia* 31 (7), 627–644.
- Owen, A.M., Roberts, A.C., Hodges, J.R., Summers, B.A., Polkey, C.E., Robbins, T.W., 1993b. Contrasting mechanisms of impaired attentional set-shifting in patients with frontal lobe damage or Parkinson's disease. *Brain* 116 (Pt 5), 1159–1175.
- Packard, M.G., 1999. Glutamate infused posttraining into the hippocampus or caudate-putamen differentially strengthens place and response learning. *Proceedings of the National Academy of Sciences of the United States of America* 96 (22), 12881–12886.
- Packard, M.G., McGaugh, J.L., 1996. Inactivation of hippocampus or caudate nucleus with lidocaine differentially affects expression of place and response learning. *Neurobiology of Learning & Memory* 65 (1), 65–72.
- Packard, M.G., Hirsh, R., White, N.M., 1989. Differential effects of fornix and caudate nucleus lesions on two radial maze tasks: Evidence for multiple memory systems. *Journal of Neuroscience* 9 (5), 1465–1472.
- Pasupathy, A., Miller, E.K., 2005. Different time courses of learning-related activity in the prefrontal cortex and striatum. *Nature* 433 (7028), 873–876.
- Pillon, B., Ertle, S., Deweer, B., Sarazin, M., Agid, Y., Dubois, B., 1996. Memory for spatial location is affected in Parkinson's disease. *Neuropsychologia* 34 (1), 77–85.
- Poldrack, R.A., Prabhakaran, V., Seger, C.A., Gabrieli, J.D., 1999. Striatal activation during acquisition of a cognitive skill. *Neuropsychology* 13 (4), 564–574.
- Poldrack, R.A., Clark, J., Pare-Blagoev, E.J., Shohamy, D., Creso Moyano, J., Myers, C., et al., 2001. Interactive memory systems in the human brain. *Nature* 414 (6863), 546–550.
- Preston, A.R., Shrager, Y., Dudukovic, N.M., Gabrieli, J.D., 2004. Hippocampal contribution to the novel use of relational information in declarative memory. *Hippocampus* 14 (2), 148–152.
- Reber, P.J., Squire, L.R., 1999. Intact learning of artificial grammars and intact category learning by patients with parkinson's disease. *Behavioral Neuroscience* 113 (2), 235–242.
- Reber, P.J., Stark, C.E., Squire, L.R., 1998. Cortical areas supporting category learning identified using functional MRI. *Proceedings of the National Academy of Sciences of the United States of America* 95 (2), 747–750.
- Reber, P.J., Gitelman, D.R., Parrish, T.B., Mesulam, M.M., 2003a. Dissociating explicit and implicit category knowledge with fMRI. *Journal of Cognitive Neuroscience* 15 (4), 574–583.
- Reber, P.J., Martinez, L.A., Weintraub, S., 2003b. Artificial grammar learning in alzheimer's disease. *Cognitive Affects and Behavioral Neuroscience* 3 (2), 145–153.
- Reed, J.M., Squire, L.R., Patalano, A.L., Smith, E.E., Jonides, J., 1999. Learning about categories that are defined by object-like stimuli despite impaired declarative memory. *Behavioral Neuroscience* 113 (3), 411–419.
- Rescorla, R.A., Wagner, A.R., 1972. A theory of pavlovian conditioning: variations in the effectiveness of reinforcement and nonreinforcement. *Appleton-Century-Crofts*.
- Robbins, T.W., 1996. Refining the taxonomy of memory. *Science* 273 (5280), 1353–1354.
- Rose, M., Haider, H., Weiller, C., Buchel, C., 2002. The role of medial temporal lobe structures in implicit learning: an event-related fMRI study. *Neuron* 36 (6), 1221–1231.
- Rumelhart, D.E., McClelland, J., Group, P.R., 1986. *Parallel distributed processing: explorations in the microstructure of cognition*. MIT press, Cambridge, MA.
- Saint-Cyr, J.A., Taylor, A.E., Lang, A.E., 1988. Procedural learning and neostriatal dysfunction in man. *Brain* 111 (Pt 4), 941–959.
- Schacter, D.L., Wagner, A.D., 1999. Medial temporal lobe activations in fmri and pet studies of episodic encoding and retrieval. *Hippocampus* 9 (1), 7–24.
- Schendan, H.E., Searl, M.M., Melrose, R.J., Stern, C.E., 2003. An fMRI study of the role of the medial temporal lobe in implicit and explicit sequence learning. *Neuron* 37 (6), 1013–1025.
- Schultz, W., 1998. The phasic reward signal of primate dopamine neurons. *Advances in Pharmacology* 42, 686–690.
- Schultz, W., 2000. Multiple reward signals in the brain. *Nature Reviews Neuroscience* 1 (3), 199–207.
- Schultz, W., Dayan, P., Montague, P.R., 1997. A neural substrate of prediction and reward. *Science* 275 (5306), 1593–1599.
- Seger, C.A., Cincotta, C.M., 2005. The roles of the caudate nucleus in human classification learning. *Journal of Neuroscience* 25 (11), 2941–2951.
- Shohamy, D., Myers, C.E., Grossman, S., Sage, J., Gluck, M.A., Poldrack, R.A., 2004a. Cortico-striatal contributions to feedback-based learning: converging data from neuroimaging and neuropsychology. *Brain* 127 (Pt 4), 851–859.
- Shohamy, D., Myers, C.E., Onlaor, S., Gluck, M.A., 2004b. Role of the basal ganglia in category learning: how do patients with Parkinson's disease learn? *Behavioral Neuroscience* 118 (4), 676–686.
- Shohamy, D., Myers, C.E., Grossman, S., Sage, J., Gluck, M.A., 2005. The role of dopamine in cognitive sequence learning: evidence from Parkinson's disease. *Behavioral Brain Research* 156 (2), 191–199.
- Shohamy, D., Myers, C.E., Geghman, K.D., Sage, J., Gluck, M.A., 2006. L-dopa impairs learning, but spares generalization, in Parkinson's disease. *Neuropsychologia* 44 (5), 774–784.
- Smith, C., 2001. Sleep states and memory processes in humans: procedural versus declarative memory systems. *Sleep Medicine Reviews* 5 (6), 491–506.
- Squire, L.R., 1987. The organization and neural substrates of human memory. *International Journal of Neurology* 21–22, 218–222.
- Squire, L.R., 1992. Memory and the hippocampus: a synthesis from findings with rats, monkeys, and humans. *Psychology Reviews* 99 (2), 195–231.
- Swainson, R., Rogers, R.D., Sahakian, B.J., Summers, B.A., Polkey, C.E., Robbins, T.W., 2000. Probabilistic learning and reversal deficits in patients with Parkinson's disease or frontal or temporal lobe lesions: possible adverse effects of dopaminergic medication. *Neuropsychologia* 38 (5), 596–612.

- Swainson, R., SenGupta, D., Shetty, T., Watkins, L.H., Summers, B.A., Sahakian, B.J., et al., 2006. Impaired dimensional selection but intact use of reward feedback during visual discrimination learning in Parkinson's disease. *Neuropsychologia* 44 (8), 1290–1304.
- Tricomi, E.M., Delgado, M.R., Fiez, J.A., 2004. Modulation of caudate activity by action contingency. *Neuron* 41 (2), 281–292.
- Wagner, A.D., Desmond, J.E., Glover, G.H., Gabrieli, J.D., 1998. Prefrontal cortex and recognition memory. Functional-MRI evidence for context-dependent retrieval processes. *Brain* 121 (Pt 10), 1985–2002.
- Wagner, A.D., Koutstaal, W., Schacter, D.L., 1999. When encoding yields remembering: insights from event-related neuroimaging. *Philosophical Transactions of the Royal Society of London Series B-Biological Sciences* 354 (1387), 1307–1324.
- White, N.M., 1997. Mnemonic functions of the basal ganglia. *Current Opinion in Neurobiology* 7 (2), 164–169.
- Whittington, C.J., Podd, J., Kan, M.M., 2000. Recognition memory impairment in parkinson's disease: power and meta-analyses. *Neuropsychology* 14 (2), 233–246.
- Wickens, J.R., Begg, A.J., Arbuthnott, G.W., 1996. Dopamine reverses the depression of rat corticostriatal synapses, which normally follows high-frequency stimulation of cortex in vitro. *Neuroscience* 70 (1), 1–5.
- Williams, Z.M., Eskandar, E.N., 2006. Selective enhancement of associative learning by microstimulation of the anterior caudate. *Nature Neuroscience* 9 (4), 562–568.
- Witt, K., Nuhsman, A., Deuschl, G., 2002. Dissociation of habit-learning in parkinson's and cerebellar disease. *Journal of Cognitive Neuroscience* 14 (3), 493–499.
- Wittmann, B.C., Schott, B.H., Guderian, S., Frey, J.U., Heinze, H.J., Düzel, E., 2005. Reward-related fMRI activation of dopaminergic midbrain is associated with enhanced hippocampus-dependent long-term memory formation. *Neuron* 45 (3), 459–467.