18

Rules through Recursion: How Interactions between the Frontal Cortex and Basal Ganglia May Build Abstract, Complex Rules from Concrete, Simple Ones

Earl K. Miller and Timothy J. Buschman

The brain has evolved to deal with two competing requirements—it must respond quickly to familiar situations while being able to adapt to novel ones and plan for the future. Quickly responding to the immediate environment in a reflexive, or habitual, fashion is relatively straightforward: Familiar stimuli activate well-established neural pathways that produce stereotyped behaviors. This is so-called "bottom-up," or "stimulus-driven," processing. These behaviors can be executed quickly and automatically because they are "concrete"; they rely on specific stimulus-response relationships, and the same cue always elicits the same response. It is an axiom of neuroscience that such reflexive reactions are formed by repeated activation of neural pathways, which strengthens their connections. Then, they can be simply triggered—fired off in an automatic fashion, with little variation and, hence, little need for internal oversight.

In contrast, truly sophisticated, goal-directed behavior requires a different mode of operation. Novel situations must be resolved, and goal direction requires the ability to act on, not just react to, a familiar environment. Navigating complex situations to achieve long-planned goals cannot rely on uncoordinated reactions. They must be orchestrated "top-down" from within oneself. By acquiring and building on knowledge of how the world works, we can predict what outcomes are desirable and determine what strategies will aid in attaining them. However, simply recording and replaying previous experiences does not suffice. Relevant relationships need to be sorted out from spurious coincidences, and smart animals get the "big picture" of the jigsaw puzzle of their experiences: They find the common structure across a wide range of experiences to form "abstract" rules—generalized principles that can be readily adapted to novel situations. These abstract rules are the overarching principles and general concepts that are the basis for high-level thought. They provide

the foresight needed for achieving distant goals; because abstract rules, by definition, are generalized across many past experiences, they provide the basis for generalizing to (predicting) future events.

The goal of this chapter is to review evidence that goal-directed behavior depends on interactions between two different "styles" of learning mechanisms in different frontal lobe systems. Specifically, we propose that ever more complex thoughts and actions can be bootstrapped from simpler ones through recursive interactions between fast, reward-based plasticity in the basal ganglia (BG) and slower, more Hebbian-based plasticity in the frontal cortex. By having these two systems interact in recursive processing loops, the brain can learn new concrete relationships quickly, but also can take the time to link in more experiences and more gradually build up abstract, big-picture thoughts and sophisticated actions.

ABSTRACT RULES AND THE PREFRONTAL CORTEX

Abstract rules lie at the center of the ability to coordinate thought and action and direct them toward a goal. Virtually all long-term, goal-directed behaviors are learned, and thus depend on a cognitive system that can acquire the rules of the game: what outcomes are possible, what actions might be successful at achieving them, what the costs of those actions might be, etc. Consider the set of rules invoked when we dine in a restaurant, such as "wait to be seated," "order," and "pay the bill." These rules are long divorced from the specific circumstances in which they were learned and thus give us an idea about what to expect (and what is expected of us) when we try a new restaurant. We have learned to generalize beyond specific experiences and construct a set of abstract rules that direct behavior. These rules orchestrate processing in diverse brain regions along a common, internal theme. It is widely accepted that the prefrontal cortex (PFC)—a neocortical region that finds its greatest elaboration in humans—is centrally involved in this process.

The PFC is situated at the anterior end of the brain and reaches its greatest elaboration and relative size in the primate, especially human, brain (Fuster, 1995). Thus, it is presumably involved in our advanced cognitive capabilities and goal-directed behaviors. Indeed, recent imaging work has suggested that the size of the PFC is directly correlated with intelligence in adult humans (Haier et al., 2004). The PFC seems anatomically well situated to play a role in the creation and implementation of abstract rules. As shown in Figure 18–1, the PFC receives and sends projections to most of the cerebral cortex (with the exception of primary sensory and motor cortices), as well as all of the major subcortical systems, such as the hippocampus, amygdala, cerebellum, and most importantly for this chapter, the BG (Porrino et al., 1981; Amaral and Price, 1984; Amaral, 1986; Selemon and Goldman-Rakic, 1988; Barbas and De Olmos, 1990; Eblen and Graybiel, 1995; Croxson et al., 2005). The PFC seems to be a hub of cortical processing, able to synthesize a wide range of external and internal information and also exert control over much of the cortex. Although

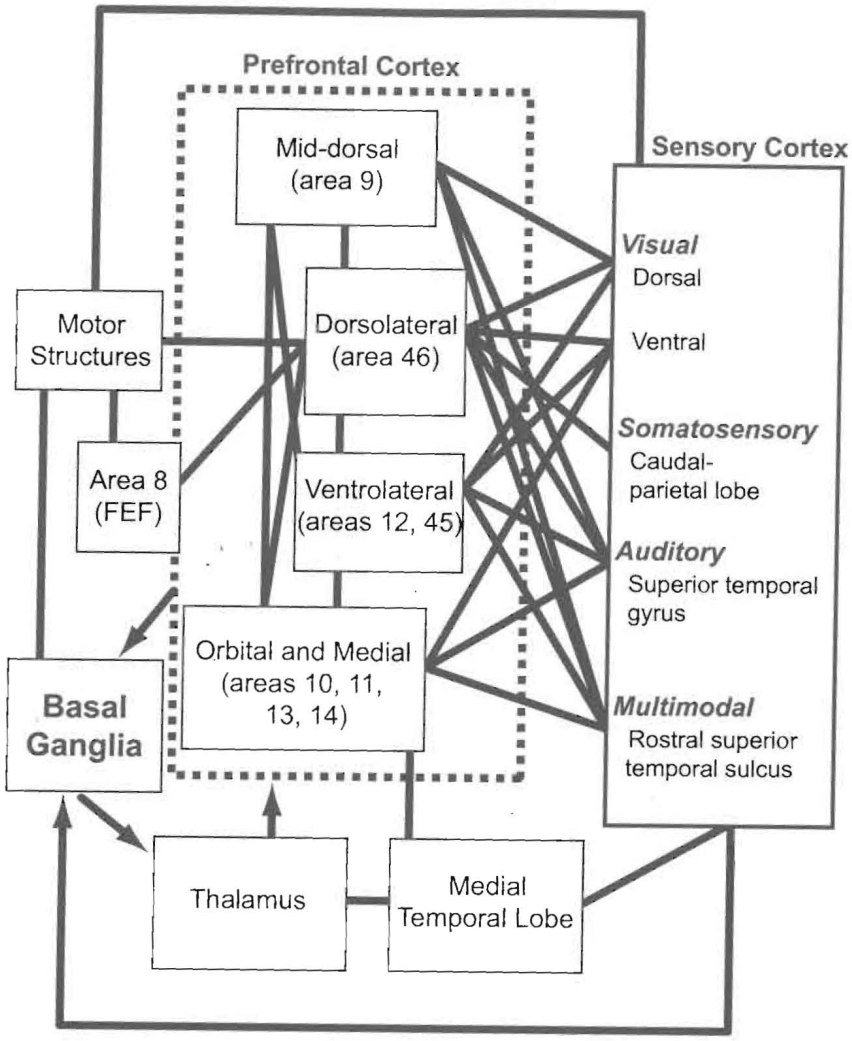

Figure 18–1 Schematic diagram of some of the extrinsic and intrinsic connections of the prefrontal cortex. The partial convergence of inputs from many brain systems and internal connections of the prefrontal cortex (PFC) may allow it to play a central role in the synthesis of diverse information needed for complex behavior. Most connections are reciprocal; the exceptions are indicated by *arrows*. The frontal eye field (FEF) has variously been considered either adjacent to, or part of, the PFC. Here, we compromise by depicting it as adjacent to, yet touching, the PFC. (Adapted from Miller and Cohen, *Annual Review of Neuroscience*, 24, 167–202.)

different PFC subdivisions have distinct patterns of interconnections with other brain systems (e.g., lateral—sensory and motor cortex; orbital—limbic), there are prodigious connections both within and between PFC subdivisions, ensuring a high degree of integration of information (Pandya and Barnes, 1987; Barbas and Pandya, 1989; Pandya and Yeterian, 1990; Barbas et al., 1991; Petrides and Pandya, 1999). Additionally, the heavy reciprocal interconnections between regions provide an infrastructure ideal for abstract learning—one that can act as a large associative network for detecting and storing associations between diverse events, experiences, and internal states. After learning, such a network can complete or "recall" an entire pattern given a subset of its inputs, an ability that may allow for a given situation to be recognized as a specific instance of an internal model of a more abstract one.

In addition to the anatomical evidence, there is a large amount of psychological, lesion, and neurophysiological evidence supporting the role of the frontal cortex in learning abstract rules (also see Chapter 2). Indeed, neurophysiological studies in animals and imaging studies in humans have shown that the PFC has many of the attributes necessary for representing abstract rules (Miller, 2000). First, the neurons sustain their activity across short, multisecond memory delays (Pribram et al., 1952; Fuster and Alexander, 1971; Fuster, 1973; Funahashi et al., 1989; Miller et al., 1996). This property is crucial for goal-directed behavior, which, unlike "ballistic" reflexes, typically extends over time. Second, neurons within the PFC are highly multimodal, representing a wide range of information, and the cells are plastic—with training, they learn to represent task-relevant information. For example, after training on a wide range of operant tasks, many PFC neurons (typically one-third to one-half of the population) reflect the learned task contingencies—the logic or rules of the task (White and Wise, 1999; Asaad et al., 2000; Wallis et al., 2001; Mansouri et al., 2006). For example, neurons have been found to represent visual categories (see Chapter 17) and small numbers (Nieder et al., 2002), whereas some neurons might activate in anticipation of a forthcoming expected reward or a relevant cue (Watanabe, 1996; Rainer et al., 1999; Wallis and Miller, 2003; Padoa-Schioppa and Assad, 2006). In short, the PFC does, indeed, act like a brain area that absorbs and reflects the abstract rules needed to guide goal-directed, volitional behavior.

Based on this evidence, Miller and Cohen (2001) argued that the cardinal PFC function is to acquire and actively maintain patterns of activity that represent goals and the means to achieve them ("rules") *and* the cortical pathways needed to perform the task ("maps"—together, "rulemaps") [Fig. 18–2]. Under this model, activation of a PFC rulemap sets up bias signals that propagate throughout much of the rest of the cortex, affecting sensory systems as well as systems responsible for response execution, memory retrieval, and emotional evaluation. The aggregate effect is to guide the flow of neural activity along pathways that establish the proper mappings between inputs, internal states, and outputs to best perform the task. Establishing the proper mapping is especially important whenever stimuli are ambiguous (i.e., they activate more

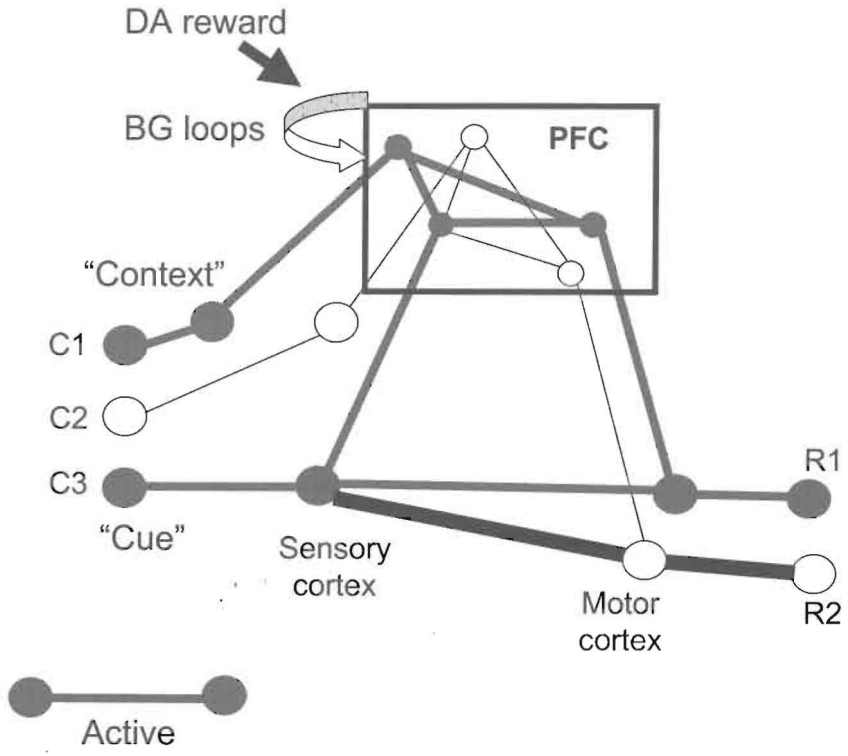

Figure 18–2 Schematic diagram illustrating the suggested role for the prefrontal cortex (PFC) in cognitive control adapted from Miller and Cohen (2001). Shown are processing units representing cues, such as sensory inputs, current motivational state, memories, and so on (C1, C2, and C3), and those representing two voluntary actions (e.g., "responses" R1 and R2). Also shown are internal, or "hidden," units that represent more central stages of processing. The PFC is not heavily connected with primary sensory or motor cortices, but instead is connected with higher-level "association" and premotor cortices. Via interactions with the basal ganglia (BG) [see text], dopaminergic (DA) reward signals foster the formation of a task model, a neural representation that reflects the learned associations between task-relevant information (as shown by the recursive arrow). A subset of the information (e.g., C1 and C2) can then evoke the entire model, including information about the appropriate response (e.g., R1). Thus, the PFC can coordinate processing throughout the brain and steer processing away from a prepotent (reflexive) response (C3 to R2) toward a weakly established, but more goal-relevant, response (C3 to R1). Excitatory signals from the PFC feed back to other brain systems to enable task-relevant neural pathways. *Thick lines* indicate well-established pathways mediating a prepotent behavior. *Solid circles* indicates active units or pathways.

than one input representation), or when multiple responses are possible and the task-appropriate response must compete with stronger, more habitual alternatives. In short, task information is acquired by the PFC, which provides support to related information in posterior brain systems, effectively acting as a global attentional controller.

However, as noted earlier, the PFC is heavily interconnected and does not work in isolation. Later in the chapter, we will review evidence that the PFC works in close collaboration with the BG in the learning of goal-directed behaviors. Specifically, we will argue that, through reciprocal connections between the PFC and BG, increasingly complex rules can be constructed.

CONCRETE RULES AND THE BASAL GANGLIA

The BG is a collection of subcortical nuclei that, similar to the PFC, have a high degree of cortical convergence. Cortical inputs arrive largely via the striatum (which includes both the caudate and the putamen); are processed through the globus pallidus, the subthalamic nucleus (STN), and the substantia nigra; and are then directed back into the cortex via the thalamus (Fig. 18–3). Although the PFC is believed to be involved in the creation and implementation of abstract rules, the BG is believed to be involved in the formation of concrete habits. We will review some anatomical and physiological evidence in support of this theory.

Early evidence about the function of the BG came from human patients with damage or dysfunction to this area. For example, both Parkinson's disease and Huntington's disease cause profound behavioral deficits, ranging from motor (e.g., difficulty initiating volitional movement) to cognitive (e.g., difficulty switching tasks) [Taylor et al., 1986; Cronin-Golomb et al., 1994; Lawrence et al., 1998]. Animal models of lesions of the striatum produce impairments in learning new operant behaviors (or concrete rules) and show that damage to different parts of the striatum generally causes deficits similar to those caused by lesions of the area of the cortex that loop with the affected region of the striatum (Divac et al., 1967; Goldman and Rosvold, 1972). For example, lesions of the regions of the caudate associated with the frontal cortex result in cognitive impairments, suggesting that the reciprocal connections between the BG and cortex play a significant role in the functioning of that cortical area.

Projections from the striatum are distributed along two parallel routes: the "direct" and "indirect" pathways (Fig. 18–3) [Mink, 1996; Graybiel, 2000]. The direct pathway leads from the striatum into the globus pallidus internal (GPi) and the substantia nigra pars reticulata (SNpr). These regions directly project onto the thalamus. All projections from the striatum release gamma-aminobutyric acid (GABA); therefore, they inhibit downstream neurons in the GPi/SNpr. Neurons in the GPi/SNpr inhibit the thalamus, making the direct pathway effectively excitatory—activity in the striatum releases inhibition on the thalamus. The indirect pathway involves striatal projections to the globus pallidus external (GPe), which in turn, projects to the STN, which

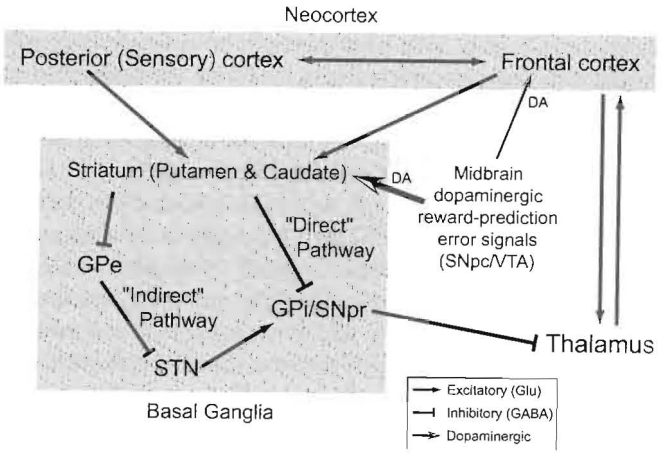

Figure 18–3 Simplified circuit diagram for the basal ganglia, illustrating the loops it makes with the frontal cortex. See text for explanation. The *heavy arrow* illustrates the much heavier projection of midbrain dopaminergic neurons to the striatum than to the cortex. DA, dopamine; GPe, globus pallidus external; STN, subthalamic nucleus; GPi, globus pallidus internal; SNpr, substantia nigra pars reticulata; SNpc, substantia nigra, pars compacta; VTA, ventral tegmental area; Glu, glutamate; GABA, gamma-aminobutyric acid.

projects onto the GPi/SNpr. Similar to the other connections in the BG, GPe inputs into the STN are inhibitory, but the STN provides glutamatergic, excitatory input into the GPi/SNpr. Due to the added inhibitory synapse, the indirect pathway increases inhibition on the thalamus. These two pathways are believed to exist in an equilibrium that allows for the release of desired patterns, while inhibiting unintended ones. Although cortical inputs into the striatum had a divergent nature, connections between the striatum and the GPi/SNpr and GPe are believed to be highly convergent (Flaherty and Graybiel, 1993, 1994; Parent and Hazrati, 1993). This convergence of inputs is effectively a reduction in the dimensionality of the patterns, and may allow for a certain degree of integration and generalization across specific cortical inputs.

Similar to the PFC, the structure of the BG is ideal for integrating information. Most of the cortex projects directly onto the striatum (Kemp and Powell, 1970; Kitai et al., 1976) in a divergent manner, so that cortical afferents make connections to multiple striatal neurons (Flaherty and Graybiel, 1991). The striatum is believed to be subdivided into striosomes and the matrix (Graybiel and Ragsdale, 1978), with striosomes preferentially receiving inputs from the entire cerebral cortex and the matrix primarily receiving inputs from the limbic and hippocampal systems and from the PFC (Donoghue and Herkenham, 1986; Gerfen, 1992; Eblen and Graybiel, 1995). Anatomical tracing techniques have suggested that functionally similar cortical areas project

into the same striosome (Yeterian and Van Hoesen, 1978; Van Hoesen et al., 1981; Flaherty and Graybiel, 1991). For example, both sensory and motor areas relating to the arm seem to preferentially innervate the same striosome. The segregated nature of BG inputs are maintained throughout the different nuclei such that the output from the BG (via the thalamus) is largely to the same cortical areas that gave rise to the initial inputs into the BG (Selemon and Goldman-Rakic, 1985; Parasarathy et al., 1992). Additionally, the frontal cortex receives the largest portion of BG outputs, suggesting a close collaboration between these structures (Middleton and Strick, 1994, 2000, 2002).

The majority of neurons found in both the striosome and the matrix are spiny cells (as high as 90%) [Kemp and Powell, 1971]. These neurons are so named for the high density of synaptic boutons along their dendritic arbor, due to the convergent nature of cortical inputs. Along with the cortical inputs, spiny cells receive a strong dopaminergic (DA) input from neurons in the midbrain. These DA neurons have been suggested to provide a reward-based "teaching signal" that gates plasticity in the striatum. All of this has suggested that the striatum has an ideal infrastructure for rapid, supervised learning (i.e., the quick formation of connections between cortical inputs that predict reward). This is exactly the type of learning that supports the imprinting of specific stimulus-response pairing that supports concrete rules. Finally, it is important to note that there are functional and anatomical differences between the dorsal and ventral striatum. The dorsal striatum is more associated with the PFC and the stimulus-response-reward learning that is the subject of this chapter. The ventral striatum is more connected with the sensory cortex and seems to be more involved in learning the reward value of stimuli (see O'Doherty et al., 2004).

DOPAMINERGIC TEACHING SIGNALS

The formation of rules requires guidance. Concrete rules are formed, through feedback, to actively bind neural representations that lead to reward and break associations that are ineffective. This direct form of plasticity can pair coactivated neurons to form specific rules and predictions. Abstract rules are also guided by feedback so that relevant events and predictive relationships can be distinguished from spurious coincidences. Although the form of plasticity is different for concrete and abstract rules, both need be guided by information about which associations are predictive of desirable outcomes. This guidance appears to come in the form of a "reinforcement signal" and is suggested to be provided by DA neurons in the midbrain.

Dopaminergic neurons are located in both the ventral tegmental area and the substantia nigra, pars compacta (Schultz et al., 1992, 1997; Schultz, 1998), and show activity that directly corresponds to the reward prediction error signals suggested by models of animal learning. These neurons increase activity whenever the animal receives an unexpected reward and will reduce activity if an expected reward is withheld. When active, these neurons release dopamine

onto downstream targets. Dopamine is a neuromodulator that has been suggested to regulate plasticity at the innervated site.

Midbrain DA neurons send heavy projections into both the frontal cortex and the striatum. The projections into the frontal cortex show a gradient connectivity with heavier inputs anteriorly that drop off posteriorly, suggesting a preferential input of reward information into the PFC (Thierry et al., 1973; Goldman-Rakic et al., 1989). However, the midbrain input of DA into the striatum is much heavier than that of the PFC, by as much as an order of magnitude (Lynd-Balta and Haber, 1994). Furthermore, recent evidence suggests that neither strengthening nor weakening of synapses in the striatum by long-term depression or potentiation can occur without DA input (Calabresi et al., 1992, 1997; Otani et al., 1998; Kerr and Wickens, 2001).

After training, DA neurons in the midbrain will learn to increase activity to an unexpected stimulus that directly predicts a reward: The event "stands in" for the reward (Schultz et al., 1993). DA neurons will now respond to the predictive event when it is unexpected, but will no longer respond to the actual, now expected, reward event. In short, the activity of these neurons seems to correspond to a teaching signal that says, "Something good happened and you did not predict it, so remember what just happened so you can predict it in the future." Alternatively, if a reward is expected, but not received, the signal provides feedback that whatever behavior was just taken is not effective in getting rewarded. If these reward signals affect connections within the PFC and BG that were recently active, and therefore likely involved in recent behavior, then the result may be to help to strengthen reward-predicting associations within the network, while reducing associations that do not increase benefits. In this way, the brain can learn what rules are effective in increasing desirable outcomes.

"FAST," SUPERVISED BASAL GANGLIA PLASTICITY VERSUS "SLOWER," LESS SUPERVISED CORTICAL PLASTICITY

One might expect that the greatest evolutionary benefit would be gained from learning as quickly as possible, and there are obvious advantages to learning quickly—adapting at a faster rate than competing organisms lends a definite edge, whereas missed opportunities can be costly (even deadly). However, there are also disadvantages to learning quickly because one loses the ability to integrate across multiple experiences to form a generalized, less error-prone prediction. Take the classic example of one-trial learning: conditioned taste aversion. Many of us have had the experience of eating a particular food and then becoming ill for an unrelated reason. However, in many cases, the person develops an aversion to that food, even though the attribution is erroneous. Extending learning across multiple episodes allows organisms to detect the regularities of predictive relationships and leave behind spurious associations and coincidences. In addition to avoiding errors, slower, more deliberate learning also provides the opportunity to integrate associations across many different experiences to detect common structures.

It is these regularities and commonalities across specific instances that form abstractions, general principles, concepts, and symbolisms that are the medium of the sophisticated, "big-picture" thought needed for truly long-term goals. Indeed, this is fundamental to proactive thought and action. Generalizing among many past experiences gives us the ability to generalize to the future, to imagine possibilities that we have not yet experienced—but would like to—and given the generalized rules, we can predict the actions and behaviors needed to achieve our goal. In addition, abstraction may aid in cognitive flexibility, because generalized representations are, by definition, concise because they lack the details of the more specific representations. Based on the compressed representations, it is probably easier to switch between, and maintain, multiple generalized representations within a given network than to switch between representations when they are elaborate and detailed.

Networks that learn at a slower rate also tend to be more stable. It is believed that fast versus slow learning correlates with large versus small changes in synaptic weights, respectively. Artificial neural networks with small changes in synaptic weights at each learning episode converge very slowly, whereas large synaptic weight changes can quickly capture some patterns, the resulting networks tend to be more volatile and exhibit erratic behavior. This is due to the fact that a high learning rate can overshoot minima in the error function, even oscillating between values on either side of the minima, but never reaching the minima (for more information on artificial neural networks, see Hertz et al., 1991; Dayan and Abbott, 2001).

Given the advantages and disadvantages associated with both forms of learning, the brain must balance the obvious pressure to learn as quickly as possible with the advantages of slower learning. One possible solution to this conundrum comes from O'Reilly and colleagues, who suggested that fast learning and slow learning systems interact with one another (McClelland et al., 1995; O'Reilly and Munakata, 2000). Studying the consolidation of long-term memories, McClelland et al. (1995) specifically suggested that fast plasticity mechanisms within the hippocampus are able to quickly capture new memories while "training" the slower-learning cortical networks. In this way, the brain is able to balance the need to initially grasp new memories with the advantages of a generalized, distributed representation of long-term memories. The idea is that the hippocampus is specialized for the rapid acquisition of new information; each learning trial produces large weight changes. The output of the hippocampus will then repeatedly activate cortical networks that have smaller weight changes per episode. Continued hippocampal-mediated reactivation of cortical representations allows the cortex to gradually connect these representations with other experiences. That way, the shared structure across experiences can be detected and stored, and the memory can be interleaved with others so that it can be readily accessed.

We propose that a similar relationship exists between the PFC and BG. A recent experiment by our laboratory provides suggestive evidence (Pasupathy and Miller, 2005) [see Fig. 18–4]. Monkeys were trained to associate a visual

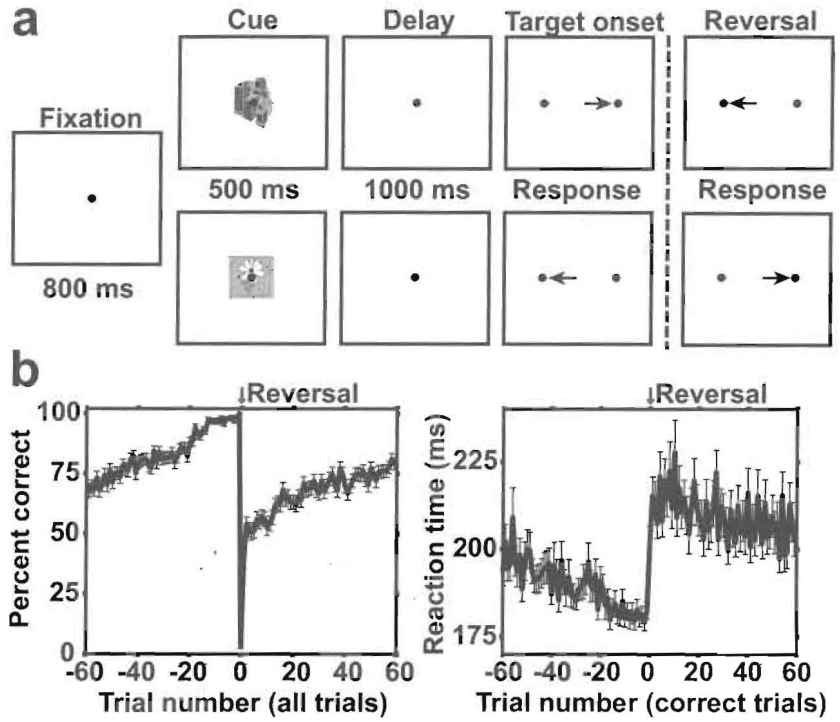

Figure 18–4 *A.* One of two initially novel cues was briefly presented at the center of gaze, followed by a memory delay and then presentation of two target spots on the right and left. Saccade to the target associated with the cue at that time was rewarded (as indicated by arrow). After this was learned, the cue-saccade associations were reversed and relearned. *B.* Average percentage of correct performance on all trials (*left*) and average reaction time on correct trials (*right*) across sessions and blocks as a function of trial number during learning for two monkeys. Zero (*downward arrow*) represents the first trial after reversal. *Error bars* show standard error of the mean.

cue with a directional eye movement over a period of trials (Fig. 18–4*A*). Once performance reached criterion and plateaued, the stimulus-response associations were reversed and the animals were required to relearn the pairings (Fig 18–4*B*). During the task, single neurons were recorded in both the PFC and the BG to determine the selectivity for the cue-direction association in each area. Over the period of a few tens of trials, the animals quickly learned the new cue-direction pairing (Fig 18–4*B*), and selectivity in both the striatum and PFC increased. As can be seen in Figure 18–5*A*, neural activity in the striatum showed rapid, almost bistable, changes in the timing of selectivity. This is in contrast to the PFC, where changes were much slower, with selective responses slowly advancing across trials (Fig 18–5*B*). Interestingly, however, the slower PFC seemed to be the final arbiter of behavior; the monkeys'

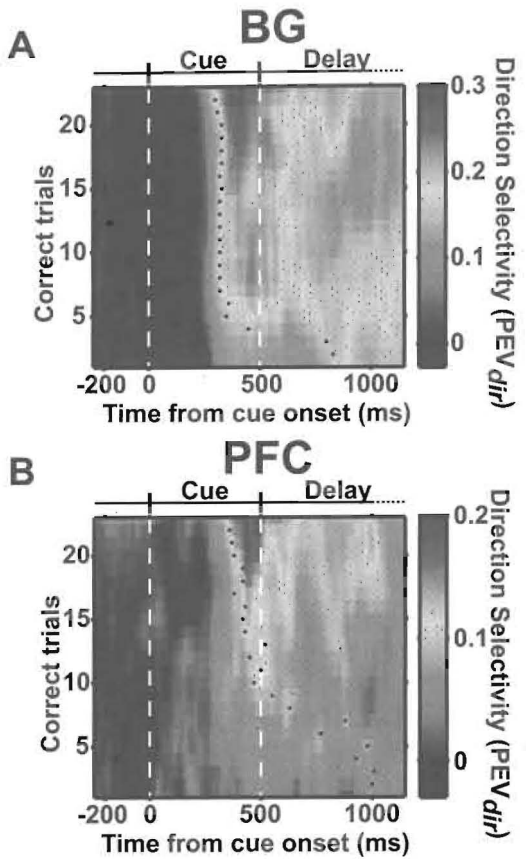

Figure 18–5 *A and B.* Selectivity for the direction of eye movement associated with the presented cue. Selectivity was measured as the percent of explained variance by direction (PEV_{dir}), and is shown in the color gradient across time for both the basal ganglia (BG) [*A*], and prefrontal cortex (PFC) [*B*]. *Black dots* show the time of rise, as measured by the time to half-peak.

improvement in selecting the correct response more closely matched the timing of PFC changes than striatum changes.

These results may reflect a relationship between the BG and PFC that is similar to the relationship between the hippocampus and cortex, as suggested by O'Reilly. As the animals learned specific stimulus-response associations, these changes are quickly represented in the BG, which in turn, slowly trains the PFC. In this case, the fast plasticity in the striatum (strong weight changes) is better suited to the rapid formation of concrete rules, such as the associations between a specific cue and response. However, as noted earlier, fast

learning tends to be error-prone, and indeed, striatal neurons began predicting the forthcoming behavioral response early in learning, when that response was often wrong. By contrast, the smaller weight changes in the PFC may have allowed it to accumulate more evidence and arrive at the correct answer more slowly and judiciously. Interestingly, during this task, behavior more closely reflected the changes in the PFC, possibly due to the fact that the animals were not under enough pressure to change its behavior faster, choosing instead the more judicious path of following the PFC.

The faster learning-related changes in the striatum reported by Pasupathy and Miller (2005) are consistent with our hypothesis that there is stronger modulation of activity in the striatum than in the PFC during performance of these specific, concrete rules. But what about abstracted, generalized rules? Our model of fast BG plasticity versus slower PFC plasticity predicts the opposite, namely, that abstract rules should have a stronger effect on PFC activity than on BG activity because the slower PFC plasticity is more suited to this type of learning. A recent experiment by Muhammad et al. (2006) showed just that. Building on the work of Wallis et al. (2001), in this experiment, monkeys were trained to apply the abstract rules "same" and "different" to pairs of pictures. If the "same" rule was in effect, monkeys responded if the pictures were identical, whereas if the "different" rule was in effect, monkeys responded if the pictures were different. The rules were abstract because the monkeys were able to apply the rules to novel stimuli—stimuli for which there could be no pre-existing stimulus-response association. This is the definition of an abstract rule. Muhammad et al. (2006) recorded neural activity from the same PFC and striatal regions as Pasupathy and Miller (2005), and found that, in contrast to the specific-cue response associations, the abstract rules were reflected more strongly in PFC activity (more neurons with effects and larger effects) than in BG activity, the opposite of what Pasupathy and Miller (2005) reported for the specific cue-response associations.

In fact, this architecture (fast learning in more primitive, noncortical structures training the slower, more advanced, cortex) may be a general brain strategy; in addition to being suggested for the relationship between the hippocampus and cortex, it has also been proposed for the cerebellum and cortex (Houk and Wise, 1995). This makes sense: The first evolutionary pressure on our cortex-less ancestors was presumably toward faster learning, whereas only later did we add on a slower, more judicious and flexible cortex. These different styles of plasticity in the striatum versus PFC might also be suited to acquiring different types of information beyond the distinction between concrete and abstract discussed so far. This is illustrated in a recent proposal by Daw et al. (2005).

THE PREFRONTAL CORTEX AND STRIATUM: MODEL-BUILDING VERSUS "SNAPSHOTS"

Daw et al. (2005) proposed functional specializations for the PFC and BG (specifically, the striatum) that may be in line with our suggestions. They

suggested that the PFC builds models of an entire behavior—it retains information about the overall structure of the task, following the whole course of action from initial state to ultimate outcome. They liken this to a "tree" structure for a typical operant task: Behaviors begin in an initial state, with two or more possible response alternatives. Choosing one response leads to another state, with new response alternatives, and this process continues throughout the task, ultimately leading to a reward. The PFC is able to capture this entire "tree" structure, essentially providing the animal with an internal model of the task. By contrast, the striatum is believed to learn the task piecemeal, with each state's response alternatives individually captured and separate from the others. This "caching reinforcement learning" system retains information about which alternative is "better" in each state, but nothing about the overall structure of the task (i.e., the whole "tree").

This is believed to explain observations of tasks that use reinforcer devaluation. In such tasks, you change the value of the reward by saturating the animal on a given reward (e.g., overfeeding on chocolate if chocolate is a reward in that task). This has revealed two classes of behavior. Behaviors that are affected by reinforcer devaluation are considered goal-directed because changing the goal changes the behavior. As mentioned earlier, goal-directed behaviors depend on the PFC. By contrast, overlearned behaviors whose outcomes remain relatively constant can become habits, impervious to reinforcer devaluation. Because these behaviors are not affected by changing the goal, they seem to reflect control by a caching system in which the propensity for a given alternative in each situation is stored independently of information about past or future events (states). Habits have long been considered a specialization of the BG. Daw et al. (2005) proposed that there is arbitration between each system based on uncertainty; whichever system is most accurate is the one deployed to control behavior.

We believe that this maps well onto our notion of the fast, supervised, BG plasticity versus slow, more-Hebbian, PFC plasticity. Fast plasticity, such as the nearly bistable changes that Pasupathy and Miller (2005) observed in the striatum, would seem ideal for learning the reinforcement-related snapshots that capture the immediate circumstances and identify which alternative is preferable for a particular state. The slow plasticity in the PFC seems more suited for the linking in of additional information about past states that is needed to learn and retain an entire model of the task and thus predict future states.

The interactions of these systems might explain several aspects of goal-directed learning and habit formation. The initial learning of a complex operant task invariably begins with the establishment of a simple response immediately proximal to reward (i.e., a single state). Then, as the task becomes increasingly complex as more and more antecedents and qualifications (states and alternatives) are linked in, the PFC shows greater involvement. It facilitates this learning via its slower plasticity, allowing it to stitch together the relationships between the different states. This is useful because uncertainty about the

correct action in a given state adds up across many states in a complex task. Thus, in complex tasks, the ability of the reinforcement to control behavior would be lessened with the addition of more and more states. However, model-building in the PFC may provide the overarching infrastructure—the thread weaving between states—that facilitates learning of the entire course of action. This may also explain why, when complex behaviors are first learned, they are affected by reinforcer devaluation and susceptible to disruption by PFC damage. Many tasks will remain dependent on the PFC and the models it builds, especially those requiring flexibility (e.g., when the goal often changes or there are multiple goals to choose among), or when a strongly established behavior in one of the states (e.g., a habit) is incompatible with the course of action needed to obtain a specific goal. However, if a behavior, even a complex one, is unchanging, then all of the values of each alternative at each juncture are constant, and once these values are learned, control can revert to a piecemeal caching system in the BG. That is, the behavior becomes a "habit," and it frees up the more cognitive PFC model-building system for behaviors requiring the flexibility it provides.

Note that this suggests that slower plasticity in the PFC might sometimes support relatively fast learning on the behavioral level (i.e., faster than relying on the BG alone) because it is well suited to learning a complex task. This distinction is important, because thus far, we have been guilty of confusing learning on the neuronal level and learning on the behavioral level. Although it is true that small changes in synaptic weights might often lead to slow changes in behavior and vice versa, this is too simplistic. Certain tasks might be learned better and faster through the generalized, model-based learning seen in the PFC than through the strict, supervised learning observed in the striatum.

RECURSIVE PROCESSING AND BOOTSTRAPPING IN CORTICO-GANGLIA LOOPS

"Bootstrapping" is the process of building increasingly complex representations from simpler ones. The recursive nature of the anatomical loops between the BG and PFC may lend itself to this process. As described earlier, anatomical connections between the PFC and BG seem to suggest a closed loop—channels within the BG return outputs, via the thalamus, into the same cortical areas that gave rise to their initial cortical input. This recursive structure in the anatomy may allow for learned associations from one instance to be fed back through the loop for further processing and learning. In this manner, new experiences can be added onto previous ones, linking in more and more information to build a generalized representation. This may allow the bootstrapping of neural representations to increasing complexity, and with the slower learning in the PFC, greater abstractions.

A hallmark of human intelligence is the propensity for us to ground new concepts in familiar ones because it seems to ease our understanding of novel ideas. For example, we learn to multiply through serial addition and we begin

to understand quantum mechanisms through analogies to waves and particles. The recursive interactions between the BG and PFC may support this type of cognitive bootstrapping—initial, simple associations (or concrete rules) are made in the BG and fed back into the PFC. This feedback changes the representation of the original association in the PFC, helping to encode the concrete rule in both the BG and PFC. Additional concrete associations through different experiences can also be made and modified in a similar manner. The associative nature of the PFC will begin to bind across experiences, finding similarities in both the cortical inputs into the PFC as well as the looped inputs from the BG. This additional generalization is the basis for the formation of abstract rules based on the concrete rules that are first learned in the BG. As this process continues, new experiences begin to look "familiar" to the PFC, and a more generalized representation of a specific instance can be constructed. This generalized representation can now be looped through the BG to make reliable predictions of associations based on previously learned concrete rules.

Reward processing is a specific instance where recursive processing might provide the framework necessary for the observed neuronal behavior. As previously described, midbrain DA neurons respond to earlier and earlier events in a predictive chain leading to a reward. Both the frontal cortex and the striatum send projections into the midbrain DA neurons, possibly underlying their ability to bootstrap to early predictors of reward. However, although this is suggestive, it is still unknown whether these descending projections are critical for this behavior.

Additionally, the PFC-BG loops suggest an autoassociative type of network, similar to that seen in the CA3 of the hippocampus. The outputs looping back on the inputs allow the network to learn to complete (i.e., recall) previously learned patterns, given a degraded version or a subset of the original inputs (Hopfield, 1982). In the hippocampus, this network has been suggested to play a role in the formation of memories; however, BG-PFC loops are heavily influenced by DA inputs, and therefore may be more goal-oriented.

An intriguing feature of autoassociative networks is their ability to learn temporal sequences of patterns and thus make predictions. This feature relies on feedback of the activity pattern into the network with a temporal delay, allowing the next pattern in the sequence to arrive as the previous pattern is fed back, building an association between the two (Kleinfeld, 1986; Sompolinsky and Kanter, 1986).

The PFC-BG loops have two mechanisms by which to add this lag in feedback. One possibility is through the use of inhibitory synapses, which are known to have a slower time constant than excitatory ones. The "direct" pathway has two inhibitory synapses, the result being a net excitatory effect on the cortex via disinhibition of the thalamus, whereas the "indirect" one has three inhibitory synapses, making it net inhibitory. These two pathways are believed to exist in balance—activity in the indirect pathway countermands current processing in the direct loop. But why evolve a loop out of inhibitory synapses? First, it can prevent runaway excitation and thus allow greater control

over processing (Wong et al., 1986; Connors et al., 1988; Wells et al., 2000), but it is also possible that inhibitory synapses are used to slow the circulation of activity through the loops and allow for the binding of temporal sequences. Many inhibitory synapses are mediated by potassium channels with slow time courses (Couve et al., 2000). A second way to add lag to the recursion is through a memory buffer. The PFC is well known for this type of property; its neurons can sustain their activity to bridge short-term memory delays. This can act as a bridge for learning contingencies across several seconds, or even minutes. The introduction of lag into the recursive loop through either mechanism (or both) may be enough to tune the network for sequencing and prediction.

After training, a lagged autoassociative network that is given an input will produce, or predict, the next pattern in the sequence. This is a fundamentally important feature for producing goal-directed behaviors, especially as they typically extend over time. Experimental evidence for the role of the BG in sequencing and prediction comes from neurophysiological observations that striatal neural activity reflects forthcoming events in a behavioral task (Jog et al., 1999) and that lesions of the striatum can cause a deficit in producing learned sequences (Miyachi et al., 1997; Bailey and Mair, 2006).

SUMMARY: FRONTAL CORTICAL–BASAL GANGLIA LOOPS CONSTRUCT ABSTRACT RULES FOR COGNITIVE CONTROL

In this chapter, we have proposed that the learning of abstract rules occur through recursive loops between the PFC and BG. The learning of concrete rules, such as simple stimulus-response associations, is more a function of the BG, which—based on anatomical and physiological evidence—is specialized for the detection and storage of specific experiences that lead to reward. In contrast, abstract rules are better learned slowly, across many experiences, in the PFC. The recursive anatomical loops between these two areas suggest that the fast, error-prone learning in the BG can help train the slower, more reliable, frontal cortex. Bootstrapping from specific instances and concrete rules represented and stored in the BG, the PFC can construct abstract rules that are more concise, more predictive, and more broadly applicable; it can also build overarching models that capture an entire course of action. Note that we are not suggesting that there is serial learning between the BG and PFC; we are not suggesting that the BG first learns a task and then passes it to the PFC. Goal-directed learning instead depends on a highly interactive and iterative processing between these structures, working together and in parallel to acquire the goal-relevant information.

The result of this learning can be thought of as creating a "rulemap" in the PFC that is able to capture the relationships between the thoughts and actions necessary to successfully achieve one's goals in terms of which cortical pathways are needed (Miller and Cohen, 2001) [see Fig. 18–2]. The appropriate rulemap can be activated when cognitive control is needed: in situations in

which the mapping between sensory inputs, thoughts, and actions either is weakly established relative to other existing ones or is rapidly changing. Activation of the PFC rulemaps establishes top-down signals that feed back to most of the rest of the cortex, dynamically modulating information flow through the brain to best regulate important information and generate appropriate goal-directed thoughts and actions.

ACKNOWLEDGMENTS This work was supported by the National Institute of Mental Health, the National Institute of Neurological Disorders and Stroke, and the RIKEN-MIT Neuroscience Research Center. The authors thank M. Wicherski for helpful comments.

REFERENCES

Amaral DG (1986) Amygdalohippocampal and amygdalocortical projections in the primate brain. Advances in Experimental Medicine and Biology 203:3–17.

Amaral DG, Price JL (1984) Amygdalo-cortical projections in the monkey (*Macaca fascicularis*). Journal of Comparative Neurology 230:465–496.

Asaad WF, Rainer G, Miller EK (2000) Task-specific activity in the primate prefrontal cortex. Journal of Neurophysiology 84:451–459.

Bailey KR, Mair RG (2006) The role of striatum in initiation and execution of learned action sequences in rats. Journal of Neuroscience 26:1016–1025.

Barbas H, De Olmos J (1990) Projections from the amygdala to basoventral and mediodorsal prefrontal regions in the rhesus monkey. Journal of Comparative Neurology 300:549–571.

Barbas H, Henion TH, Dermon CR (1991) Diverse thalamic projections to the prefrontal cortex in the rhesus monkey. Journal of Comparative Neurology 313:65–94.

Barbas H, Pandya DN (1989) Architecture and intrinsic connections of the prefrontal cortex in the rhesus monkey. Journal of Comparative Neurology 286:353–375.

Calabresi P, Maj R, Pisani A, Mercuri NB, Bernardi G (1992) Long-term synaptic depression in the striatum: physiological and pharmacological characterization. Journal of Neuroscience 12:4224–4233.

Calabresi P, Saiardi A, Pisani A, Baik JH, Centonze D, Mercuri NB, Bernardi G, Borrelli E, Maj R (1997) Abnormal synaptic plasticity in the striatum of mice lacking dopamine D2 receptors. Journal of Neuroscience 17:4536–4544.

Connors B, Malenka R, Silva L (1988) Two inhibitory postsynaptic potentials, and GABAA and GABAB receptor-mediated responses in neocortex of the rat and cat. Journal of Physiology (London) 406:443–468.

Couve A, Moss SJ, Pangalos MN (2000) GABAB receptors: a new paradigm in G protein signaling. Molecular and Cellular Neuroscience 16:296–312.

Cronin-Golomb A, Corkin S, Growdon JH (1994) Impaired problem solving in Parkinson's disease: impact of a set-shifting deficit. Neuropsychologia 32:579–593.

Croxson PL, Johansen-Berg H, Behrens TEJ, Robson MD, Pinsk MA, Gross CG, Richter W, Richter MC, Kastner S, Rushworth MFS (2005) Quantitative investigation of connections of the prefrontal cortex in the human and macaque using probabilistic diffusion tractography. Journal of Neuroscience 25:8854–8866.

Dayan P, Abbott L (2001) Theoretical neuroscience: computational and mathematical modeling of neural systems. Cambridge: MIT Press.

Daw ND, Niv Y, Dayan P (2005) Uncertainty-based competition between the prefrontal and dorsolateral striatal systems for behavioral control. Nature Neuroscience 8:1704–1711.

Divac I, Rosvold HE, Szwarcbart MK (1967) Behavioral effects of selective ablation of the caudate nucleus. Journal of Comparative and Physiological Psychology 63:184–190.

Donoghue JP, Herkenham M (1986) Neostriatal projections from individual cortical fields conform to histochemically distinct striatal compartments in the rat. Brain Research 365:397–403.

Eblen F, Graybiel A (1995) Highly restricted origin of prefrontal cortical inputs to striosomes in the macaque monkey. Journal of Neuroscience 15:5999–6013.

Flaherty AW, Graybiel AM (1991) Corticostriatal transformations in the primate somatosensory system: projections from physiologically mapped body-part representations. Journal of Neurophysiology 66:1249–1263.

Flaherty A, Graybiel A (1993) Output architecture of the primate putamen. Journal of Neuroscience 13:3222–3237.

Flaherty A, Graybiel A (1994) Input-output organization of the sensorimotor striatum in the squirrel monkey. Journal of Neuroscience 14:599–610.

Funahashi S, Bruce CJ, Goldman-Rakic PS (1989) Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex. Journal of Neurophysiology 61:331–349.

Fuster JM (1973) Unit activity in prefrontal cortex during delayed-response performance: neuronal correlates of transient memory. Journal of Neurophysiology 36: 61–78.

Fuster JM (1995) Memory in the cerebral cortex. Cambridge: MIT Press.

Fuster JM, Alexander GE (1971) Neuron activity related to short-term memory. Science 173:652–654.

Gerfen CR (1992) The neostriatal mosaic: multiple levels of compartmental organization. Trends in Neurosciences 15:133–139.

Goldman PS, Rosvold HE (1972) The effects of selective caudate lesions in infant and juvenile rhesus monkeys. Brain Research 43:53–66.

Goldman-Rakic PS, Leranth C, Williams SM, Mons N, Geffard M (1989) Dopamine synaptic complex with pyramidal neurons in primate cerebral cortex. Proceedings of the National Academy of Sciences U S A 86:9015–9019.

Graybiel AM (2000) The basal ganglia. Current Biology 10:R509–R511.

Graybiel AM, Ragsdale CW Jr (1978) Histochemically distinct compartments in the striatum of humans, monkeys, and cats demonstrated by acetylthiocholinesterase staining. Proceedings of the National Academy of Sciences U S A 75:5723–5726.

Haier RJ, Jung RE, Yeo RA, Head K, Alkire MT (2004) Structural brain variation and general intelligence. Neuroimage 23:425–433.

Hertz J, Palmer R, Krogh A (1991) Introduction to the theory of neural computation. Redwood City, Calif.: Addison-Wesley Pub. Co.

Hopfield JJ (1982) Neural networks and physical systems with emergent collective computational abilities. Proceedings of the National Academy of Sciences U S A 79: 2554–2558.

Houk JC, Wise SP (1995) Distributed modular architectures linking the basal ganglia, cerebellum, and cerebral cortex: their role in planning and controlling action. Cerebral Cortex 5:95–110.

Jog MS, Kubota Y, Connolly CI, Hillegaart V, Graybiel AM (1999) Building neural representations of habits. Science 286:1745–1749.

Kemp JM, Powell TP (1970) The cortico-striate projection in the monkey. Brain 93: 525–546.
Kemp JM, Powell TP (1971) The structure of the caudate nucleus of the cat: light and electron microscopy. Philosophical Transactions of the Royal Society of London B 262:383–401.
Kerr JND, Wickens JR (2001) Dopamine D-1/D-5 receptor activation is required for long-term potentiation in the rat neostriatum in vitro. Journal of Neurophysiology 85:117–124.
Kitai ST, Kocsis JD, Preston RJ, Sugimori M (1976) Monosynaptic inputs to caudate neurons identified by intracellular injection of horseradish peroxidase. Brain Research 109:601–606.
Kleinfeld D (1986) Sequential state generation by model neural networks. Proceedings of the National Academy of Science USA 83:9469–9473.
Lawrence AD, Hodges JR, Rosser AE, Kershaw A, French-Constant C, Rubinsztein DC, Robbins TW, Sahakian BJ (1998) Evidence for specific cognitive deficits in preclinical Huntington's disease. Brain 121:1329–1341.
Lynd-Balta E, Haber SN (1994) The organization of midbrain projections to the ventral striatum in the primate. Neuroscience 59:609–623.
Mansouri FA, Matsumoto K, Tanaka K (2006) Prefrontal cell activities related to monkeys' success and failure in adapting to rule changes in a Wisconsin Card Sorting Test analog. Journal of Neuroscience 26:2745–2756.
McClelland J, McNaughton B, O'Reilly R (1995) Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. Psychological Review 102:419–457.
Middleton FA, Strick PL (1994) Anatomical evidence for cerebellar and basal ganglia involvement in higher cognitive function. Science 266:458–461.
Middleton FA, Strick PL (2000) Basal ganglia and cerebellar loops: motor and cognitive circuits. Brain Research Reviews 31:236–250.
Middleton FA, Strick PL (2002) Basal-ganglia 'projections' to the prefrontal cortex of the primate. Cerebral Cortex 12:926–935.
Miller EK (2000) The prefrontal cortex and cognitive control. Nature Reviews Neuroscience 1:59–65.
Miller EK, Cohen JD (2001) An integrative theory of prefrontal function. Annual Review of Neuroscience 24:167–202.
Miller EK, Erickson CA, Desimone R (1996) Neural mechanisms of visual working memory in prefrontal cortex of the macaque. Journal of Neuroscience 16:5154–5167.
Mink J (1996) The basal ganglia: focused selection and inhibition of competing motor programs. Progress in Neurobiology 50:381–425.
Miyachi S, Hikosaka O, Miyashita K, Karadi Z, Rand MK (1997) Differential roles of monkey striatum in learning of sequential hand movement. Experimental Brain Research 115:1–5.
Muhammad R, Wallis JD, Miller EK (2006) A comparison of abstract rules in the prefrontal cortex, premotor cortex, inferior temporal cortex, and striatum. Journal of Cognitive Neuroscience 18:974–989.
Nieder A, Freedman DJ, Miller EK (2002) Representation of the quantity of visual items in the primate prefrontal cortex. Science 297:1708–1711.

O'Doherty J, Dayan P, Schultz J, Deichmann R, Friston K, Dolan RJ (2004) Dissociable roles of ventral and dorsal striatum in instrumental conditioning. Science 16: 452–454.
O'Reilly RC, Munakata Y (2000) Computational explorations in cognitive neuroscience: understanding the mind. Cambridge: MIT Press.
Otani S, Blond O, Desce JM, Crepel F (1998) Dopamine facilitates long-term depression of glutamatergic transmission in rat prefrontal cortex. Neuroscience 85:669–676.
Padoa-Schioppa C, Assad JA (2006) Neurons in the orbitofrontal cortex encode economic value. Nature 41(7090):223–226.
Pandya DN, Barnes CL (1987) Architecture and connections of the frontal lobe. In: The frontal lobes revisited (Perecman E, ed.), pp 41–72. New York: IRBN Press.
Pandya DN, Yeterian EH (1990) The prefrontal cortex in relation to other cortical areas in rhesus monkey: architecture and connections. Progress in Brain Research 85:63–94.
Parasarathy H, Schall J, Graybiel A (1992) Distributed but convergent ordering of corticostriatal projections: analysis of the frontal eye field and the supplementary eye field in the macaque monkey. Journal of Neuroscience 12:4468–4488.
Parent A, Hazrati LN (1993) Anatomical aspects of information processing in primate basal ganglia. Trends in Neurosciences 16:111–116.
Pasupathy A, Miller EK (2005) Different time courses of learning-related activity in the prefrontal cortex and striatum. Nature 433:873–876.
Petrides M, Pandya DN (1999) Dorsolateral prefrontal cortex: comparative cytoarchitectonic analysis in the human and the macaque brain and corticocortical connection patterns. European Journal of Neuroscience 11:1011–1036.
Porrino LJ, Crane AM, Goldman-Rakic PS (1981) Direct and indirect pathways from the amygdala to the frontal lobe in rhesus monkeys. Journal of Comparative Neurology 198:121–136.
Pribram KH, Mishkin M, Rosvold HE, Kaplan SJ (1952) Effects on delayed-response performance of lesions of dorsolateral and ventromedial frontal cortex of baboons. Journal of Comparitive and Physiological Psychology. 45:565–575.
Rainer G, Rao SC, Miller EK (1999) Prospective coding for objects in the primate prefrontal cortex. Journal of Neuroscience 19:5493–5505.
Schultz W (1998) Predictive reward signal of dopamine neurons. Journal of Neurophysiology 80:1–27.
Schultz W, Apicella P, Ljungberg T (1993) Responses of monkey dopamine neurons to reward and conditioned stimuli during successive steps of learning a delayed response task. Journal of Neuroscience 13:900–913.
Schultz W, Apicella P, Scarnati E, Ljungberg T (1992) Neuronal activity in monkey ventral striatum related to the expectation of reward. Journal of Neuroscience 12: 4595–4610.
Schultz W, Dayan P, Montague PR (1997) A neural substrate of prediction and reward. Science 275:1593–1599.
Selemon L, Goldman-Rakic P (1988) Common cortical and subcortical targets of the dorsolateral prefrontal and posterior parietal cortices in the rhesus monkey: evidence for a distributed neural network subserving spatially guided behavior. Journal of Neuroscience 8:4049–4068.
Selemon LD, Goldman-Rakic (1985) Longitudinal topography and interdigitation of corticostriatal projections in the rhesus monkey. Journal of Neuroscience 5:776–794.

Sompolinsky H, Kanter II (1986) Temporal association in asymmetric neural networks. Physical Review Letters 57:2861–2864.

Taylor AE, Saint-Cyr JA, Lang AE (1986) Frontal lobe dysfunction in Parkinson's disease: the cortical focus of neostriatal outflow. Brain 109:845–883.

Thierry AM, Blanc G, Sobel A, Stinus L, Glowinski J (1973) Dopaminergic terminals in the rat cortex. Science 182:499–501.

Van Hoesen GW, Yeterian EH, Lavizzo-Mourey R (1981) Widespread corticostriate projections from temporal cortex of the rhesus monkey. Journal of Comparative Neurology 199:205–219.

Wallis JD, Miller EK (2003) Neuronal activity in the primate dorsolateral and orbital prefrontal cortex during performance of a reward preference task. European Journal of Neuroscience 18:2069–2081.

Wallis JD, Anderson KC, Miller EK (2001) Single neurons in the prefrontal cortex encode abstract rules. Nature 411:953–956.

Watanabe M (1996) Reward expectancy in primate prefrontal neurons. Nature 382: 629–632.

Wells JE, Porter JT, Agmon A (2000) GABAergic inhibition suppresses paroxysmal network activity in the neonatal rodent hippocampus and neocortex. Journal of Neuroscience 20:8822–8830.

White IM, Wise SP (1999) Rule-dependent neuronal activity in the prefrontal cortex. Experimental Brain Research 126:315–335.

Wong RK, Traub RD, Miles R (1986) Cellular basis of neuronal synchrony in epilepsy. Advances in Neurology 44:583–592.

Yeterian EH, Van Hoesen GW (1978) Cortico-striate projections in the rhesus monkey: the organization of certain cortico-caudate connections. Brain Research 139:43–63.