

Conditional Routing of Information to the Cortex: A Model of the Basal Ganglia's Role in Cognitive Coordination

Andrea Stocco, Christian Lebiere, and John R. Anderson
Carnegie Mellon University

The basal ganglia play a central role in cognition and are involved in such general functions as action selection and reinforcement learning. Here, we present a model exploring the hypothesis that the basal ganglia implement a conditional information-routing system. The system directs the transmission of cortical signals between pairs of regions by manipulating separately the selection of sources and destinations of information transfers. We suggest that such a mechanism provides an account for several cognitive functions of the basal ganglia. The model also incorporates a possible mechanism by which subsequent transfers of information control the release of dopamine. This signal is used to produce novel stimulus–response associations by internalizing transferred cortical representations in the striatum. We discuss how the model is related to production systems and cognitive architectures. A series of simulations is presented to illustrate how the model can perform simple stimulus–response tasks, develop automatic behaviors, and provide an account of impairments in Parkinson's and Huntington's diseases.

Keywords: basal ganglia, neural networks, computational modeling, procedural learning

This article puts forward a hypothesis on how the transfer of information among cortical regions is organized. According to this hypothesis, the transfer of information in the brain is primarily directed by the basal ganglia. We present a computational model that has a series of advantages over previous models, including a solution to the information-routing problem that directly maps onto brain physiology.

The organization of information transfer is important for many reasons. First, it is a common problem that every complex model of the brain must ultimately resolve. Improved understanding of individual neural circuits has spawned a number of ambitious attempts to model the basic workings of the brain (Arbib, 2003; Granger, 2006; Hawkins & Blakeslee, 2004; Houk, 2005). Although these attempts differ from each other, they all face the problem of how to control the transfer of information from one processing site to another. This issue is particularly crucial for achieving information integration, because it enables special-

ized processing circuits to access their proper inputs. It is also ultimately decisive in defining cognitive control, because deciding where individual representations are delivered shapes how behavior will be carried out and which actions will be taken.

In the field of cognitive science, *cognitive architectures* are probably the most ambitious examples of general-purpose, integrated cognitive models, aimed at providing a set of primitive functions upon which cognitive behaviors can be built (Anderson, 1983; Newell, 1973). Different architectures have been proposed over the years (Anderson, 1983; Just & Carpenter, 1992; Meyer & Kieras, 1997a, 1997b; Newell, 1990). In a few cases (Anderson, 2007; Anderson, Fincham, Qin, & Stocco, 2008) specific mappings between computational components and brain regions have been developed, giving the architecture a biological substrate and shedding light on the large-scale design of the brain. These modeling attempts solve the problem of how information is transferred in different ways. For instance, Just and Varma (2007) hypothesized that a number of cortical processors exist and share information dynamically with each other. In contrast, Anderson (1983, 2007) and Meyer and Kieras (1997a, 1997b) proposed that information can only pass through a central system.

This article is organized into four sections. The first section provides a brief overview of the basal ganglia, their functional anatomy, their contributions to cognition, and previous modeling attempts of their function. The second section introduces our conditional routing model of the basal ganglia, its functions, and its capabilities. The third section discusses how learning occurs in the model and how these computations are related to a number of cognitive functions like skill acquisition and reward learning. Finally, implications and alternative accounts of basal ganglia function are discussed.

Andrea Stocco, Christian Lebiere, and John R. Anderson, Department of Psychology, Carnegie Mellon University.

This research was supported by Grant R305H030016 from the Institute of Education Sciences to John R. Anderson and Grant FA9550-08-1-0404 from the Air Force Office of Scientific Research to John R. Anderson, Randall C. O'Reilly, and Christian Lebiere. We would like to express our gratitude to Stanislas Dehaene of Collège de France; Jon M. Fincham, Ion Juvina, and Chantel Prat of Carnegie Mellon University; Randall O'Reilly of the University of Colorado; Jelmer Borst and Niels Taatgen of the University of Groningen; and Julie Fiez of the University of Pittsburgh for their help in the development of the model and their suggestions at different stages of the article.

Correspondence concerning this article should be addressed to Andrea Stocco, Department of Psychology, BP 354-Q, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213. E-mail: stocco@cmu.edu

Anatomy of the Basal Ganglia

The basal ganglia (see Figure 1) are a set of subcortical nuclei located in the midbrain, around the thalamus. The major nuclei of the basal ganglia are the striatum, which is composed of the caudate nucleus and the putamen; the internal and external parts of the globus pallidus (henceforth GPi and GPe, respectively); the pars reticulata and the pars compacta of the substantia nigra (SNr and SNc, respectively); and the subthalamic nucleus (STN; Chesselet & Delfs, 1996). Figure 1 illustrates their reciprocal connections.

The striatum is the entry point of the circuit (Albin, Young, & Penney, 1989; Chesselet & Delfs, 1996). Virtually all the cortical mantle projects to the striatum (Carman, Cowan, & Powell, 1963; Kemp & Powell, 1970); these projections reflect cortical topography (Albin et al., 1989; Bolam, Hanley, Booth, & Bevan, 2000; Kemp & Powell, 1970). Other projections to the striatum come from the STN (e.g., Parent & Hazrati, 1995), the ventral tegmental area, and the SNc (e.g., Haber, 2003).

On the other end of the circuit are the SNr and GPi, which constitute the output nuclei of the ganglia. Despite being physically separated by white matter, they are made of similar cells and are often considered as a single entity (Albin et al., 1989; Bolam et al., 2000). This article also considers them a unit and refers to them as the *output nuclei* or SNr/GPi. The main target of these output nuclei, and therefore of the entire circuit, are the ventral anterior, ventral lateral, and the medial dorsal nuclei of the thalamus (Bolam et al., 2000; McFarland & Haber, 2002; A. D. Smith & Bolam, 1990). Most of the known projections originating in the thalamic relay nuclei target the frontal lobe (e.g., Hoover & Strick, 1993; Middleton & Strick, 1994). Smaller projections to temporal (Middleton & Strick, 1996) and parietal areas (Clower, Dum, & Strick, 2005), however, have been described. In addition, the output nuclei also project to the superior colliculus, the reticular

formation, and the pedunculopontine nucleus (Gerfen, Staines, Arbuthnott, & Fibiger, 1982), three subcortical structures that play a substantial role in modulating behavior (Bolam et al., 2000).

Projections to the striatum are organized topologically (Alexander, DeLong, & Strick, 1986; Kemp & Powell, 1970), with afferents from different cortical areas mainly targeting separate areas of the striatum. Thalamo-cortical projections are also organized topologically, with projections to different parts of the cortex originating in specific parts of the thalamus and the output nuclei of the basal ganglia (Middleton & Strick, 2000). This topological organization has initially been thought to reflect the existence of separate parallel circuits (Alexander et al., 1986; Hoover & Strick, 1993). The separation between the circuits has been questioned, and there seem to exist various places of the circuit where multiple inputs from different areas converge (Haber, 2003). Also, experimental evidence suggests that cortical projections, besides targeting their correspondent striatal region, innervate to other parts of the striatum (Parthasarathy, Schall, & Graybiel, 1992).

Most of the projections between the various nuclei of the basal ganglia are inhibitory (see Figure 1). In fact, all the projection neurons in the striatum, in the output nuclei, and in the GPe are GABAergic (Albin et al., 1989; Bolam et al., 2000; A. D. Smith & Bolam, 1990). The output nuclei have elevated tonic activity; therefore, under resting conditions, the net output to the thalamus is inhibitory. The basal ganglia circuit works by modulating this inhibition (Chevalier & Deniau, 1990; DeLong, 1990).

Major Pathways Within the Circuit

There are two main routes that connect the striatum to the output nuclei. One is the *direct* pathway, made of inhibitory projections from the striatum to the SNr/GPi. The other is the *indirect* pathway that proceeds through the GPe and STN (Albin et al., 1989; DeLong, 1990; Penney & Young, 1986). These two pathways have opposite effects on thalamic activity: Whereas the direct pathway disinhibits thalamic activity by inhibiting the output nuclei, the indirect pathway inhibits the thalamic output by exciting the output nuclei (Albin et al., 1989; DeLong, 1990).

A third pathway exists that proceeds directly from the cortex to the STN and from there to the output nuclei. This route is known as the *hyperdirect* pathway (Rouzaire-Dubois & Scarnati, 1987). The importance of this pathway has recently gained prominence. The large number of cortical afferents makes the STN a second input nucleus of the circuit, after the striatum (Mink & Thach, 1993). In addition, cortical signals traveling along the hyperdirect pathway reach the output nuclei faster than do those traveling on either the direct or indirect routes (Nambu et al., 2000). On the basis of these and other findings, Nambu, Tokuno, and Takada (2002) have proposed that the hyperdirect pathway performs a preliminary diffuse inhibition of motor programs, enhancing and sharpening the disinhibitory signals from the direct pathway.

Besides these major routes, a number of smaller descending pathways within the circuit have been identified. For instance, GPe neurons also project directly to the output nuclei (Bevan, Booth, Eaton, & Bolam, 1998; Bolam et al., 2000; see Figure 1).

In addition to pathways that proceed “forward” along the cortex-to-thalamus direction, the circuit includes a number of other pathways that proceed “backward” and may serve feedback and control functions (Bolam et al., 2000; Redgrave, Prescott, & Gurney,

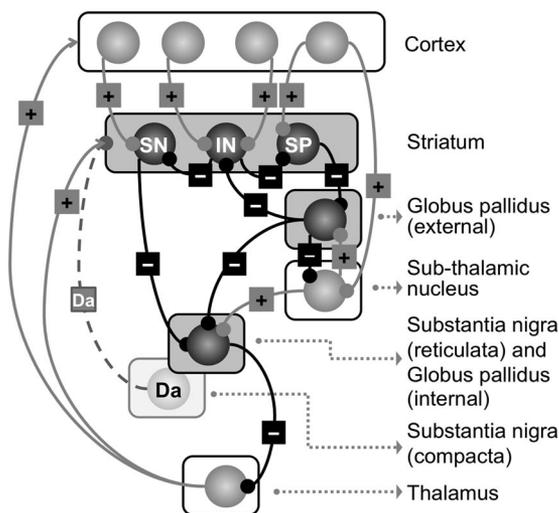


Figure 1. A simplified representation of the basal ganglia circuit and its connectivity (adapted with modifications from Graybiel, 2000, and Bolam et al., 2000). Excitatory projections are marked with a “+”; inhibitory projections with a “-.” Glutamatergic neurons and projections are represented in medium gray; GABAergic, in dark gray; dopaminergic, in light gray. SP = striatopallidal projection neuron; SN = striatonigral projection neuron; IN = striatal interneuron; Da = dopaminergic neuron.

1999). These include projections from the thalamus to the striatum (Y. Smith, Raju, Pare, & Sidibe, 2004); from the STN to the GPe (Shink, Bevan, Bolam, & Smith, 1996); and from the GPe back to the striatum (Bevan et al., 1998). Of particular importance is the nigrostriatal pathway that connects the SNc to the striatum, which is formed by the axons of dopaminergic cells in the SNc (Haber, 2003; Hajos & Greenfield, 1994). Dopamine plays an important regulatory role on striatal neurons (Joel & Weiner, 2000; Schultz, 1998, 2002) and plays a central role in our model as well.

Microcircuitry of the Striatum

The striatum is the largest nucleus of the circuit. To understand its functionality, its internal architecture is as important as its connections with the remaining nuclei. Therefore, this section presents a succinct overview of striatal microcircuitry, which is visually summarized in Figure 2.

Projection Neurons in the Striatum

Most of the striatum is composed of projection neurons. Projection neurons are medium spiny neurons; they account for up to 96% of the striatal cells in rodents (Rymar, Sasseville, Luk, & Sadikot, 2004; Yelnik, Francois, Percheron, & Tande, 1991) and up to 77% in primates (Graveland, Williams, & DiFiglia, 1985). As previously noted, all projection neurons are GABAergic and, therefore, inhibitory (Gerfen, 1992; Kita & Kitai, 1988). They are the main targets of cortical projections, and their axons originate the direct and indirect pathways of the basal ganglia (Tepper &

Bolam, 2004). Projection neurons also make extensive inhibitory connections with each other (Somogyi, Bolam, & Smith, 1981; Wilson & Groves, 1980); the strength and importance of their reciprocal inhibitory connections, however, have been demoted over time (Jaeger, Kita, & Wilson, 1994; Koos, Tepper, & Wilson, 2004; Plenz, 2003; Tunstall, Oorschot, Kean, & Wickens, 2002).

Projection neurons divide into two subpopulations: *striatonigral* (SN) neurons projecting to the SNr/GPi and originating the direct pathway and *striatopallidal* (SP) neurons projecting to the GPe and originating the indirect pathway (Bolam et al., 2000; Kawaguchi, 1997). These two populations of neurons have different chemical and physiological properties. Most importantly, they express different types of dopamine receptors (e.g., Gerfen et al., 1990; Tepper & Bolam, 2004). As a result, dopamine has opposite effects on them: It excites SN neurons but inhibits SP neurons.

Within projection neurons, it is also possible to make a distinction between small groups that do not stain for acetylthiocholinesterase (called *striosomes* or *patches*) located among the larger amount of neurons that do (collectively known as the *matrix*; Gerfen, 1992; Graybiel & Ragsdale, 1978). Striosome neurons are special in that they project preferentially to the dopaminergic neurons of the SNc (Gerfen, 1984; Gerfen, Baimbridge, & Miller, 1985).

Due to the dominant inhibitory pressure exerted by striatal interneurons (which is described in the next section), projection neurons are predominantly silent and inactive at rest (Wilson & Groves, 1981). Cortical activity is not sufficient, per se, to trigger action potentials in projection neurons. Cortical signals, however, are capable of leading them into the *up state*, a condition where a burst of action potentials can be easily triggered by either an additional increase in excitation or a drop in inhibition (Bolam et al., 2000; Tepper & Bolam, 2004; Wilson, 1993). Therefore, despite receiving most of the cortical afferents, the response of projection neurons is controlled and modulated by other factors, which include the inhibition from interneurons and the local release of important neurotransmitters like dopamine.

Interneurons in the Striatum

Striatal interneurons compose up to 23% of the cells in primates (Graveland et al., 1985). There are two main types of interneurons: GABAergic (which can be further classified in different subtypes; see Tepper & Bolam, 2004, for a review) and cholinergic interneurons. The relationship between projection neurons, GABAergic interneurons, and cholinergic interneurons is summarized in Figure 2.

GABAergic interneurons target mainly projection neurons; in fact, they are the projection neurons' main source of inhibition (Koos & Tepper, 1999; Koos et al., 2004; Mallet, Le Moine, Charpier, & Gonon, 2005; Tepper & Bolam, 2004). They also receive widespread cortical projections. Whereas cortical afferents to projection neurons are organized topologically, afferents on interneurons are more convergent, with inputs from different cortical areas converging on the same interneurons. This makes them an ideal site for integrating information from different areas (Bolam et al., 2000).

Cholinergic interneurons target both interneurons and projection neurons. They have an excitatory effect on GABAergic interneurons (Tepper & Bolam, 2004) and a powerful inhibitory effect on

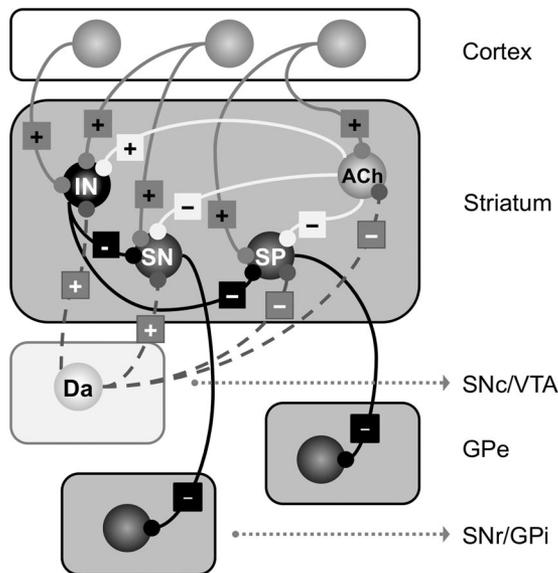


Figure 2. A simplified representation of the connections between different types of striatal neurons (adapted with modifications from Tepper & Bolam, 2004). SP = striatopallidal projection neuron; SN = striatonigral projection neuron; Da = dopaminergic neuron; ACh = cholinergic interneuron; IN = GABAergic interneuron; "+" = excitatory connection; "-" = inhibitory projection; SNc = pars compacta of the substantia nigra; VTA = ventral tegmental area; GPe = external part of the globus pallidus; SNr = pars reticulata of the substantia nigra; GPi = internal part of the globus pallidus.

projection neurons (Pakhotin & Bracci, 2007). These interneurons are believed to be tonically active neurons (TANs). TANs have elevated tonic activity at rest, but their firing rate drops suddenly at the onset of behaviorally significant stimuli (Apicella, 2002, 2007; Morris, Arkadir, Nevet, Vaadia, & Bergman, 2004). These neurons are also thought to receive widespread cortical connectivity and have been suggested to play an important role in detecting and encoding contextual information for action application (Apicella, 2007). The suspension of cholinergic response is known to be crucial to reduce the activity of GABA interneurons (Koos & Tepper, 2002), which in turn permits projection neurons to fire.

It is crucial to understand what causes the pause in the TANs' response. Research suggests that these neurons become inactive after an initial excitatory activity. Thus, an initial cortical stimulation briefly excites TANs, and this excitation determines a subsequent transient pause that allows projection neurons to fire (Reynolds, Hyland, & Wickens, 2004; Reynolds & Wickens, 2004). This provides a possible mechanism by which cortical signals can be gated to projection neurons and, hence, proceed through the circuit.

In summary, striatal interneurons are in an ideal position for integrating signals for diverse cortical areas. By maintaining or releasing their inhibitory output on SN and SP cells, striatal interneurons are capable of modulating the response of large ensembles of projection neurons and ultimately of influencing the entire output of the circuit.

Cognitive Functions of the Basal Ganglia

The most obvious symptoms of basal ganglia pathologies are movement disorders. Therefore, their contributions were initially thought to be limited to motor control. Converging evidence from single-cell recordings, lesion studies in humans and animals, and brain imaging studies in humans have challenged this view and have made it clear that the basal ganglia play important roles outside the motor sphere. In fact, the range of their contributions spans many different cognitive faculties. These include procedural memory (Packard & Knowlton, 2002), habit and skill learning (Knowlton, Mangels, & Squire, 1996), attention (Ravizza & Ivry, 2001; Teicher et al., 2000), perception (L. L. Brown, Schneider, & Lidsky, 1997), and language (Prat, Keller, & Just, 2007; Teichmann, Dupoux, Kouider, & Bachoud-Levi, 2006; Ullman et al., 1997). Experimental evidence suggests that the basal ganglia contribute to even higher level cognitive functions, like planning (Anderson, Albert, & Fincham, 2005; Dagher, Owen, Boecker, & Brooks, 2001), syllogistic reasoning (Goel, Buchel, Frith, & Dolan, 2000), and mathematical problem solving (Anderson, 2005; Stocco & Anderson, 2008).

Stimulus–Response Association Learning in Animals

Many authors have suggested that the basal ganglia circuitry is the key component of a specialized memory subsystem. This subsystem mediates the acquisition of stimulus–response associations (Knowlton et al., 1996; Mishkin & Petri, 1984; Packard, Hirsh, & White, 1989; Packard & Knowlton, 2002). Lesion studies in rats (e.g., Kesner, Bolland, & Dakis, 1993) and monkeys (e.g., Fernandez-Ruiz, Wang, Aigner, & Mishkin, 2001; Teng, Stefanacci, Squire, & Zola, 2000) have provided evidence that basal

ganglia lesions impair the acquisition of motor responses that are conditional on discriminating stimuli. Additionally, *in vivo* electrophysiological recordings in animals have shown that the acquisition of new habits results in major changes in the spiking patterns of the basal ganglia neurons (Jog, Kubota, Connolly, Hillegeart, & Graybiel, 1999).

In rats, this stimulus–response memory system can be dissociated from the hippocampus system, which mediates the acquisition of spatial memories. A comparison of lesion effects in the two structures highlights the differences of the corresponding memory subsystems (see Packard & Knowlton, 2002, for a review). For instance, rats can be trained to visit all the arms in a radial maze in order to receive food. This task requires keeping track of the visited locations, and rats with hippocampal lesions are unable to perform it correctly. Basal ganglia lesions, on the other hand, do not affect performance on this task. However, when the task is changed so that rats have to visit all the lit arms of a maze, one can observe the opposite pattern: Basal ganglia lesions, but not hippocampal lesions, impair performance (Packard et al., 1989). Similarly, lesions of the basal ganglia impair rats in finding a platform in a water tank when the platform location varies across trials but is consistently predicted by a visual cue. Lesions to the hippocampus, on the other hand, impair rats in finding the platform when it is kept in the same location but the visual cues are varied (Packard & McGaugh, 1992).

In a particularly elegant experiment, Packard and McGaugh (1996) put rats in one arm of a plus-maze. Food was consistently positioned at the end of either the left or the right arm. After rats were successfully trained to find the food, the authors chemically knocked out either the rats' hippocampus or caudate nucleus. The animals were then put in the opposite maze arm, while the food was kept in the same location. Rats with impaired caudate nucleus relied on the hippocampus-mediated place memory and went back to the food location. This required them to make a turn in the opposite direction to what they had previously learned. On the contrary, rats with impaired hippocampus fell back on their preserved stimulus–response memory, turning in the same direction as they were trained, and thus ending up opposite to the food location (Packard & McGaugh, 1996).

Nondeclarative Learning in Humans

Many authors think that, in humans, the basal ganglia mediate the acquisition of nondeclarative knowledge, whereas the declarative subsystem is supported by the hippocampus (Knowlton et al., 1996; Squire, 1992; Squire & Zola, 1996; but see Reder, Park, & Kieffaber, 2009, for a critique of this taxonomy). However, it is often difficult to dissociate these two types of knowledge. One way to isolate the nondeclarative subsystem is to use tasks that require the acquisition of complex sensorimotor skills. It has been shown, for example, that patients with basal ganglia disorders, such as Parkinson's disease, are impaired in learning mirror reading, whereas amnesic patients with hippocampal damage are not (Cohen & Squire, 1980). Similarly, patients affected by a basal ganglia disorder known as Huntington's disease are impaired at adapting to wearing prism goggles, whereas patients affected by Alzheimer's disease (which affects the frontotemporal regions) perform as well as controls (Paulsen, Butters, Salmon, Heindel, & Swenson, 1993).

Another way to isolate nondeclarative learning is to use probabilistic classification tasks. One example of these paradigms is the weather prediction task (Gluck & Bower, 1988; Knowlton, Squire, & Gluck, 1994). In this task, participants have to predict “sun” or “rain” on the basis of drawings on four different cards. Each configuration of cards has a different probability of being associated with either of the two outcomes. The underlying assumption is that complex arrangements of stimuli and nondeterministic association between stimuli and response make this type of task impossible to solve by relying on declarative memory. This assumption is somewhat simplistic, and at least some participants do indeed use declarative strategies (Gluck, Shohamy, & Myers, 2002). However, there is evidence that performance in the weather prediction task depends at least partially on the basal ganglia. Patients with Parkinson’s disease, for instance, are unable to perform above chance level (Knowlton et al., 1994). Furthermore, a neuroimaging investigation showed that performance in this task correlates with activity in the caudate nucleus (Seger & Cincotta, 2005).

Another common nondeclarative learning paradigm is the serial reaction time task and its variants. In a typical serial reaction time task (e.g., Nissen & Bullemer, 1987), participants continuously track the position of a dot on a screen by pressing the finger corresponding to the dot’s location. The order of dot locations repeats according to some pattern or rule. Participants’ latencies usually decrease for those screen locations that are predictable, thus exhibiting some form of learning. However, they are typically unaware of the underlying organization of the sequence (e.g., Destrebecqz & Cleeremans, 2001), suggesting that their learning was nondeclarative in nature. As in the case of probabilistic classification tasks, it is hard to rule out subtle declarative influences (see Shanks & St. John, 1994, for a review). Nonetheless, experimental evidence suggests that this task does tap into an implicit learning system underpinned by the basal ganglia. Neuropsychological studies have reported that patients with Huntington’s (Willingham & Koroshetz, 1993) or Parkinson’s disease (Jackson, Jackson, Harrison, Henderson, & Kennard, 1995) are typically unable to learn the hidden sequence. Also, neuroimaging studies of serial reaction time tasks have reported that activity in the striatum is correlated with implicit sequence learning (Doyon, Owen, Petrides, Sziklas, & Evans, 1996; Peigneux et al., 2000; Rauch et al., 1997).

In summary, experimental evidence suggests that the basal ganglia circuit is part of a distinctive learning system. This system underlies the acquisition of stimulus–response associations in animals, and some forms of skills and procedures in humans.

Working Memory

The basal ganglia circuit is connected to prefrontal regions, such as BA 46, that are also important for working memory (Middleton & Strick, 1994). As expected, patients with either Parkinson’s or Huntington’s disease are impaired in tasks tapping different forms of working memory (Gabrieli, 1998; Gabrieli, Singh, Stebbins, & Goetz, 1996; Lawrence, Sahakian, & Robbins, 1998; Owen, Iddon, Hodges, Summers, & Robbins, 1997). Conversely, the administration of dopamine, which is used to ameliorate the conditions of Parkinson’s patients, enhances their working memory performance (Cooper et al., 1992). Working memory-related activity in the

basal ganglia has been reported in a number of neuroimaging studies (Braver et al., 1997; Lewis, Dove, Robbins, Barker, & Owen, 2004; Owen, Doyon, Petrides, & Evans, 1996; Rypma, Prabhakaran, Desmond, Glover, & Gabrieli, 1999).

Recently, McNab and Klingberg (2008) directly investigated the involvement of the basal ganglia in working memory in a neuroimaging experiment. They used a spatial working memory task, where participants had to memorize the location of three or five circles on a screen. The circles could be either red or yellow. Each trial was preceded by a cue, which instructed participants to treat the yellow stimuli as targets or distractors. The authors found that basal ganglia activity in the period between cue and stimuli presentation was larger in the distractor trials than in the target trials. This suggests a role for the basal ganglia in filtering out stimuli that should not be included in the active set maintained in working memory.

Besides McNab and Klingberg (2008), other studies have found a correlation between individual differences in working memory and basal ganglia activation. In two investigations of the neural basis of individual differences in language comprehension, Prat and colleagues found that individuals with high working memory capacity showed greater recruitment of the caudate nucleus with increasing task demands (Prat & Just, 2010; Prat et al., 2007).

Zhang et al. (2007) examined individuals expressing two genes that determine the presence of D2 dopamine receptors in the striatum. Participants were divided into those who carry only one copy of the genes (heterozygotes) and those who carry two copies (homozygotes). Heterozygotes produce smaller amounts of D2 receptors than do homozygotes. Correspondingly, their behavioral performance in a working memory task was inferior to that of homozygotes. The authors also recorded participants’ brain activity in a functional magnetic resonance imaging (fMRI) scanner. Subsequent analysis showed that heterozygotes’ activity in the striatum was significantly smaller than that of homozygotes, suggesting less resource utilization and more efficient processing in the striatum.

Reward Processing and Decision Making

The basal ganglia also figure prominently among those structures responsible for reward processing and reward-based learning. A number of landmark single-cell recording studies in monkeys have shown that the spiking rates of dopamine neurons projecting to the striatum encode different aspects of reward-related information (Schultz, Apicella, & Ljungberg, 1993; Tobler, Fiorillo, & Schultz, 2005). In these studies, monkeys undergo a conditional learning procedure by which a conditional stimulus (e.g., a light) comes to be associated with a reward (i.e., orange juice). Dopamine neurons projecting to the striatum are initially responsive to the juice administration. During the conditioning phase, however, dopamine bursts appear when the conditioned stimulus is presented and not when the actual reward is given. Furthermore, delaying or omitting a reward causes a decrease in the response of dopamine neurons at the time when the reward was expected.

This pattern clearly does not reflect reward per se, but rather the difference between the actual and the expected reward. In fact, the activity of dopamine neurons can be accurately predicted on the basis of an algorithm known as temporal difference (TD) learning (Sutton, 1988). This algorithm estimates future rewards

on the basis of the difference between reward predictions at two consecutive moments in time. The dopamine response closely mirrors the error term that encodes this difference (Barto, 1995; Schultz, Dayan, & Montague, 1997). Neuroimaging techniques made it possible to replicate these learning tasks in humans. Patterns of activation that reflect the reward expectancy error (and thus the modulatory effect of dopamine signals) were found in the striatum (O'Doherty, Deichmann, Critchley, & Dolan, 2002).

Taken together, these studies show that striatal activity is modulated by reward signals. One might wonder, however, how exactly these signals are employed by the basal ganglia and translated into behavior. Given their role in learning, one possibility is that the basal ganglia use the difference between predicted and actual rewards as a feedback signal to direct learning from one's own errors. If this is the case, basal ganglia disorders should also impair learning depending on performance feedback. In fact, patients with Parkinson's disease show abnormal learning and error patterns in a binary choice task where the two options have different reward probabilities. In particular, their deprivation of dopamine signal has the effect of making patients more sensitive to errors than to correct trials, and therefore more able to discriminate between low-reward stimuli than between high-reward stimuli. Conversely, patients under dopamine medication show an opposite pattern of response, discriminating better between high-rewarding stimuli (Frank, Seeberger, & O'Reilly, 2004).

Because reward signals in the basal ganglia reflect the predicted value of possible options, they provide a natural basis for decision making. Consistent with this hypothesis, neuroimaging studies have provided mounting evidence that the striatum is one of the pivotal regions that activates in decision-making tasks (Montague, King-Casas, & Cohen, 2006; O'Doherty, 2004). For example, striatum activation correlates with a decision's expected reward and risk (Preuschoff, Bossaerts, & Quartz, 2006). Also, individual differences in loss aversion in decision-making tasks correlate with individual differences in striatal activity (Tom, Fox, Trepel, & Poldrack, 2007).

Computational Models of the Basal Ganglia

Interest in the basal ganglia has spawned many modeling attempts. Different models focus on different characteristics of this circuit. For instance, many important models deal with the role of the basal ganglia in motor programming (e.g., Humphries, Stewart, & Gurney, 2006). Other models simulate specific aspects of the basal ganglia physiology, such as the interactions of different neurotransmitters (Kötter & Wickens, 1995; Wickens, Alexander, & Miller, 1991; Wickens & Arbuthnott, 1993) or the generation of signature spike patterns (Humphries et al., 2006; Terman, Rubin, Yew, & Wilson, 2002). This section briefly reviews only those models that deal with the cognitive functions of the basal ganglia. This section loosely follows a taxonomy introduced by Gillies and Arbuthnott (2000) and outlines the differences between existing models and our approach.

Reinforcement Learning

The involvement of the basal ganglia in reward-based learning has attracted many modeling attempts (see Joel, Niv, & Ruppin, 2002, for an in-depth review). These models implement possible

ways by which the basal ganglia learn to correctly predict the expected value of an action on the basis of its previous rewards. They usually import methods and algorithms from the reinforcement learning literature in artificial intelligence (e.g., Sutton & Barto, 1998). The connection between reinforcement learning and the basal ganglia relies on the striking similarity between the intensity of the dopamine signal that innervates the striatum, on one side, and the error term in reinforcement learning. This error term, known as the *temporal difference* (TD), is used to update an action's predicted value in the TD-learning algorithm (Sutton, 1988). Although there are basal ganglia models that are based on different reinforcement algorithms (e.g., Dominey, Arbib, & Joseph, 1995; O'Reilly & Frank, 2006), TD-learning-based models are the most common and possibly the most successful.

One influential account (Barto, 1995) uses the error term within a computational architecture composed of an *actor* and a *critic*. The actor learns to execute the action that maximizes the predicted reward. The critic learns to correctly predict the value of each action by correcting its estimated reward by the error term. Barto (1995) proposed that the actor and the critic correspond to the matrix and patch compartments of the striatum, respectively, and the error term is conveyed by SNc dopamine neurons. Thus, patch neurons learn to associate stimuli with their expected rewards, and matrix neurons learn to associate stimuli with rewarding actions.

Barto's (1995) original framework is still influential (Montague, Hyman, & Cohen, 2004; Schultz, Dayan, & Montague, 1997), and many researchers have elaborated upon it. For example, Suri and Schultz (1998) introduced a more robust actor based on competing units, and Suri, Bargas, and Arbib (2001) introduced a more complex mechanism for selecting actions based on the opposition of the direct and indirect pathways.

Other models have improved the critic component. In the model by Houk, Adams, and Barto (1995), striosomes project to the SNc through the direct and indirect pathways. The signal carried by the direct pathway is the current reward prediction. The indirect pathway, because of its longer path and additional inhibitory synapses, carries the previous reward prediction (with opposite sign). The sum of these two quantities is the prediction error, as used in the TD-learning algorithm. It is also the input of dopamine neurons, which feed the error signals back to matrix and striosome neurons in the striatum.

The model by Houk et al. (1995) depends on bidirectional connectivity between striosomes and the SNc. This assumption does not match with physiology (see Joel et al., 2002, for a review of this problem), and subsequent models relaxed it. In these models, striosomes deliver only an inhibitory signal that reflects the opposite of predicted reward. The current reward, instead, is provided by excitatory connections from prefrontal cortex (Contreras-Vidal & Schultz, 1999) or from the pedunculopontine nucleus (J. Brown, Bullock, & Grossberg, 1999).

Some authors have questioned the equivalence of dopamine signals with reward predictions errors. It has been noted that dopamine neurons also fire under circumstances that lie beyond the scope of reinforcement. For instance, they respond to uncertainty (Fiorillo, Tobler, & Schultz, 2003) and novelty (Redgrave & Gurney, 2006) and unexpected rewards (Schultz, 1998). Finally, some authors have argued that the dopamine response is really too rapid to reflect adequate estimates of prediction errors (Redgrave & Gurney, 2006).

TD-based models tend to cover only a rather specific portion of the spectrum of the cognitive functions of the basal ganglia. For instance, none of the models reviewed above are concerned with the issues of skill acquisition or working memory update; these functions are the focus of this article.

On the other hand, these models have worked out in remarkable detail the connection between the dopamine system and the basal ganglia (see Montague et al., 2004). This makes it easier to integrate reinforcement techniques within larger models of the basal ganglia (e.g., Dominey et al., 1995; O'Reilly & Frank, 2006). These integrated approaches have proven both useful and powerful. Our model also follows this path, providing a detailed account of how dopamine release and caudate activity are affected by both reward and learning.

Serial Processing

Some functions of the basal ganglia are serial in nature. This is the case, for instance, of motor skill acquisition or sequence learning. This fact has spurred a number of models that investigated the possible neural substrates of sequence learning in the basal ganglia. Some of them (e.g., Beiser & Houk, 1998; Dominey et al., 1995; Suri & Schultz, 1998, 1999) rely on reinforcement learning techniques, so that the basal ganglia could learn to associate the proper action with a stimulus representing the proper position within a sequence. In some models, serial positions are represented externally, in the form of consecutive environmental cues (Suri & Schultz, 1998, 1999). In other models, positional cues are represented internally, in the form of representations in the prefrontal cortex (Beiser & Houk, 1998; Dominey et al., 1995).

Berns and Sejnowski (1998) followed a different approach, where the serial order of actions is directly encoded in the basal ganglia. Their model assumes that processing of stimuli is delayed in time along the indirect pathway. Therefore, this pathway can hold a temporary representation of the previous action. This temporary representation is fed back to the GPe and used to associate two consecutive steps in a sequence of operations.

Our model does not include any special mechanism for sequence learning. However, as long as external or internal representations can discriminate between different serial positions, the model is able to execute operations in an ordered sequence.

Gating Functions and Working Memory

Redgrave et al. (1999) argued that many aspects of the basal ganglia physiology could be understood in terms of a general device that allows selection of information and inhibition of competing programs. Gurney, Prescott, and Redgrave (2001) provided an original account of this gating mechanism. In their model, an action is generally defined as the opening or closing of particular input *channels* that flow from the cortex through the striatum. The activation of a neuron within a channel indicates its salience. Two competing pathways that converge on the output nuclei concur in selecting the appropriate channels. These two routes only marginally overlap with the direct and indirect pathways. In the model, the first route proceeds from the cortex to the striatum and the output nuclei. Here, lateral inhibition among striatal units is instrumental in enhancing the difference in salience among the active input channels. A second pathway proceeds from the cortex

through the STN and the output nuclei; it computes the sum of saliencies of all the competing actions. Active channels are compared to this background activity, and only the most active ones are disinhibited.

In Gurney et al.'s (2001) model the basal ganglia gate signals to the frontal cortex. This view offers a natural connection with the role of the basal ganglia in working memory, where several studies (e.g., Lewis et al., 2004; McNab & Klingberg, 2008) have shown that the basal ganglia control the access of new information to a short-term store. Many subsequent models, including ours, have capitalized on this idea. Other related models include those of Amos (2000) and Monchi, Taylor, and Dagher (2000).

The FROST (FRONTAL, STRIATAL, THALAMIC) model by Ashby, Ell, Valentin, and Casale (2005) takes an original stance on the role of the basal ganglia in working memory. In contrast to the majority of working memory models (e.g., O'Reilly, Braver, & Cohen, 1999), FROST does not assume that items in working memory are actively maintained by recurrent connections within prefrontal cortex. Instead, activation is maintained by three parallel loops connecting prefrontal cortex with parietal cortex, the basal ganglia, and the thalamus directly. The FROST model is thus capable of explaining the existence of sustained activation in cells outside the prefrontal cortex, like in the head of the caudate or the thalamus. Additionally, the FROST model includes a sophisticated architecture of how the different cortical layers are connected to the cortical, striatal, and thalamic loops. Our model also pays attention to the role of different cortical layers, and it was influenced by FROST's assumptions.

The FROST model, however, does not address specific issues of the basal ganglia physiology, such as the existence of the direct and indirect pathways and the internal architecture of the striatum. These physiological constraints play an important role in the PBWM (Prefrontal cortex–Basal ganglia Working Memory) model by Frank, Loughry, and O'Reilly (2001). In contrast to FROST, the PBWM model assumes that active maintenance of items in working memory occurs spontaneously through recurrent connections in prefrontal regions and that the basal ganglia modulate this process of active maintenance. In turn, their activity is the result of the opposing effects of the direct and indirect pathways. When the direct pathway is the more active, it transmits a “go” signal the corresponding region. This signal has the effect of temporarily disabling the recurrent connections and interrupting the maintenance of representations. And when the recurrent connections are disabled, new contents from posterior regions can be copied in, letting them overwrite the existing representation. The indirect pathway, on the other hand, encodes a “no-go” signal that protects prefrontal representations by preventing new contents from being copied.

The PBWM model is arguably the closest to our approach. There are, however, three significant differences. First, although PBWM does include a gating mechanism, it does not include full routing capabilities. When a group of cortical neurons receives a “go” signal, any incoming representation can be gated. Ultimately, the model depends on appropriate cortical connectivity to make sure that representations always reach the correct destination. In our model, on the other hand, striatal subdivisions contain a representation of their corresponding cortical region's connectivity and are able to tune the routing of information more precisely.

A second point of disagreement concerns the computational role of striatal interneurons. PBWM works virtually in absence of any significant contribution from these cells. On the contrary, they are known to play an important computational role in silencing the spiny neurons (Koos et al., 2004; Tepper & Bolam, 2004). Our model acknowledges the principal role of interneurons in determining striatal responses.

Finally, PBWM differs from our model in term of learning capabilities. Although the PBWM is able to learn new gating patterns by means of reinforcement learning, it is unable to further simplify the sequence of learned procedures. As such, it does not possess the skill-learning capabilities that we demonstrate in our model. This difference is a consequence of the underlying assumptions about the capabilities of striatal projection neurons. In PBWM, they simply encode a binary go/no-go signal. In the routing model, on the other hand, they can incorporate larger amounts of information, which can be later used as a substitute for cortical representations.

An important, general-purpose role of the basal ganglia has been put forward within the ACT-R (Adaptive Control of Thought—Rational) cognitive architecture (Anderson, 2007; Anderson et al., 2008). ACT-R claims a functional correspondence between a specific component of the architecture (the *procedural* module) and the activity of the basal ganglia. The procedural model executes operations in the form of condition–action rules. Each rule responds to specific contents of module *buffers*, which correspond to consciously available representations held in the cortex (Anderson et al., 2008). This correspondence has led to a number of interesting and hitherto untested predictions in the past. It has been shown that activity in the caudate nucleus, as indexed by oxygen consumption in fMRI experiments, increases linearly with the number of individual cognitive steps in a task (Anderson, 2005), decreases with practice (Anderson, 2005), and increases with the amount of information that is transferred by each individual operation (Stocco & Anderson, 2008).

The ACT-R view of the basal ganglia had a significant influence on our model. There are, however, a number of differences between our model and ACT-R. First, in ACT-R the basal ganglia system provides a much more central role than in our proposal. This is because ACT-R lacks direct module-to-module communication, and the transfer of information can happen only within the basal ganglia. Our model, on the other end, assumes the existence of cortico–cortical connectivity and describes the role of the basal ganglia in the context of these existing pathways.

Another difference concerns the role of the basal ganglia as a bottleneck for the entire system. ACT-R assumes that only one specific production rule can be executed at any given moment—that there is no parallel processing in the basal ganglia. Although such a commitment can be incorporated in our model, it does not emerge from internal constraints. The number of independent actions that can be performed at any instant is ultimately a function of the striatal resources each of them consumes. Therefore, the limitation is on the overall complexity of actions, and not on their number.

Learning and Skill Acquisition

One crucial contribution of the basal ganglia is their role in habit learning and skill acquisition. Unlike habit learning, which de-

pends on reward, skill acquisition is based on practice and results in the development of automatic procedures (e.g., Shiffrin & Schneider, 1977). In contrast to habit learning, skill learning is not frequently addressed by computational models. Among the exceptions is the SPEED (Subcortical Pathways Enable Expertise Development) model by Ashby, Ennis, and Spiering (2007). SPEED was designed to account for automaticity in categorization. The processes it models, however, generalize easily to both habit and skill acquisition. SPEED assumes that the striatum is needed for the acquisition of associations between a sensory stimulus and a task-dependent motor response. This association is initially established in the basal ganglia circuit by a dopamine burst occurring when a correct categorization response is made. The basal ganglia subsequently bias the cortex toward the response that is more closely associated with a positive feedback. By means of Hebbian learning, the stimulus–response association is eventually encoded directly within the cortico–cortical pathways. This learning mechanism is relevant for skill acquisition because it provides a detailed account of how an initially goal-directed behavior becomes automatized. Other important models (e.g., Frank & Claus, 2006) also proposed similar mechanisms where the basal ganglia facilitate the creation of stimulus–response associations between cortical areas. Our routing model differs from them because it postulates an intermediate phase where the skills are encoded in the striatum. Also, these models assume that the necessary dopamine burst is triggered by a performance reward for the correct response. Our model, on the other hand, includes an internal mechanism that modulates dopamine to learn new skills.

The ACT-R cognitive architecture possesses an original skill acquisition mechanism. This algorithm, called production compilation (Taatgen & Lee, 2003), consists of merging two production rules that fire consecutively into a single-step rule. In contrast to SPEED, in ACT-R all the learned skills remain confined to the procedural module, and learning never involves the establishment of direct transmission of information between cortical areas. This algorithm has been influential in shaping the skill-learning mechanism described in this article, and our model can be considered the first effort to provide a biological account for production compilation. However, our model differs from production compilation because it does not assume the procedural knowledge is permanently stored in the basal ganglia and acknowledges the important role of cortico–cortical pathways as the final repository of skills.

The Routing Model

As anticipated above, this article proposes that the main function of the basal ganglia is to direct the flow of information processing in the cortex. Specifically, we propose that the “actions” performed by the basal ganglia can be thought of as *routing operations*. Routing operations define which signals are transferred between pairs of cortical regions. In particular, routing operations trigger one (or more) target cortical area, called the *destination region*, to accept and process information from one (or more) other region, called the *source region*. Such a mechanism enables the basal ganglia to perform general-purpose computations. This framework takes the form of a computational model, which we refer to as the *routing model*, and whose structure and principles of operation reflect known facts about the basal ganglia. The next

sections introduce the model, present a number of simulations, and demonstrate how routing operations can be easily shaped and modified by practice, providing insights into the skill-learning mechanisms of the basal ganglia.

The routing model is implemented as a neural network.¹ The model architecture mirrors the most important features of basal ganglia anatomy and is visually represented in Figure 3.² Before going into the implementation detail, we begin with a brief overview.

In the model, cortical areas are constantly communicating with each other along cortico–cortical pathways (see the top layer of Figure 3). In most cases, cortico–cortical signals are not strong enough to significantly affect the representations held in a region. In addition to communicating with each other, each cortical area transmits signals to a corresponding subdivision of the striatum. Each striatal subdivision receives afferents from one corresponding cortical region. Within a subdivision there is an internal organization of ensembles of neurons corresponding to the cortical destinations that the contents of the source region might be sent to. Thus, striatal subdivisions reflect cortical topology at two levels. At a macro level, they mirror the organization of cortex into specific regions (see the striatum in Figure 3). At a lower level, each of them reflects the organization of cortical connectivity of the projecting cortical region. This two-level organization is mirrored in the other nuclei of the circuit as well. If one assumes that each model cortical region is connected to every other one, this two-level organization can be easily visualized in form of a matrix, as in Figure 3. Note that a real biological system does not need to represent all the possible source–destination pairs and saves considerable resources by representing only those that correspond to existing cortico–cortical pathways.

In the striatum, the activation of SN neurons determines which contents will be selected, and the activation of SP neurons determines where they will be routed. As explained above, the opposite contributions of SN and SP neurons travel along the direct and indirect pathways and eventually sum up in the SNr/GPi. Their sum determines which destination the selected cortical representation will be transferred to. When the activity of an SP ensemble is not sufficient to contrast the effect of the corresponding SN group, the SN signal is transmitted through the thalamus and henceforth to the cortex.

Each ensemble of cells in the output nuclei SNr/GPi projects to the corresponding ensemble of units in the thalamus. We refer to each ensemble of neurons encoding a specific source–destination pairing as a thalamic subdivision. Through the thalamus, the signal reaches the appropriate destination regions, providing the necessary local bias to let the source region’s signal come in. It is important to note that this architecture does not imply that thalamic subdivisions project to different cortical regions. On the contrary, projections from a thalamic subdivision are specific and focal to their cortical targets and do not diffuse over different regions. Within the same thalamic subdivision, different cells simply project to cortical neurons within the same area that receive afferents from different cortical regions. It is only because of this reason that they can modulate the different contributions of different cortical afferents.

The state of activation in the model thalamus provides a useful, condensed representation of the contents that have been selected and where they are being routed to. In fact, this article often uses

the state of the thalamus as an indicator of the routing operations that have been performed by the basal ganglia. Routing operations can take fairly complex forms. Figure 4 illustrates three such cases, corresponding to the routing of one (top), two (middle), or three (bottom) contents from two different regions. As shown in the figure, the model is capable of performing transfers that require a convergence of multiple sources over the same destination.

Model Implementation

As it is common in connectionist models (Amit, 1993; O’Reilly & Munakata, 2000; Rolls & Treves, 1998), the model neurons consist of simple computational units that calculate an activation value from their inputs. The activation value is always in the range [0, 1] and represents the normalized firing rate of that neuron. The neuron’s input consists of the sum of all the projecting neurons’ activations, weighted by their corresponding synaptic strengths. Synaptic strengths are represented as scalar values, called *weights*. Weights are negative for inhibitory synapses and positive for excitatory ones. Appendix A provides a detailed overview of the rules that govern the behavior of the model neurons, and Table A1 offers a concise summary of the equations and parameters that govern each type of neuron’s response.

With respect to basal ganglia physiology, a number of simplifying assumptions have been made. They were needed to keep the model simple and to avoid introducing ad hoc mechanisms and parameters. The assumptions are the following:

Simplifying Assumptions

Simplifying Assumption 1. The SNr and the GPi have been treated as a single nucleus, which is referred to as the SNr/GPi. This is consistent with the anatomical and functional similarity between the two nuclei (e.g., Albin et al., 1989; Bolam et al., 2000) and is a common practice in basal ganglia models (Ashby et al., 2005; Frank et al., 2001; Monchi et al., 2000; O’Reilly & Frank, 2006).

Simplifying Assumption 2. The model’s indirect pathway includes only projections from the GPe to the SNr/GPi, omitting the intermediate station of the STN. This is due to the fact that, in the model’s architecture, the longer projections that proceed through the STN eventually convey a similar signal. This simplifying assumption is common to other models (e.g., Frank et al., 2001; O’Reilly & Frank, 2006). Note that projections from the STN to the SNr/GPi are actually included in the model as part of the hyperdirect pathway (see Figure 3).

Simplifying Assumption 3. The “backward” projections that reach back to the striatum from the GPe and from the thalamus have been left out of the model. It has been often suggested that these projections have important control and feedback functions (e.g., Redgrave et al., 1999). However, feedback and processing control was not necessary for our model to run.

¹ The entire model code is available for download at <http://act-r.psy.cmu.edu/publications/pubinfo.php?id=828>

² For clarity’s sake, the depicted model contains fewer cells and fewer cortical regions than actually used in most of the simulations. Additionally, simulated patterns are represented as bidimensional matrices instead of unidimensional vectors.

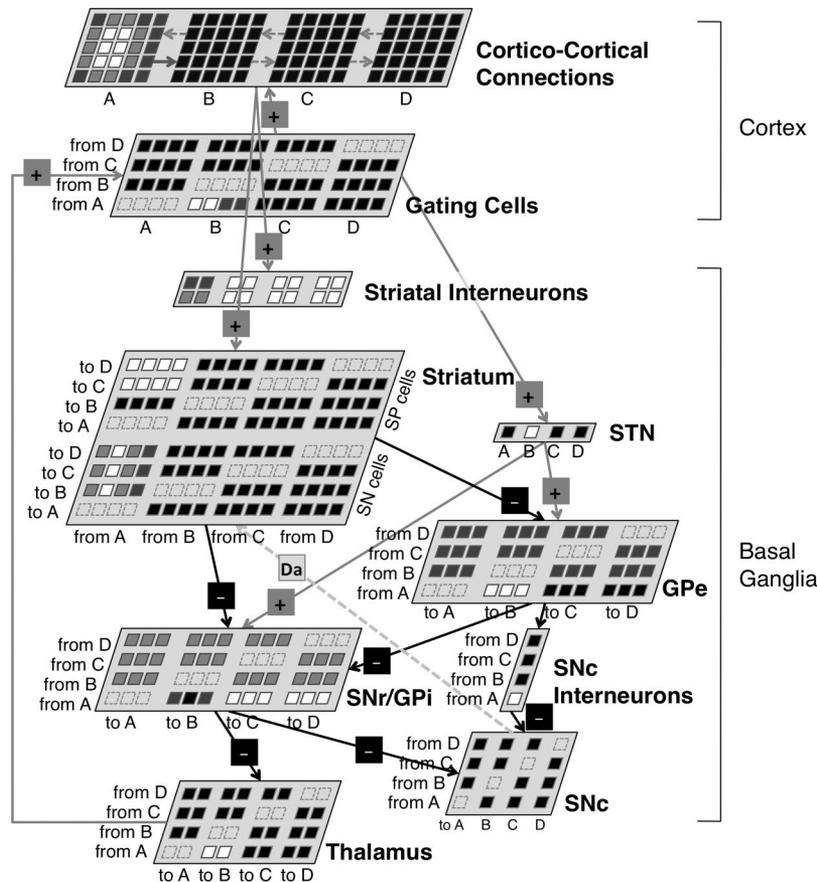


Figure 3. Architecture of the model. The different cell colors reflect different degrees of activation, from black (no activation) to white (maximum activation). The solid cortical arrow in the top layer illustrates which source–destination binding has been performed by the basal ganglia; it corresponds to the transfer of a representation from cortical region *A* to cortical region *B*. Excitatory projections are marked with a “+”; inhibitory projections with a “-”; dopamine projections with a “Da.” STN = subthalamic nucleus; GPe = external part of the globus pallidus; SNc = pars compacta of the substantia nigra; SNr = pars reticulata of the substantia nigra; GPi = internal part of the globus pallidus.

Simplifying Assumption 4. The direct cortico–thalamic projections have been omitted from our model. This omission is also common to other models (e.g., Ashby et al., 2007; Gurney et al., 2001). Cortico–thalamic connections have been sometimes included as a means of maintaining activation in prefrontal cortex through a cortico–thalamo–cortical loop (e.g., Ashby et al., 2005; Beiser & Houk, 1998) or as a means of ensuring proper gating of signals between cortical regions (e.g., Frank et al., 2001). However, our model does not deal with the dynamics of working memory maintenance and does not require cortical inputs to ensure proper gating. Therefore, the omission of cortico–thalamic pathways does not undermine the model’s validity.

Architecture of the Nuclei

In the model, four nuclei in the circuit (the striatum, GPe, SNr/GPi, and the thalamus) have a two-level organization. At the first level, their cells are organized into subdivisions that receive afferents from the same cortical area. Throughout the article, we

refer to this area as the *source* region. At a second level, neurons within each subdivision are divided into groups that reflect the projections of the source region to other areas.

Many authors have advocated that the basal ganglia are organized into segregated loops that run parallel to each other (e.g., Alexander et al., 1986). Other authors have argued, on the contrary, that there is room for open-ended loops within the circuit (Haber, 2003; Joel & Weiner, 1994; McFarland & Haber, 2002). It should be noted that, although it permits open-loop exchange, the proposed two-level architecture remains consistent with much of the experimental evidence that supports the closed-loop view of the basal ganglia. This is because, if one follows the projections from a cortical area throughout the circuit, some of them will form a separate loop that runs in parallel with other channels. To examine the degree to which a loop is segregated, one applies both anterograde and retrograde tracing methods to the same cortical region and examines the overlap of their targets. Kelly and Strick (2004) used this procedure and found that, in the monkey putamen, the cells identified by anterograde tracing and retrograde tracing

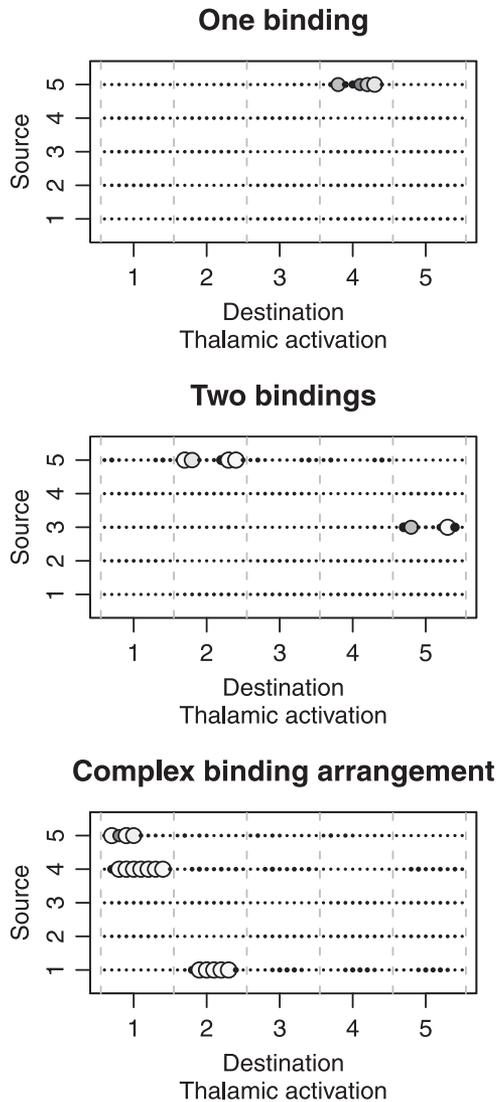


Figure 4. A “balloon” representation of the activity of neurons in the model thalamus in three example cases. In the plots, each circle represents the activation of an individual neuron in the model thalamus. The amount of activation is reflected in both its circle size (from small to large) and color (from black to white). In the model, the thalamus receives projections from the output nuclei (SNr/GPi) of the basal ganglia and maintains the same topological organization. Therefore, it is convenient to illustrate the activity of its model neuron in source–destination matrix arrangement. The pattern of neuron activity in the thalamus provides a concise representation of the operation that has been executed last. The three graphs illustrate three example cases. Top: in this case, only one representation is transferred, from Region 5 to Region 4. Middle: In this case, two transfers occur at the same time: From Region 5 to Region 2 and from Region 3 to Region 5. Bottom: This is a complex case. Three transfers occur at the same time. The contents originate from three different regions (Regions 1, 4, and 5) but affect only two destinations (Regions 1 and 2). Notice that Region 1 receives representations from two sources (Regions 4 and 5) at the same time.

shared a common center but were only partially overlapping. This is exactly what would be expected for a two-level, “switchboard” organization as described here. This organization is also consistent with previous findings by Parthasarathy et al. (1992), who found that the degree of overlap between the striatal subdivisions targeted by pairs of cortical regions reflects the degree of cortical connectivity between the two.

Information Processing in the Striatum

The striatal matrix is the heart of the model. The matrix is composed of SP and SN neurons, organized in parallel according to the two-level system outlined above (see Figure 3). Active neurons in SN ensembles signal that the corresponding cortical region is a source region and that its contents have been picked up for routing. Active SP ensembles within a subdivision, on the other hand, prevent the corresponding destinations from receiving the source contents. That is, the SP ensembles veto or mask undesired destinations where the selected contents should not be routed. Source and destination information travels separately along the direct and the indirect pathways.

Cortical areas are larger and have more neurons than the striatal subdivision they project to. Therefore, whatever patterns of information they contain need to be compressed. Some authors have actually hypothesized that one of the most important functions of the basal ganglia is actually the functional reduction of cortical representations (Bar-Gad, Havazelet-Heimer, Goldberg, Rupp, & Bergman, 2000). In the presented model, each receiving neuron covers a limited part of the projecting region. Its receptive field is modeled as a Gaussian function, with maximum sensitivity to those cortical neurons that occupy the same relative position within the cortical array of neurons, as discussed in Appendix A. The ratio of cortical projection neurons to striatal projection neurons has been estimated to be about 10:1 (Zheng & Wilson, 2002). Therefore, we adopted this ratio in the model implementation. In addition to projecting to the direct and indirect pathways, each projection neuron also makes inhibitory connections with its local neighbors. As a result of their internal processing, SN neurons contain a compressed version of the original cortical signal. When transmitted through the basal ganglia, this compressed representation provides the necessary redundant signal that forces the destination region to receive information from the source region.

Under normal conditions, the model SN and SP ensembles are inactive. Their lack of activity is due to the strong inhibitory effect of striatal interneurons, which have an elevated baseline activity (see Figure 3). The proper set of representations in the cortex causes interneurons to momentarily cease firing. A sufficient release of inhibition makes it possible for projection neurons to be excited by the incoming cortical signals. The precise pattern of activation depends on the strength of the synapses between interneurons and projection neurons, as well as the excitatory thresholds of projection neurons (see Appendix A for details). In general, the inhibition of different combinations of interneurons will activate different ensembles of SN and SP neurons and ultimately will determine which cortical representation will be routed.

The model striatal interneurons were designed to capture the contributions of both GABAergic and cholinergic interneurons (see Figure 2). As such, they share features of both types of cells. Like GABAergic interneurons, they do exert a constant inhibitory

pressure on projection neurons and prevent them from responding to cortical stimulation under normal conditions. Like cholinergic interneurons, they are tonically active until specific cortical patterns of behavioral significance are detected. Interneurons are capable of detecting complex patterns of cortical activity because they receive widespread connections from different cortical areas. To model the characteristic pause in TANs' activity that follows excitatory cortical inputs, the interneurons' activation function was modeled as an inverse sigmoid, which decreases as its input increases (see Appendix A for details). Therefore, when proper cortical inputs are detected, interneurons release their inhibition on projection neurons, permitting the incoming cortical signals to be processed. A later section illustrates how interneurons can be trained to respond with the precise cortical pattern by means of Hebbian learning.

Many models of the basal ganglia do not simulate the contribution of striatal interneurons (e.g., Ashby et al., 2005, 2007; Frank et al., 2001; O'Reilly & Frank, 2006). However, experimental evidence suggests that they play a crucial role in modulating the response of striatal projection neurons (Tepper & Bolam, 2004). Furthermore, their activity has been the subject of many experimental studies (see Apicella, 2002, for a review).

Direct and Indirect Pathways

Projections from the SN and SP neurons originate the direct and indirect pathways. Projections from the SN terminate on the SNr/GPi—that is, on the output nuclei of the basal ganglia. The SNr/GPi shares the same two-level organization as the striatal matrix (see Figure 3). In particular, their organization closely mirrors the organization of SN ensembles. Neurons in the SNr/GPi have a sustained tonical activity, unless they are depressed by afferents from the striatum or the GPe.

The model takes an original stance on the role of the indirect pathway. According to the most common view of the basal ganglia (Albin et al., 1989; DeLong, 1990), the direct pathway pushes for the execution of a particular action by disinhibiting the thalamus, whereas the indirect pathway exerts the opposite function by maintaining the tonic activity of output nuclei cells. This influential hypothesis, commonly known as the *brake–accelerator* view (e.g., Graybiel, 2000), is often assumed in computational models of the basal ganglia (e.g., Frank et al., 2001). In the proposed model, the two pathways carry different components of an action's representation—namely, the sources and the destinations. Despite the apparent dissimilarity, this view can be seen as a generalization of the brake–accelerator model. The signal coming from the direct pathway can be seen as a command to transfer a representation to all the cortical areas connected to the source region. The signal carried by the indirect pathway, on the other hand, can be seen as a stop signal that prevents the transfer to all but the destination region. As in the brake–accelerator model, the correct execution of an action depends on the balance between these two forces. The similarity becomes more apparent if one considers the possible consequence of an imbalance in the strength between the two pathways. If the direct signal is weak, or the indirect signal is too strong, no signal can be transferred. And if the direct signal is too strong, or the indirect signal too weak, then the same representation is transferred inappropriately to undesired destinations. A later section on how the model can simulate two common disorders of

the basal ganglia (Parkinson's and Huntington's diseases) illustrates these two cases.

The STN and the Hyperdirect Pathway

The model's hyperdirect pathway proceeds from the cortex through the STN to the SNr/GPi (see Figure 3). Additionally, STN projections also target the GPe. Differently than the striatum, GPe, SNr/GPi, or the thalamus, the model STN has only a simple, first-level organization, where each subdivision corresponds to a source cortical region. This fact is consistent with the closer overlap between anterograde and retrograde tracings in the STN than in the GPe (Kelly & Strick, 2004), which suggest that the STN, contrary to other nuclei, is indeed organized in closed loops.

In the model, the tonic excitatory output of the STN contributes to the tonic activity of inhibitory neurons in the pallidus and substantia nigra, as previously suggested by Bevan and Wilson (1999) and Nakanishi, Kita, and Kitai (1987). While a cortical region is receiving signals from another region and updating its representation, it also sends excitatory inputs to the corresponding STN units. Inputs to the STN prevent the release of inhibitory pressure to the thalamus. Therefore, excitation from the cortex prevents signals from being relayed through the basal ganglia system while a cortical region is still processing its own inputs. This process paces the relaying of information from the striatum to the output nuclei (e.g., Plenz & Kital, 1999) and provides a cortical way to inhibit a preponderant response until realized. A cortical region that is exciting the STN is sending a message that can be translated as “hold your horses” (as suggested by Frank, Samanta, Moustafa, & Sherman, 2007). This function is consistent with the views of other authors (e.g., Nambu et al., 2002).

The Role of Cortico–Cortical Connectivity

The model relies on cortico–cortical connectivity for the last step of the information-routing process. Since a realistic modeling of cortico–cortical connectivity was beyond our modeling effort, the organization of cortical connectivity information exchange has been simplified. Essentially, a cortical area consists of a dedicated group of neurons that stores and processes its internal representation and a set of units that receive incoming projections from other cortical areas. These units work as an internal “gate” to the cortical region. Thalamo–cortical projections (see Figure 3) work by targeting and activating these units. In particular, selective activation of one ensemble of units favors the transmission of information coming from the corresponding cortical region. These “gate” neurons also make excitatory synapses to the corresponding subdivisions of the STN, thus temporarily restoring the output nuclei inhibition of the thalamus (see Figure 3). Therefore, their activation permits the routing of new representations, and, at the same time, prevents the delivery of new signals as long as the new representation is being processed.

Relationship to Production Systems

From a purely computational point of view, routing operations can be seen as a neural network analog to production rules in production systems. Production rules are control statements expressed in the form of condition–action clauses (“if . . . then . . .”).

The condition (or “right-hand side”) specifies when the rule can be applied, whereas the action (or “left-hand side”) specifies what is executed. The similarity between the conditional routing model and a production system can be seen if one assumes the following: A rule is embedded in the incoming and outgoing synaptic matrices of a set of striatal interneurons. The condition part of the rule is represented by incoming synapses. These, in turn, encode the specific cortical representation that will trigger the interneuron to fire. The action is encoded in the outgoing synapses to the striatal projection neurons. An action corresponds to the activation of particular ensembles of SN or SP neurons, which, in turn, trigger the transmission information from the source region of the cortex to the destination region.

Part of the flexibility of production systems originates from the use of variables in the production rules. This is because variables make it possible to use the same rule for different representations, therefore capturing a general pattern of behavior. However, variables are not easily dealt with in neural networks. To overcome this problem, a number of procedures have been proposed over the years, like a special binding space (Touretzky & Hinton, 1988); tensor product variable binding (Smolensky, 1990); temporal-dependent binding (Ajjanagadde & Shastri, 1991); and holographic reduced representation (Plate, 1995; Stewart & Eliasmith, 2008), a representation format where single values can be unambiguously extracted from a combined form.

The routing model provides an alternative solution to the variable-binding problem. In the model, a *variable* corresponds to a specific cortical location. During the execution of routing operations, moving content from a source region to a destination region corresponds to binding the variable in a pattern and using it to create a new structure. Production rules also specify constants in the structures they build, and this corresponds to transferring a fixed content to a destination region. This case is illustrated in one of the forthcoming sections on learning. This solution has the advantage of being tied to a specific neurological substrate.

The similarity between routing operations and production rules is important, because production systems have been proposed as a general means to model cognition (Newell, 1973). In fact, production systems have been successfully used as general models of cognition. In particular, they have provided the framework for *cognitive architectures*. A cognitive architecture specifies the primitive elements of a cognitive system, so that each task can be modeled simply by providing the system with appropriate task knowledge (Anderson, 1983).

Different cognitive architectures have been proposed (Anderson, 1983, 2007; Just & Varma, 2007; Laird, Newell, & Rosenbloom, 1987; Meyer & Kieras, 1997a, 1997b). Among the existing architectures, an identical perspective on the functions of the basal ganglia is assumed in the ACT-R cognitive architecture. In ACT-R, the selection and execution of rules is managed by a specific module that is explicitly identified with the basal ganglia and whose temporal course of activity has been successfully employed to predict the hemodynamic response in the head of the caudate nucleus (Anderson et al., 2008). In fact, the model presented here was explicitly developed using ACT-R’s procedural module as a reference for the functional properties to implement in the circuit. It is interesting that a neural network implementation was made of an earlier version of ACT-R (Lebiere & Anderson, 1993). Although not explicitly addressing issues of biological

plausibility, the original implementation anticipated some ideas that have been developed here, as well as in other computational models of the basal ganglia (e.g., the gating functions of the circuit, as in Frank et al., 2001, and Gurney et al., 2001).

Model Performance

In developing our model, our focus was on how the basal ganglia can perform general-purpose routing operations and how their function can provide a flexible system and organize the flow of processing within the cortex. Therefore, we mainly concentrate on testing its capabilities and performance. We first illustrate how the basal ganglia model can be used to perform a simple task. Then, we test the model across a range of different configurations, to show its robustness. Then, we show how the model can reproduce the effects of basal ganglia damage. Finally, we show how Hebbian learning can enable the model to perform powerful computations on its routing operations and how this results in changes in activity in both striatal and cortical areas.

An Example Task

This section provides an example of how the model coordinates a series of routing operations to perform a task. The example paradigm is an aural discrimination task that has been used as part of a dual-task experiment by Schumacher et al. (2001) and Hazeltine, Teague, and Ivry (2002). In this task, participants respond to the presentation of a tone. Tones could have three different pitches (220, 880, and 3,520 Hz), to which participants had to respond “one,” “two,” or “three,” respectively.

This task requires assembling a number of basic cognitive functions in a novel and arbitrary way and, therefore, depends on controlling the flow of information among cortical areas. It is also simple enough that its modeling requires very few assumptions. With some differences in the details, various authors (Anderson, Taatgen, & Byrne, 2005; Hazeltine et al., 2002; Schumacher et al., 2001) have agreed that three basic processing steps are taking place: (a) stimulus classification, during which the stimulus is presented and appropriately encoded; (b) response selection, during which the appropriate response is selected from the set of possible options; and (c) response execution, where the chosen response is eventually vocalized.

It is rather uncontroversial that the first and the third step rely on the auditory and motor cortices, respectively (see Anderson, 2007, for an fMRI investigation that confirmed this fact). More uncertain is the localization of response selection. Anderson, Taatgen, and Byrne (2005) proposed an ACT-R model that can successfully reproduce most of the experimental findings. Following ACT-R’s mapping of cognitive process onto brain regions (see Anderson et al., 2008), the model implies that response selection recruits the left lateral inferior prefrontal cortex. This interpretation is consistent with the established role of this region in selecting among competing responses in word generation and pair-associate tasks (Danker, Gunn, & Anderson, 2008; Sohn, Goode, Stenger, Carter, & Anderson, 2003; Thompson-Schill, D’Esposito, & Kan, 1999). The specific involvement of this region has been confirmed by an fMRI investigation of this task reported in Anderson (2007, Figure 4.15c).

A simple cortico-basal ganglia circuit was generated to simulate the task. The circuit was simplified to contain only the three

cortical regions required by the task. Correspondingly, the striatum contained only three main subdivisions. It was further assumed that each region was connected to the other two. In the model, response selection was simulated as a two-phase step, where the cortical region first attends to the encoded tone from the aural region and then uses it as a cue to select the appropriate response. To simulate the selection process, the model prefrontal region was connected to a data structure (perhaps corresponding to the hippocampus) that could hold the long-term representations of the three possible responses. The prefrontal region sends its internal representations to this structure and receives back the response pattern that is associated with the best-matching input representations. Each cortical region contained 100 artificial neurons.

Figure 5 illustrates how the model performs such tasks. The figure reads top to bottom, left to right. The four panels on the left-hand side represent the activation of the cortical units, divided into areas, at the three stages of task execution. Note that the two middle panels represent the two phases of response selection. Two routing operations are required to perform this task: They are represented in the two right panels. The two routing operations are required to connect the three task phases. Their implementation follows the model rules described in Anderson, Taatgen, and Byrne (2005). These rules reflect an initial level of task exposure, before participants' performance has been optimized by practice. The model was trained to perform these two routing operations with a contrastive Hebbian learning (CHL) procedure. The detailed procedure and the reasons why it was chosen are described in Appendix B.

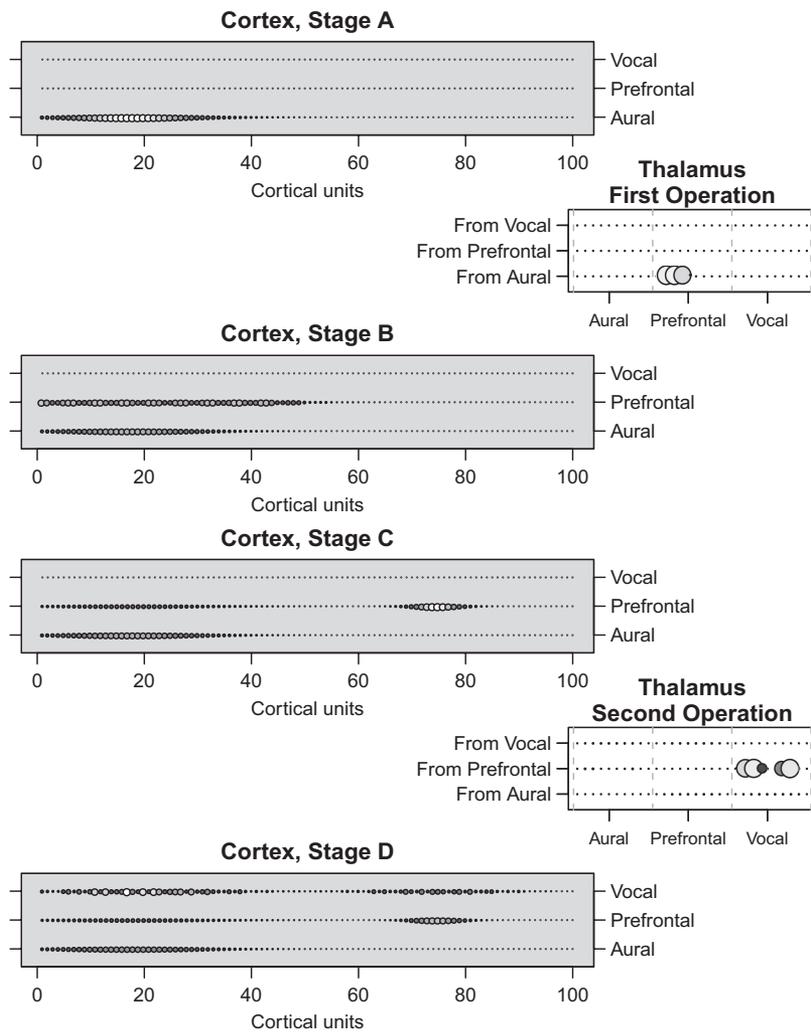


Figure 5. A schematic illustration of how the basal ganglia perform the simple stimulus–response task chosen as an example. In the figure, time flows vertically, from the top plots to the bottom. The four plots on the left represent the states of the cortical regions at the four stages of the task (see main text for further details). The two plots on the right illustrate the activation pattern in the thalamus at the moments where the two crucial routing operations are executed. These patterns provide a representation of the operation performed by the basal ganglia. The first operation causes the second region (*retrieval*) to receive the inputs from the first region (*aural*). The second operation causes the retrieved response from the second region to be routed to the third region (*vocal*).

In Figure 5, Panel A in the top left corner corresponds to the state of the cortex when the auditory signal is first encoded. The first routing operation is applied at this stage and consists of directing the transfer of the tone representation to the prefrontal region. The top right panel represents this routing operation. This panel shows the state of activation of thalamic subdivisions, organized as a source–destination matrix. The active cells are located in the subdivision that projects to the second (prefrontal) region from the source region (aural).

Activation of these thalamic terminals determines the transition to the second step, which is represented in Panel B. When the prefrontal region has received the auditory cue, it responds by selecting a pattern corresponding to the response associated with the tone. This phase is represented in Panel C. Note that the model assumes that this operation occurs within the cortex, and the basal ganglia are not involved. The second routing operation (illustrated in the bottom right panel) is triggered at this point and routes the retrieved response to the vocal region, where it can be executed as a vocal program. This corresponds to the final stage, illustrated in Panel D.

General Performance

Having the model reproduce a particular task does not provide sufficient information on its generality as an information-routing device. A series of simulations were therefore carried out to investigate the model's performance. During the simulations, three factors were varied parametrically. Two factors affect the model's configuration: the numbers of cortical regions (3, 6, 9, 12, or 15 regions) and the size of each cortical region (containing 50, 100, 150, 200, or 250 model neurons). The third factor was the number of operations learned before being tested (5, 10, 15, 20, or 25 operations). This factor was chosen to examine the interference among different possible courses of actions. Although the size of cortical regions was varied parametrically, the size of each striatal subdivision was kept constant across simulations. In particular, each striatal subdivision contained 20 SN and 20 SP neurons. Also, the number of neurons in each thalamic and SNr/GPI subdivision was kept equal to 10. These values were kept constant so that each different cortical region's size would correspond to a different ratio of cortex-to-striatum size.

For each training test, a specified number of operations were generated randomly. Each operation was to be performed in response to a different, randomly generated pattern of cortical representations, and the model was trained on each of them. Training was performed by means of the same modified version of CHL that was used for the example task, and that is outlined in Appendix B.³ During testing, one of the operations was then selected at random, and its cortical pattern was presented to the model. The pattern was propagated through the circuit, and the state of the thalamic subdivisions was compared against the desired response. The model was tested 100 times for each level of the three factors (number of regions, cortical size, and number of operations). Trial performance was assessed by comparing the state of the thalamic subdivisions against the desired response. A trial counted as incorrect whenever (a) there were active cells that did not correspond to a proper source–destination binding or (b) the desired cells were not active.

The number of incorrect trials was counted for each combination of factor levels. Each factor and each two-factor interaction was then analyzed independently, using a fixed-effects statistical model. The size of cortical regions did not have any significant effect on the model's performance, $F(4, 120) = 0.23, p = .91$, and did not interact with the other factors, $F(16, 100) < 0.55, p > .92$. On the other hand, the number of cortical regions, $F(4, 120) = 23.14, p < .0001$, of routing operations, $F(4, 120) = 2.65, p = .03$, and their interaction, $F(16, 100) = 6.38, p < .0001$, were all significant.

The left panel of Figure 6 illustrates the percentage of errors for each combination of number of regions and operations, collapsed across different sizes of cortical regions. It can be seen that the probability of making an error increased with the number of possible operations and decreased as the number of regions increased. This increase in errors can be due to the fact that, as the number of regions decreases, the routing operations' patterns become increasingly similar. Under such circumstances, undesired source–destination bindings, which were supposed to be the response of a different operation, might show up in addition to those of the executed operations. To examine this possibility, we analyzed a different measure of model's performance: the number of *additional bindings*. An additional binding was defined as a thalamic subdivision that contains active neurons but that does not belong to the desired source–destination bindings. The analysis confirmed our prediction. The number of additional bindings was not affected by the cortical size or its interaction. However, like the percentage of correct trials, this measure decreased with the number of regions, $F(4, 120) = 11.52, p < .0001$, increased with the number of operations, $F(4, 120) = 2.27, p = .06$, and was affected by their interaction, $F(16, 100) = 2.17, p = .01$ (see Figure 6, right panel).

Summary

This section has presented an overview of the model's capabilities. The model's capabilities were tested in two different ways. First, it was shown how the model performs a simple stimulus–response task. Second, it was shown how robust the model's performance is when a number of changes are made to its configuration (e.g., increasing size or increasing number of cortical regions) and when the number of available responses is progressively increased. Overall, the model's robustness in face of large changes in its structure confirms the efficiency of the basal ganglia architecture for routing information.

Skill Acquisition

Computational models have shown that reward-based learning is powerful enough to enable the acquisition of action sequences (Berns & Sejnowski, 1998; Dominey et al., 1995; Suri & Schultz, 1998, 1999) and even the complex temporal dependencies of a continuous working memory task (O'Reilly & Frank, 2006). This

³ To ensure that eventual errors made by the model in the test simulation were not due to our learning algorithm, preliminary tests were performed in which the model was tested under each combination of size and number of regions after being trained to perform a single operation. The model was always able to perform the learned operation accurately.

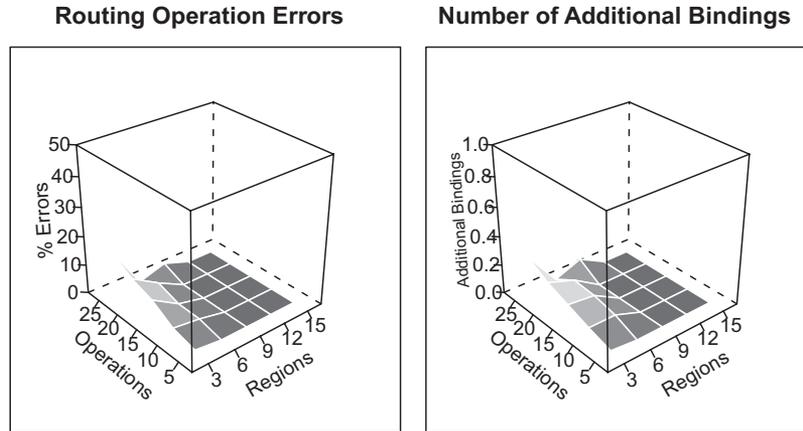


Figure 6. Performance of the model in the general simulations. Left: Mean number of incorrect trials by number of regions and number of competing operations. Right: Mean number of additional bindings by number of regions and number of competing operations. Additional bindings are source–destination bindings that are found in the model thalamic subdivision but do not belong to the desired model response.

article, however, focuses on the acquisition of new skills. The main characteristics of this type of learning is that it is mediated by practice and that it results in the development of automaticity, that is, the ability to perform the acquired skills without the need for central cognitive control (Shiffrin & Schneider, 1977).

In the model, procedural learning occurs because of practice-related changes in the routing operations. In particular, repeated transfers of the same signal between two regions causes the signal to be encoded within the basal ganglia circuit. As a result, routing operations can be ultimately performed in absence of the source representation in the cortex, saving a number of intermediate computations. The lack of need for the intermediate cortical representations is the basis of automaticity and is consistent with the large drops in activation that are consistently found in neuroimaging investigations of practice, especially in the fronto–parietal areas that are thought to underpin central cognitive resources (see Hill & Schneider, 2006, for a review). The next sections describe the computational mechanisms by which the model achieves automaticity.

Hebbian Learning

All the learning that occurs in the model is due to changes in the strength of synapses between neurons. Computationally, these changes follow simple Hebbian rules. Hebbian algorithms are regarded as a plausible approximation to the biological dynamics of synaptic long-term potentiation (LTP) and long-term depression (T. H. Brown, Kairiss, & Keenan, 1990). In Hebbian learning, changes in synaptic weights (indicated as $\Delta w_{i,j}$) are proportional to the product of pre- and postsynaptic activations. Many variations of this principle have been proposed, differing in mathematical properties such as long-term stability and convergence (Dayan & Abbott, 2001; Gerstner & Kistler, 2002). In our model, the Hebbian rule was implemented as follows:

$$\Delta w_{i,j} = r(x_i - \langle x_i \rangle)(x_j - \langle x_j \rangle), \quad (1)$$

where r is the learning rate and $\langle x_i \rangle$ denotes neuron i 's baseline activity. This rule states that the synapses between two neurons are

strengthened whenever their firing rates conjointly exceed or fall below their baseline activation. A negative value of r was used for inhibitory projections. This turns the rule into an anti-Hebbian algorithm, which maintains the correct direction of LTP in inhibitory synapses.

Striatal interneurons are dealt with in a special way. As described above, interneurons are characterized by a high baseline activity $\langle x \rangle$, and their firing rate decreases when significant cortical patterns are detected (see Appendix A for the exact mathematical implementation). To account for this asymmetry, the opposite term $\langle x \rangle - x$ (instead of $x - \langle x \rangle$) was used whenever it referred to striatal interneurons.

Specialization in the Example Task

Within our model, this simple form of Hebbian learning is sufficient by itself to capture certain features of skill acquisition, namely, *specialization* and *automaticity*. Figure 7 illustrates the changes in the response of SN projection neurons after a different amount of repetitions of the task. The SN neurons belong to the striatal subdivision that receives projections from the prefrontal region (source) and transmits information to the vocal region (destination). In the both panels, time flows horizontally. The left panel details the activation of SN neurons at different levels of practice, corresponding to 0, 3, 6, or 9 repetitions of the same trial. Their activation reflects both the excitatory input from the cortex and the decreased inhibition from interneurons. The right panel reflects the contribution of striatal interneurons only, without the cortical component. The figure shows that, with practice, the pattern that is embedded in the synapses between interneurons and projection neurons comes to resemble the incoming cortical input.

These changes provide a preliminary basis for the development of automaticity. As long as the same information is available in the striatal interneurons' projections, the original cortical representation is not needed, and the same pattern can be used without the need for cortical processing. This fact is consistent with the drop in cortical activation that can be experimentally observed with

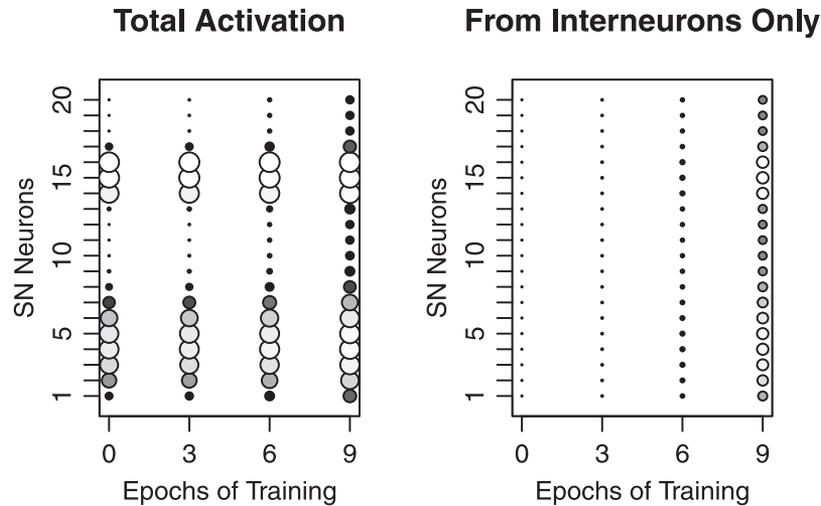


Figure 7. An example of Hebbian learning in the model. Left: Activation of an ensemble of striatonigral (SN) neurons during different levels of practice in the aural–vocal task. In the panels, time flows horizontally, and the different vertical patterns correspond to the neurons’ activations after 0, 3, 6, and 9 repetitions of the same task with the same stimulus. Right: The relative contribution of interneurons to the overall activation. With repetition, synapses between interneurons and SN units encode more of the transferred representation.

practice (Chein & Schneider, 2005; Hill & Schneider, 2006; Qin et al., 2003; Raichle et al., 1994).

Dopamine and Skill Learning

This simple associative learning mechanism cannot go very far. The modulation of dopamine in the striatum, however, can strategically direct Hebbian learning, significantly increasing the model’s learning capabilities. One important way in which practice can improve performance is by eliminating intermediate processing steps that require cognitive control. An example of such a processing step occurs between Stages B and C in the example task (see Figure 5). In this step, the prefrontal region uses the auditory stimulus as a cue to retrieve an associated response from long-term memory. With practice, this extra step can be replaced by a specialized routing operation that binds the initial auditory stimuli with their associated responses. Computationally, the idea of producing novel knowledge by creating direct stimulus–response mappings and skipping intermediate steps has been exploited in a number of production system learning algorithms. These algorithms include powerful techniques like chunking (Laird, Rosenbloom, & Newell, 1986) and production compilation (Taatgen & Lee, 2003). All these techniques have a long record of successes in modeling human learning. Furthermore, they can be seen as examples of skill learning, which is one of the memory functions of the basal ganglia (Packard & Knowlton, 2002). Therefore, it is important to show that the conditional routing model possesses similar learning capabilities.

In the model this kind of learning takes place when the basal ganglia harvest in one cycle a representation from the same cortical source (in the example, the lateral inferior prefrontal region) that was the target destination in the previous cycle. The fact that a region that has been recently used as a destination now figures among the source of routed representations is an important cue that this region has been used for some intermediate processing.

Learning to skip steps exceeds the capabilities of the simple Hebbian dynamics presented in the previous section. It can be accomplished, however, by strategically guiding the Hebbian rules. In the model, this is accomplished by the intervention of dopamine, a neurotransmitter that plays a crucial role in changing synaptic plasticity in the striatum (Calabresi et al., 2000; Wickens, Begg, & Arbuthnott, 1996). Biologically, the striatum receives dopamine from two major pathways, the mesolimbic pathway from the ventral tegmental area (VTA) and the nigrostriatal originating from the SNc.

Much is known about the response of dopamine neurons to unexpected rewards and how their bursts closely reflect the reward prediction error (Schultz, 1998, 2002). Although the routing model can, in principle, also learn from these reward-related responses, this article focuses on a different, practice-related form of learning. This type of learning depends on the release of dopamine under specific additional circumstances. In particular, we hypothesized that procedural skill acquisition is mediated by a dopamine signal carried by the nigrostriatal dopamine pathway originating in the SNc. This pathway is essential for habit formation (Faure, Haberland, Conde, & El Massioui, 2005) and has been previously included in other models of the basal ganglia (e.g., Ashby et al., 2007).

The SNc receives direct projections from the striatum, as well as indirect projection through the external pallidum and the SNr (Haber, 2003; Haber, Fudge, & McFarland, 2000; Hajos & Greenfield, 1994). Thus, its activity can be modulated by the other nuclei of the circuit. Among these afferents, there is evidence that the influence of direct SN projections is rather weak. For instance, the spontaneous activity of dopamine neurons does not change when the SN connections are removed (Hajos & Greenfield, 1994). Both the projections from the GPe (Hattori, Fibiger, & McGeer, 1975; A. D. Smith & Bolam, 1990) and the SNr (Tepper, Martin, & Anderson, 1995), on the other hand, have significant effects on the output of dopamine neurons. The model dopamine system, whose architecture is shown in Figure 8, is also controlled by the SNr/GPI

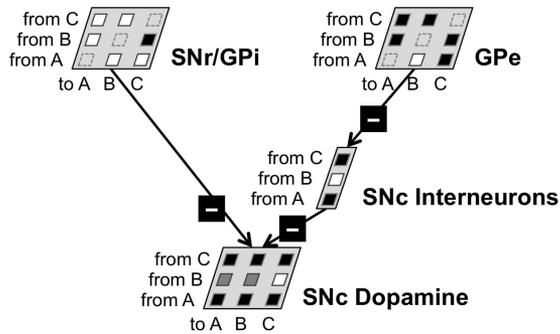


Figure 8. A close-up of the pathways controlling dopamine release in the SNc; inhibitory projections are marked with a “-.” Dopamine neurons receive two concurrent inputs, one from the SNr/GPi (left) and one from the GPe (right). This second signal is delayed by the activity of SNc interneurons, which maintain a temporary memory of the previous state. The sum of these two signals eventually activates dopamine neurons in the striatal region where a processing step can be automatized. SNr = pars reticulata of the substantia nigra; GPi = internal part of the globus pallidus; GPe = external part of the globus pallidus; SNc = pars compacta of the substantia nigra.

and GPe projections. Dopamine projection neurons in the SNc receive direct input from the SNr/GPi as well as from local SNc interneurons (Hajos & Greenfield, 1994; Juraska, Wilson, & Groves, 1970). The proposed mechanism additionally assumes that the SNc interneurons receive projections from the GPe and that their activations persist for a sufficient time to provide a delayed memory of the destinations that were used in the previous transfers (see Figure 3 and also the Discussion). This delay is crucial because it allows the SNc dopamine neurons to directly compare the sources of the representations that are being transferred (SNr/GPi) with the destinations of the previous transfers (the delayed signal from the GPe). Note that the model does not assume that the GPe signal itself is delayed with respect to the SNr; it simply assumes that the GPe signal can be temporarily maintained for comparison with the subsequent patterns of activity from the SNr. The Discussion examines other possible mechanisms that produce the same effect. Because of this delay, when the current sources figure among the previous destinations, the sum of inputs from the SNr and GPe triggers an increase in dopamine output to the striatum. This case is visually represented in Figure 8.

There are different ways of modeling the effects of dopamine. Our model adopts a solution proposed by Ashby et al. (2007), who captured the role of dopamine by adding a third term to the Hebbian rule for striatal synapses. This third term reflects the activity of dopamine neurons and expresses the biological fact that learning in the striatum is due to the interaction of pre- and postsynaptic neurons with dopamine (Ashby et al., 2007; Miller, Sanghera, & German, 1981). Within our simple Hebbian framework, this three-way interaction can be easily reproduced by including the activation x_d of dopamine neurons to Equation 1:

$$\Delta w_{i,j} = r(x_i - \langle x_i \rangle)(x_j - \langle x_j \rangle)(x_d - \langle x_d \rangle). \quad (2)$$

If we indicate $(x_d - \langle x_d \rangle)$ with the symbol d , we can rewrite Equation 2 as

$$\Delta w_{i,j} = dr(x_i - \langle x_i \rangle)(x_j - \langle x_j \rangle). \quad (3)$$

Equation 3 makes it apparent that increases and decreases in dopamine modulate synaptic plasticity by increasing or decreasing the learning rate. In fact, when dopamine falls below baseline (i.e., $d < 0$), the direction of learning can even be inverted. Although this equation does not capture all the subtleties of learning in the basal ganglia, it has the advantages of being simple and free of additional assumptions. Therefore, dopamine effects on learning were modeled by increasing or decreasing the learning rate term dr .

In addition to modulating the learning rate, dopamine directly affects the activity of striatal cells. In particular, it excites SN neurons and inhibits SP cells (Bolam et al., 2000; Nicola, Surmeier, & Malenka, 2000; see Figure 2). These differential effects permit a fine modulation of the direct and indirect pathways, which have often been included in basal ganglia models (Frank et al., 2001; O’Reilly & Frank, 2006). Less frequently modeled, but equally important, are the opposing effects of dopamine on GABAergic and cholinergic interneurons (Tepper & Bolam, 2004). Our model contains one single type of interneuron that captures properties of both. Because cholinergic interneurons also control the fast-spiking GABAergic interneurons (see Figure 2), their reaction to dopamine was taken as the dominant one. Therefore, dopamine inhibits the model interneurons. Excitatory and inhibitory effects of dopamine neurons were modeled by simply using excitatory or inhibitory projections from SNc dopamine neurons to striatal cells.

Skill acquisition depends on the dynamics between all these effects. LTP in the cortico-striatal projections increases the probability of SN neurons firing and of interneurons to deactivate in the presence of a similar pattern of cortical activity. After repeated exposures, the synapses between interneurons and projection neurons encoded the transferred representation (see Figure 7); therefore, this representation can be imposed to SN neurons even in absence of the original cortical input from the corresponding region.

Skill Learning in the Example Task

The simple aural-vocal task described in the previous section is useful for demonstrating these effects of learning. In the simple model outlined above, the correct response to a tone had to be selected from long-term memory (see Figure 5, Stages B and C). This intermediate step can be omitted with practice. The lateral prefrontal cortex figures as the destination of the first routing operation (see Figure 5, top right panel) and as the source of the second (see Figure 5, bottom right panel). Therefore, the redundant step can be detected in the convergent pathways on the SNc neurons. In turn, this triggers dopamine release in the striatum, initiating the learning process described above. Figure 9 illustrates how the model performs the task after the learning step has happened a sufficient number of times to allow the newly learned routing operation to fire. The figure illustrates how the new operation routes an immediate response to the vocal region when presented with the original stimulus. Thus the model transitions from the initial stage to the final stage in Figure 9 without the intermediate stages in Figure 5.

It is interesting that this new operation can be fired at the same time as the original response to the stimulus. This is shown in the right panel of Figure 9. This panel depicts the new pattern of thalamic activation. The new thalamic pattern includes activations in the subdivisions that were previously found in two separate

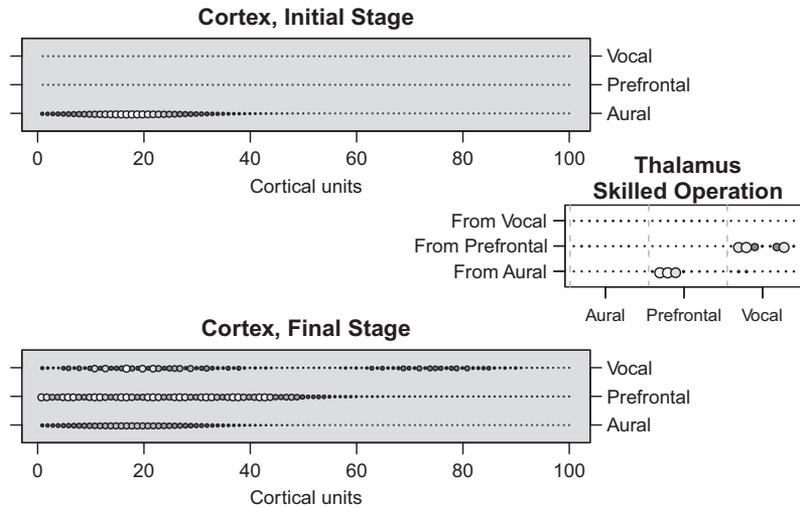


Figure 9. Model performance in the example task, after dopamine-mediated learning has occurred. Synaptic plasticity in the striatum resulted in anticipated activation of the thalamic sector in response to the presentation of the auditory stimulus. Activation in this sector was previously triggered by a vocal response retrieved in the second region. As a consequence, a response is transmitted to third cortical region (*vocal*) at the same time the cue is routed to the second region (*retrieval*).

routing operations (compare to the right panels of Figure 5). In a way, learning had the effect of “shifting back” the original second operation so that it can be applied in advance. The anticipation of striatal activation is consistent with the reorganization of firing patterns in the striatum following habit learning (Jog et al., 1999).

Note that reorganization of striatal activity did not completely remove the transfer to prefrontal cortex. Rather, it reorganized it so that it can occur in parallel with the speeded-up vocal response. Since the prefrontal contribution in response selection is eventually irrelevant to the vocal response, further practice and the establishment of cortico–cortical connectivity will eventually eliminate it. Therefore, the model predicts that in skill acquisition the reorganization of cortical activity follows the reorganization of striatal activity. This prediction is consistent with the different rates by which striatal and prefrontal cortex cells learn to respond to stimuli with practice (Pasupathy & Miller, 2005). Correspondingly, Anderson (2007) reported data showing that, with extensive practice, task-related metabolic activity in the prefrontal region drops to baseline levels.

This form of learning relies on the strategic release of dopamine after the execution of the second operation. We have outlined one possible biological mechanism by which this could happen. It should be noted, however, that other mechanisms of dopamine control could obtain similar results. One of these mechanisms is the simple release of dopamine according to unpredicted rewards (Schultz, 1998, 2002). Thus, an increase of the d term after the second operation can also be triggered by the initial reward generated by succeeding in the task.

Time Course of Striatal and Dopamine Activity

The learning dynamics that arise in the model can be used to explain some contrasting findings in the literature. Certain lesions in the basal ganglia have been found to disrupt the execution of

highly practiced tasks, suggesting that the basal ganglia serve as the ultimate repository of skills and habits. Other studies, however, have found that basal ganglia lesions do not affect highly trained skills, suggesting that the basal ganglia are required for acquiring skills but do not constitute their ultimate repository (e.g., Hikosaka, Rand, Miyachi, & Miyashita, 1995; Miyachi, Hikosaka, Miyashita, Karadi, & Rand, 1997; Packard & McGaugh, 1996). A similar paradox arises in single-cell recordings, where some studies have found that striatal activity increases while habits are being acquired (e.g., Jog et al., 1999), whereas others have found the opposite pattern (e.g., Carelli, Wolske, & West, 1997).

One way to reconcile these results is assuming that the transition from a novel to a practiced task occurs in two steps. During an initial learning phase, a skill is temporarily encoded within the basal ganglia. With time, however, the skill can be encoded in the cortico–cortical pathways and no longer relies on the basal ganglia. This is consistent with the popular view that the basal ganglia “train” the cortex (Graybiel, 2005; Pasupathy & Miller, 2005). According to this view, an increase in basal ganglia activity should be expected during the first stage, as a result of the circuit encoding more information. A decrease, however, can be expected in the second stage, as a result of information being progressively transferred to the cortex. This pattern has been confirmed in a neuroimaging study of procedural learning (Hubert et al., 2007).

An example of the decrease of basal ganglia involvement with practice is provided by Ashby et al. (2007). Their results rely on the decay of synaptic strengths in cortico–striatal cells when dopamine ceases to be released, which occurs when transmission along cortico–cortical pathways happens faster than through the basal ganglia loops. Although their model can provide insights into this second step of proceduralization, our model can provide insights into the first. In particular, the anticipated activation of SN projection neurons (as illustrated in Figure 9) can provide the

neural basis for the metabolic increase of activation during the establishment of habits.

To investigate this possibility we ran a new series of simulations. During these simulations, the model performed the aural–vocal task until its performance moved from the original two-operation procedure (as illustrated in Figure 5) to the skilled procedure (as illustrated in Figure 9). The sum of activity across all the model striatal projection neurons was recorded at the very beginning and the very end of the learning phase. To prevent further practice from contaminating the results, Hebbian learning was disabled in those trials where striatal activation was recorded. The light gray line in Figure 10 plots the average amount of striatal activity against the experience with the task. The values in the figure have been normalized, using the average amount in the novel (i.e., initial and unskilled) condition as a baseline. The figure also plots the change of dopamine release with practice (dark gray line). This quantity was estimated by summing up the activation values of all the SNc dopamine neurons. Because the initial, unskilled execution of the task triggers dopamine release, the baseline for this parameter was taken from the skilled condition. It is clear that in the transition from novel to skilled behavior, dopamine follows a pattern that is opposite to that of striatal activity, eventually returning to the baseline level with practice.

Summary

This section has illustrated the learning capabilities of the model. In particular, it has shown how the model, with practice, can encode internally certain representations that were originally routed from the cortex. This accounts for the proceduralization and the specialization of responses. It was accomplished by means of simple Hebbian learning, which is a biologically plausible learning rule. When coupled with the effects of dopamine, the Hebbian rule triggers more complex dynamics. Eventually, these dynamics enable the acquisition of new skills that skip intermediate steps in series of information transfers. The elimination of redundant steps

during the learning phase accounts for practice speedup and automaticity. Finally, the time course of striatal activity is consistent with important established results in the field of habit learning.

Detecting the appropriate conditions for dopamine release requires a comparison of the current sources and the previous destinations in the SNc. In turn, this requires maintaining a delayed version of the previous destinations in the SNc interneurons. Although this account is partially speculative, many other models have adopted and defended similar mechanisms that compare the current signals from the direct pathway with a delayed signal from the indirect pathway (e.g., Barto, 1995; see Joel et al., 2002, for a review). It should also be noted that similar results could be accomplished by other mechanisms that regulate dopamine release. For instance, reward-related changes in dopamine also follow a pattern similar to that in Figure 10, with dopamine decreasing as a task becomes more practiced and rewards become predictable (Schultz, 1998, 2002).

Parkinson's and Huntington's Diseases

A computational model of the basal ganglia should address the two signature disorders of the circuit: Parkinson's and Huntington's diseases. Both pathologies compromise the functionality of the basal ganglia circuit but have different causes and are associated with different symptoms. This section addresses them separately.

Parkinson's Disease

Parkinson's disease is caused by the death of dopaminergic neurons in the SNc, which drastically reduces the dopamine supply to the striatum (see Jankovic, 2008, for a review). The four cardinal symptoms of this disease are tremor, rigidity, postural problems, and inability to initiate voluntary movements (akinesia), which eventually leads to paralysis (Jankovic, 2008).

As previously mentioned, SN and SP projection neurons express different types of dopamine receptors. Under normal circumstances, dopamine excites the SN cells but inhibits SP neurons (Gerfen et al., 1990; Nicola et al., 2000; see Figure 2). A decrease in striatal dopamine, therefore, has a net excitatory effect on SP neurons (because of the decreased inhibition) and a net inhibitory effect on SN cells (because of the decreased excitatory input). SN and SP neurons originate the direct and indirect pathways, respectively. This consideration underlies the development of the classic view of the basal ganglia (Albin et al., 1989; DeLong, 1990). According to this view, the output of the basal ganglia depends on the balance between the direct pathway, which disinhibits motor programs, and the indirect pathway, which inhibits them by exciting the output nuclei. Within this brake–accelerator (Graybiel, 2000) framework, the symptoms of Parkinson's disease can be explained as an imbalance between a pathologically weak direct pathway and an abnormally strong indirect pathway (Albin et al., 1989). We have argued above that our conditional routing model can be seen as a generalization of the brake–accelerator view of the basal ganglia. Therefore, it is important to show that it can provide an explanation for the same phenomena the classic framework was designed to explain.

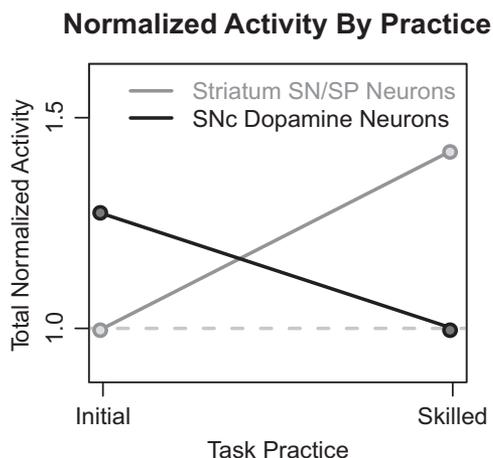


Figure 10. Relative increase of mean projection neuron activation (light gray line and circles) and corresponding decrease of dopamine neuron activation (dark gray line and circles) as the model transitions from novel to skilled performance in the aural–vocal task. SN = striatonigral; SP = striatopallidal; SNc = pars compacta of the substantia nigra.

Simulation

Because dopaminergic neurons are explicitly modeled in the SNc, a straightforward way to mimic this pathology is to simulate damage to these units. A decrease in dopamine, therefore, can be modeled as increase in the input of SN units, and a corresponding decrease in the input of SP cells. In the model, this translates into stronger inhibition of possible destinations and weaker encoding of source representations. This reduction has two important consequences. The first one is that the proper source is not strongly propagated along the direct pathway and the proper destination is still somewhat inhibited by the indirect pathway. Figure 11 illustrates such a case. In the figure, the top image represents the pattern of activation in the thalamic subdivisions when an operation is executed under normal conditions. It can be seen that only one subdivision contains active cells. When dopamine is lowered, activity of these cells is reduced because of the imbalance between the two pathways. This can be observed in the left middle panel of Figure 11. The bottom left panel represents the landscape of thalamic activation when dopamine has been drastically lowered and the same

operation is attempted. In this example, the net input to the thalamic output falls below the activation threshold, blocking the delivery of signals to the cortex. If the transferred pattern represents a specific motor command (like in the final stage of the example task; see Figure 5), then the model remains frozen in a condition resembling akinesia, unable to initiate a proper movement.

A dysfunction of the routing mechanism is not limited to motor programs; it extends to other regions, causing a wide range of nonmotor impairments that depend on the functions of the destination cortical regions. This is consistent with the existence of widespread cognitive deficits even at the early stages of the disease (Levin, Llabre, & Weiner, 1989). Given the fact that basal ganglia projections mainly target the frontal lobe, a relationship between Parkinson's disease and executive function disorders is expected and has been reported in neuropsychological studies (Muslimovic, Post, Speelman, & Schmand, 2005; Owen, 2004).

In the model, the disruption of dopaminergic input to the striatum has a second, important consequence. Dopamine depletion prevents the acquisition of new skills, which crucially rely on the

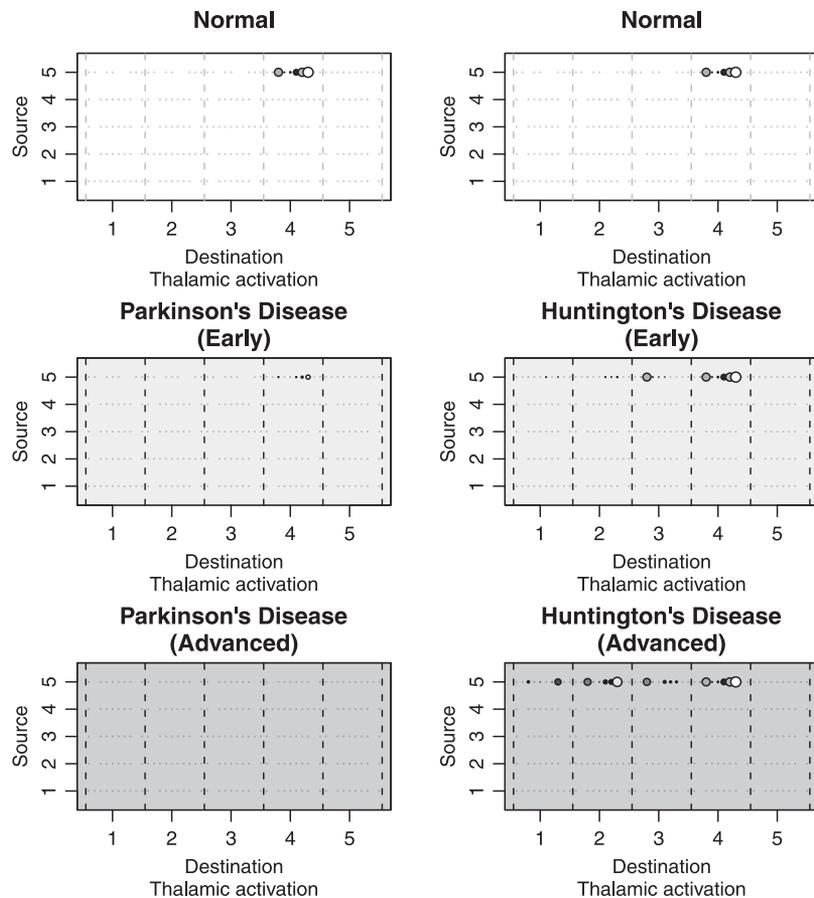


Figure 11. An illustration of how the model simulates Parkinson's and Huntington's diseases. Left: Parkinson's disease was simulated by lowering the amount of dopamine input to striatal projection neurons. As the amount of dopamine decreases, the thalamic output to the cortex becomes increasingly weaker. Right: Huntington's disease was simulated by randomly disabling striatopallidal (SP) neurons in the striatum. As a consequence, damaged representations are misrouted to other destination regions.

modulation of dopamine release to a target striatal compartment. Several studies have documented a specific inability of Parkinson's patients to acquire new skills (Jackson et al., 1995; Knowlton et al., 1996). Neuroimaging studies have provided further insight into this impairment. For example, Dagher et al. (2001) compared the brain activity patterns of healthy controls and mildly affected patients solving a set of problems with the Tower of London puzzle, a task designed to assess planning capabilities (Shallice, 1982). Participants were chosen so that the two groups had comparable behavioral performance. Compared to controls, patients exhibited less activity in the caudate nucleus but increased activity in the hippocampus. This suggests that Parkinson's patients were compensating for their inability to learn new procedures by relying on previously acquired memories of task states (Dagher et al., 2001).

Huntington's Disease

Huntington's disease is a genetic progressive neurodegenerative disorder. The responsible gene causes the death of striatal neurons, although the precise cellular mechanisms by which this happens are still poorly understood (Walker, 2007). The disease does not target all striatal cells equally. SP neurons are the most vulnerable, whereas SN cells and interneurons are less involved (Albin et al., 1992; Glass, Dragunow, & Faull, 2000; Reiner et al., 1988). The most obvious physical symptoms consist of jerky, disordered, and uncontrolled movements and tics, collectively known as chorea (Walker, 2007).

Early motor symptoms of Huntington's disease are opposite and complementary to those of Parkinson's disease. This was one of the observations that guided the brake–accelerator view of the basal ganglia, where the direct and indirect pathways exert opposite and complementary effects on the basal ganglia output nuclei. According to this view, Huntington's disease originates from an abnormal weakness of the indirect pathway, which fails to inhibit unsolicited motor or cognitive behaviors (Albin et al., 1989).

Simulation

Huntington's disease was simulated by randomly disabling SP cells in the model striatum. This decreases the inhibitory strength of the indirect pathway, the most obvious consequence of which is that transferred contents are misrouted. The bottom right image in Figure 11 illustrates such an example. In the figure, the same operation is performed that was previously used to test normal and Parkinson's conditions, only this time a number of SP neurons have been eliminated. As a result, impoverished representations are now mistakenly broadcast to different destination regions. This excess of activation provides a basis for involuntary and uncontrolled movements, whereas the degradation of transmitted patterns accounts for the deterioration of fine cognitive and motor abilities.

As in the case of Parkinson's disease, this impasse of the circuit is general and not limited to the delivery of motor programs. Correspondingly, Huntington's patients are affected by a number of cognitive as well as motor impairments. Many of the compromised cognitive abilities (e.g., planning, working memory, and set shifting) are related to the functions of the frontal lobes, which are

in fact the main target of basal ganglia projections (see Lawrence, Sahakian, & Robbins, 1998).

In addition to a difference in their motor symptoms, the model also predicts that Parkinson and Huntington's patients should differ in the domain of skill learning. Parkinson's disease is caused by the death of cells in the SNc that, according to the model, are responsible for triggering the learning signal. In the case of Huntington's, however, SNc cells are spared and the learning mechanism should be intact and available. This potential dissociation is counterbalanced by two factors. First, part of the skill acquisition process consists of the striatum learning specific patterns that were previously available in the cortex. In Huntington's patients, the striatum itself is compromised, and the extent of its damage constrains its ability to internalize and reproduce proper cortical representations. Second, skill acquisition only proceeds from an initial stage where task-relevant information is properly represented in the cortical regions (as in the simplified initial stage of the example task; see Figure 5), but the delivery of proper representations to the appropriate cortical regions is also affected in the case of striatal damage. Given these considerations, one should not expect more than a difference in degrees between the learning capabilities in the two conditions. Indeed, some experimental findings indicate that although Huntington's patients do exhibit skill-learning impairments, their deficits are less severe than in the case of Parkinson's (Sprenkelmeyer, Canavan, Lange, & Homberg, 1995). This finding is even more remarkable when one considers that the cognitive effects of the former are more severe and impairing than those of the latter.

Discussion

This article described a novel model of the basal ganglia, according to which the basal ganglia control the routing of information between cortical areas. The transfer is made possible by assuming that the basal ganglia operate by binding pairs of source and destination regions. This is made possible by the fact that the basal ganglia maintain and reflect some features of cortical topology. The outlined mechanism is general because it allows actions to be uniformly represented as source–destination bindings across different domains. It is also flexible in that it allows the basal ganglia to dynamically change the way information flows between cortical regions. Finally, the model describes a possible mechanism by which dopamine interacts with striatal neurons to enable skill acquisition. This mechanism is based on simple Hebbian learning rules and accounts for the development of new source–destination bindings and automaticity.

By deciding which information is transferred to and processed by the appropriate cortical region, the basal ganglia can articulate the ongoing processing activities of the brain in an ordered sequence. Such a function does not constitute a complete characterization of the role of the basal ganglia in human cognition. However, it is general enough to explain a number of cognitive contributions of the basal ganglia, especially in memory and higher level cognition. Furthermore, the hypothesized function emerges naturally from the structure of the circuit. Importantly, the model shows how the basal ganglia, by flexibly routing information, provide a neural instantiation for the computational power of a production system.

Relationship to Basal Ganglia Physiology

Our model stresses the importance of striatal interneurons, which have been neglected in previous modeling attempts (e.g., Ashby et al., 2005, 2007; Frank et al., 2001). In the model, they play a significant role in learning and in regulating the activity of projection neurons.

The routing model relies on the assumption that certain nuclei of the circuit, including the striatal matrix, GPe, SNr/GPi, and thalamic relay nuclei, have a two-level organization, which reflects topology of the cortical regions and their cortico-cortical projections. This organization has been visualized like a source-destination matrix in Figures 3, 4, 5, 8, 9, and 11. This organization is incompatible with the view of the basal ganglia as forming parallel closed loops that connect a frontal area to itself (e.g., Alexander et al., 1986). However, many authors have recently argued that the basal ganglia also form open loops (Joel & Weiner, 1994). It has been reported that striatum neurons receive projections from different connected areas (Parthasarathy et al., 1992) and that the thalamic relay nuclei of the basal ganglia form both reciprocal as well as nonreciprocal connections (McFarland & Haber, 2002).

In our model, closed loops represent a special case where source and destination coincide (i.e., regions along the diagonal in Figure 3). Because the focus of the article was on the transfer of information, they do not play a significant role in the simulations, and ensembles of cells belonging to closed loops have been grayed out in Figure 3. However, closed loops might play an important role in the biological circuit. In fact, they might even recruit larger neuronal ensembles than the open loops where source and destination differ. In any case, our model can accommodate the existence of closed loops without losing its functions and generality.

The Basal Ganglia and the Cortex

Central to the routing model is the idea that subcortical connections are capable of determining the state of cortical regions. The overwhelming majority of cortical afferents, however, are represented by cortico-cortical connections (Braitenberg & Schüz, 1991). This asymmetry can be reconciled with our model only if subcortical afferents are strong enough to affect the states of cortical neurons. In the case of thalamic connections, it has been shown that they are capable of driving the activity of cortical areas, even in the absence of cortical amplification of the signal (Bruno & Sakmann, 2006). As for the specific basal ganglia-thalamo-cortical connections, an indirect confirmation of their importance comes from investigations of cortical synchronization (as measured by functional connectivity between cortical regions) in Parkinson's disease. Two independent studies have shown that, as the disease progresses and basal ganglia function become more impaired, the synchronization between cortical regions increases (Moazami-Goudarzi, Sarnthein, Michels, Moukhtieva, & Jeanmonod, 2008; Stoffers et al., 2008). This suggests that, in normal conditions, the basal ganglia interfere with the spontaneous reverberations of signals across cortical regions. This hypothesis is consistent with their proposed role in transferring and superimposing specific representations onto the target regions of the cortex.

It is well known that the basal ganglia receive projections from all of the neocortex but project back almost exclusively to the

frontal areas. Such an asymmetry may seem problematic in light of the proposed information-routing function, because it implies that most source areas are not possible destinations. One possibility is that this uneven organization of the basal ganglia projections results from the functional organization of the brain. In particular, this organization might reflect the large-scale differences in functional roles of brain regions, with the frontal parts more engaged in controlling behavior by holding task-specific representations and the posterior parts predominantly engaged in processing or representing sensory and perceptual information. Assuming this distinction, it seems rational that the basal ganglia, although gathering information from all cortical regions, feed the selected signals to the prefrontal cortex in order to avoid interference with ongoing posterior processing of sensory information. In fact, Atallah, Frank, and O'Reilly (2004) have previously argued for an architectural partition of the brain based on the functional macro-organization of the cortex and its ensuing computational tradeoffs, in which the basal ganglia and the prefrontal cortex form a joint subsystem.

Current Limitations

The routing model depends on a number of assumptions about the basal ganglia physiology. The most problematic are perhaps the assumptions that underlie the dopamine-mediated skill-learning mechanism. Many authors (e.g., Frank & Claus, 2006) have considered the dopamine signal as a uniform learning signal that is best represented by a simple scalar quantity. Other models, however, do make use of a spatially differentiated dopamine signal (e.g., Frank et al., 2001; O'Reilly & Frank, 2006). In our model this signal is also spatially differentiated, and it is represented as vector of values that can be different for different subdivisions of the striatum (see Figure 8). The nigrostriatal connections are, in fact, organized into modules corresponding to different striatal subdivisions (e.g., Haber, 2003; Haber et al., 2000). This organization makes sense only when the dopamine input can vary between different striatal subdivisions. Also, most studies of dopamine modulation investigate the effects of rewards, not of practice. The skill-learning mechanism predicts that the dopamine signal needs to be inhomogeneous only for practice-related changes.

A related issue concerns the role of SNc interneurons in maintaining the signal corresponding to the previous destinations. We have no direct evidence to support the specific mechanism implemented in the model, but there are other possible mechanisms that can achieve the same function. For instance, the inhibitory signal from the GPe could be maintained longer on a dopamine cell because of the longer activity of its specific neurotransmitter (e.g., Tepper et al., 1995).

In both the example task and the general performance simulations, the model was trained to perform a number of initial routing operations. The training procedure used the CHL algorithm, which requires only local computations. The use of a special algorithm was required to provide the model with an initial set of actions. An initial set is required because the model's internal learning algorithm is based on practice and can proceed only from operations that the model already knows to execute. The model itself is agnostic on how this initial set is acquired. Two different possibilities can be suggested. First, an initial set of actions can be acquired by means of reward-related learning. This form of learn-

ing is obviously associated with the dopamine signal and has been previously integrated within gating models of the basal ganglia (Frank & Claus, 2006; O'Reilly & Frank, 2006). There is no incompatibility between a reward-based learning mechanism and our model's skill acquisition procedure. The two learning systems are, in fact, complementary. This is also the reason why we adopted the CHL algorithm as a default training procedure (see Appendix B for a discussion of the relationship between CHL and reward signals).

A second possibility is that an initial set of task-specific actions can be learned from the explicit representations of task rules that participants can hold in working memory. For instance, in the example aural-vocal task, participants might be initially rehearsing the actions they have to perform as a way to guide their performance of the novel task step by step. Consciously following these rules for a short period of time would result in acquiring task-specific rules that directly reflect the rule steps. This is another form of skill acquisition, and therefore it is within the reach of our model's learning capabilities. In fact, similar mechanisms that essentially interpret an internal task representation have been used to model the acquisition of procedural skills from instructions in production systems (e.g., Taatgen, 2005).

Relationship to Other Models

The routing model belongs to a recent strain of models that describe how the basal ganglia select and gate information from the cortex (Amos, 2000; Frank et al., 2001; Gurney et al., 2001; O'Reilly & Frank, 2006). A different tradition of models has concentrated on the role of the basal ganglia in maintaining (instead of updating) working memory (Ashby et al., 2005; Monchi et al., 2000). In our model, maintaining a piece of information in working memory can be seen as a special case of routing where source and destination coincide. Therefore, the routing model illustrates how working memory update and maintenance are not opposite functions of the basal ganglia but rather can be accounted for within a single framework.

The model uses the routing framework to account for the acquisition of skills. In particular, the model proposes that skills are acquired by temporarily encoding cortical representations within the synaptic patterns of the striatum. This particular form of learning has been seldom investigated before in biological model of the basal ganglia. Notably, Ashby et al. (2007) have proposed a model of categorization that relies on very similar principles. Their model, however, proposes that acquired skills are directly and permanently encoded in cortico-cortical representation. We propose that, in an earlier step during skill learning, intermediate representations are temporarily stored in the striatum. Within this framework, the routing model can be seen as a detailed specification of the first stage in a dual-step process, whereas Ashby et al.'s focuses on the second stage.

Many other models have dealt with a different form of learning in the basal ganglia, that is, learning that is driven by reward (Barto, 1995; Houk et al., 1995; O'Reilly & Frank, 2006; Suri & Schultz, 1998, 1999). This form of learning differs from skill acquisition in that it is driven by reward instead of practice and does not imply the development of automaticity. Also, reward-based learning is typically modeled as being modulated by the patch compartment of the striatum (Barto, 1995; Houk et al., 1995;

see also Joel et al., 2002), whereas the routing model focuses on the matrix compartment.

The routing model can be characterized in a way that is very similar to a production system. This fact is important for two reasons. First, it establishes a bridge between biologically inspired models of the brain and a common high-level characterization of cognition. Second, it provides a way to capture the functions of the basal ganglia with a formalism that is well known and widely adopted in the cognitive science community, providing a new link between structural properties of the brain and their functional characterizations in terms of computation.

References

- Ajjanagadde, V., & Shastri, L. (1991). Rules and variables in neural nets. *Neural Computation*, 3, 121–134.
- Albin, R. L., Reiner, A., Anderson, K. D., Dure, L. S. T., Handelin, B., Balfour, R., . . . Young, A. B. (1992). Preferential loss of striato-external pallidal projection neurons in presymptomatic Huntington's disease. *Annals of Neurology*, 31, 425–430.
- Albin, R. L., Young, A. B., & Penney, J. B. (1989). The functional anatomy of basal ganglia disorders. *Trends in Neurosciences*, 12, 366–375.
- Alexander, G. E., DeLong, M. R., & Strick, P. L. (1986). Parallel organization of functionally segregated circuits linking basal ganglia and cortex. *Annual Review of Neuroscience*, 9, 357–381.
- Amit, D. J. (1993). *Modeling brain function*. New York, NY: Cambridge University Press.
- Amos, A. (2000). A computational model of information processing in the frontal cortex and basal ganglia. *Journal of Cognitive Neuroscience*, 12, 505–519.
- Anderson, J. R. (1983). *The architecture of cognition*. Mahwah, NJ: Erlbaum.
- Anderson, J. R. (2005). Human symbol manipulation within an integrated cognitive architecture. *Cognitive Science*, 29, 313–341.
- Anderson, J. R. (2007). *How can the human mind occur in the physical universe?* (1st ed.). New York, NY: Oxford University Press.
- Anderson, J. R., Albert, M. V., & Fincham, J. M. (2005). Tracing problem solving in real time: fMRI analysis of the subject-paced Tower of Hanoi. *Journal of Cognitive Neuroscience*, 17, 1261–1274.
- Anderson, J. R., Fincham, J. M., Qin, Y., & Stocco, A. (2008). A central circuit of the mind. *Trends in Cognitive Sciences*, 12, 136–143.
- Anderson, J. R., Taatgen, N. A., & Byrne, M. D. (2005). Learning to achieve perfect timesharing: Architectural implications of Hazeltine, Teague, and Ivry (2002). *Journal of Experimental Psychology: Human Perception and Performance*, 31, 749–761.
- Apicella, P. (2002). Tonicly active neurons in the primate striatum and their role in the processing of information about motivationally relevant events. *European Journal of Neuroscience*, 16, 2017–2026.
- Apicella, P. (2007). Leading tonically active neurons of the striatum from reward detection to context recognition. *Trends in Neurosciences*, 30, 299–306.
- Arbib, M. A. (2003). Towards a neurally-inspired computer architecture. *Natural Computing*, 2, 1–46.
- Ashby, F. G., Ell, S. W., Valentin, V. V., & Casale, M. B. (2005). FROST: A distributed neurocomputational model of working memory maintenance. *Journal of Cognitive Neuroscience*, 17, 1728–1743.
- Ashby, F. G., Ennis, J. M., & Spiering, B. J. (2007). A neurobiological theory of automaticity in perceptual categorization. *Psychological Review*, 114, 632–656.
- Atallah, H. E., Frank, M. J., & O'Reilly, R. C. (2004). Hippocampus, cortex, and basal ganglia: Insights from computational models of complementary learning systems. *Neurobiology of Learning and Memory*, 82, 253–267.

- Bar-Gad, I., Havazelet-Heimer, G., Goldberg, J. A., Ruppin, E., & Bergman, H. (2000). Reinforcement-driven dimensionality reduction—A model for information processing in the basal ganglia. *Journal of Basic and Clinical Physiology and Pharmacology*, *11*, 305–320.
- Barto, A. G. (1995). Adaptive critics and the basal ganglia. In J. C. Houk, J. L. Davis, & D. G. Beiser (Eds.), *Models of information processing in the basal ganglia* (pp. 215–232). Cambridge, MA: MIT Press.
- Beiser, D. G., & Houk, J. C. (1998). Model of cortical–basal ganglionic processing: Encoding the serial order of sensory events. *Journal of Neurophysiology*, *79*, 3168–3188.
- Berns, G., & Sejnowski, T. J. (1998). A computational model of how the basal ganglia produce sequences. *Journal of Cognitive Neuroscience*, *10*, 108–121.
- Bevan, M. D., Booth, P. A., Eaton, S. A., & Bolam, J. P. (1998). Selective innervation of neostriatal interneurons by a subclass of neuron in the globus pallidus of the rat. *Journal of Neuroscience*, *18*, 9438–9452.
- Bevan, M. D., & Wilson, C. J. (1999). Mechanisms underlying spontaneous oscillation and rhythmic firing in rat subthalamic neurons. *Journal of Neuroscience*, *19*, 7617–7628.
- Bienenstock, E. L., Cooper, L. N., & Munro, P. W. (1982). Theory for the development of neuron selectivity: Orientation specificity and binocular interaction in visual cortex. *Journal of Neuroscience*, *2*, 32–48.
- Bolam, J. P., Hanley, J. J., Booth, P. A., & Bevan, M. D. (2000). Synaptic organisation of the basal ganglia. *Journal of Anatomy*, *196*, 527–542.
- Braitenberg, V., & Schüz, A. (1991). *Anatomy of the cortex: Statistics and geometry*. Heidelberg, Germany: Springer Verlag.
- Braver, T. S., Cohen, J. D., Nystrom, L. E., Jonides, J., Smith, E. E., & Noll, D. C. (1997). A parametric study of prefrontal cortex involvement in human working memory. *NeuroImage*, *5*, 49–62.
- Brown, J., Bullock, D., & Grossberg, S. (1999). How the basal ganglia use parallel excitatory and inhibitory learning pathways to selectively respond to unexpected rewarding cues. *Journal of Neuroscience*, *19*, 10502–10511.
- Brown, L. L., Schneider, J. S., & Lidsky, T. I. (1997). Sensory and cognitive functions of the basal ganglia. *Current Opinion in Neurobiology*, *7*, 157–163.
- Brown, T. H., Kairiss, E. W., & Keenan, C. L. (1990). Hebbian synapses: Biophysical mechanisms and algorithms. *Annual Review of Neuroscience*, *13*, 475–511.
- Bruno, R. M., & Sakmann, B. (2006, June 16). Cortex is driven by weak but synchronously active thalamocortical synapses. *Science*, *312*, 1622–1627.
- Calabresi, P., Gubellini, P., Centonze, D., Picconi, B., Bernardi, G., Chergui, K., . . . Greengard, P. (2000). Dopamine and cAMP-regulated phosphoprotein 32 kDa controls both striatal long-term depression and long-term potentiation, opposing forms of synaptic plasticity. *Journal of Neuroscience*, *20*, 8443–8451.
- Carelli, R. M., Wolske, M., & West, M. O. (1997). Loss of lever press-related firing of rat striatal forelimb neurons after repeated sessions in a lever pressing task. *Journal of Neuroscience*, *17*, 1804–1814.
- Carman, J. B., Cowan, W. M., & Powell, T. P. (1963). The organization of cortico-striate connexions in the rabbit. *Brain*, *86*, 525–562.
- Chein, J. M., & Schneider, W. (2005). Neuroimaging studies of practice-related change: fMRI and meta-analytic evidence of a domain-general control network for learning. *Cognitive Brain Research*, *25*, 607–623.
- Chesselet, M. F., & Delfs, J. M. (1996). Basal ganglia and movement disorders: An update. *Trends in Neurosciences*, *19*, 417–422.
- Chevalier, G., & Deniau, J. M. (1990). Disinhibition as a basic process in the expression of striatal functions. *Trends in Neurosciences*, *13*, 277–280.
- Clover, D. M., Dum, R. P., & Strick, P. L. (2005). Basal ganglia and cerebellar inputs to 'AIP.' *Cerebral Cortex*, *15*, 913–920.
- Cohen, N. J., & Squire, L. R. (1980, October 10). Preserved learning and retention of pattern-analyzing skill in amnesia: Dissociation of knowing how and knowing that. *Science*, *210*, 207–210.
- Contreras-Vidal, J. L., & Schultz, W. (1999). A predictive reinforcement model of dopamine neurons for learning approach behavior. *Journal of Computational Neuroscience*, *6*, 191–214.
- Cooper, J. A., Sagar, H. J., Doherty, S. M., Jordan, N., Tidswell, P., & Sullivan, E. V. (1992). Different effects of dopaminergic and anticholinergic therapies on cognitive and motor function in Parkinson's disease: A follow-up study of untreated patients. *Brain*, *115*, 1701–1725.
- Dagher, A., Owen, A. M., Boecker, H., & Brooks, D. J. (2001). The role of the striatum and hippocampus in planning: A PET activation study in Parkinson's disease. *Brain*, *124*, 1020–1032.
- Danker, J. F., Gunn, P., & Anderson, J. R. (2008). A rational account of memory predicts left prefrontal activation during controlled retrieval. *Cerebral Cortex*, *18*, 2674–2685.
- Dayan, P., & Abbott, L. F. (2001). *Theoretical neuroscience: Computational and mathematical modeling of neural systems*. Cambridge, MA: MIT Press.
- DeLong, M. R. (1990). Primate models of movement disorders of basal ganglia origin. *Trends in Neurosciences*, *13*, 281–285.
- Destrebecqz, A., & Cleeremans, A. (2001). Can sequence learning be implicit? New evidence with the process dissociation procedure. *Psychonomic Bulletin & Review*, *8*, 343–350.
- Dominey, P., Arbib, M. A., & Joseph, J. (1995). A model of corticostriatal plasticity for learning oculomotor associations and sequences. *Journal of Cognitive Neuroscience*, *7*, 311–336.
- Doyon, J., Owen, A. M., Petrides, M., Sziklas, V., & Evans, A. C. (1996). Functional anatomy of visuomotor skill learning in human subjects examined with positron emission tomography. *European Journal of Neuroscience*, *8*, 637–648.
- Faure, A., Haberland, U., Conde, F., & El Massioui, N. (2005). Lesion to the nigrostriatal dopamine system disrupts stimulus–response habit formation. *Journal of Neuroscience*, *25*, 2771–2780.
- Fernandez-Ruiz, J., Wang, J., Aigner, T. G., & Mishkin, M. (2001). Visual habit formation in monkeys with neurotoxic lesions of the ventrocaudal neostriatum. *Proceedings of the National Academy of Sciences, USA*, *98*, 4196–4201.
- Fiorillo, C. D., Tobler, P. N., & Schultz, W. (2003, March 21). Discrete coding of reward probability and uncertainty by dopamine neurons. *Science*, *299*, 1898–1902.
- Frank, M. J., & Claus, E. D. (2006). Anatomy of a decision: Striato-orbitofrontal interactions in reinforcement learning, decision making, and reversal. *Psychological Review*, *113*, 300–326.
- Frank, M. J., Loughry, B., & O'Reilly, R. C. (2001). Interactions between frontal cortex and basal ganglia in working memory: A computational model. *Cognitive, Affective & Behavioral Neuroscience*, *1*, 137–160.
- Frank, M. J., Samanta, J., Moustafa, A. A., & Sherman, S. J. (2007, October 25). Hold your horses: Impulsivity, deep brain stimulation, and medication in Parkinsonism. *Science*, *318*, 1309–1312.
- Frank, M. J., Seeberger, L. C., & O'Reilly, R. C. (2004, December 10). By carrot or by stick: Cognitive reinforcement learning in Parkinsonism. *Science*, *306*, 1940–1943.
- Gabrieli, J. D. E. (1998). Cognitive neuroscience of human memory. *Annual Review of Psychology*, *49*, 87–115.
- Gabrieli, J. D. E., Singh, J., Stebbins, G. T., & Goetz, C. G. (1996). Reduced working memory span in Parkinson's disease: Evidence for the role of a frontostriatal system in working and strategic memory. *Neuropsychology*, *10*, 322–332.
- Gerfen, C. R. (1984, October 4). The neostriatal mosaic: Compartmentalization of corticostriatal input and striatonigral output systems. *Nature*, *311*, 461–464.
- Gerfen, C. R. (1992). The neostriatal mosaic: Multiple levels of compartmental organization in the basal ganglia. *Annual Review of Neuroscience*, *15*, 285–320.

- Gerfen, C. R., Baimbridge, K. G., & Miller, J. J. (1985). The neostriatal mosaic: Compartmental distribution of calcium-binding protein and parvalbumin in the basal ganglia of the rat and monkey. *Proceedings of the National Academy of Sciences, USA*, *82*, 8780–8784.
- Gerfen, C. R., Engber, T. M., Mahan, L. C., Susel, Z., Chase, T. N., Monsma, F. J., Jr., & Sibley, J. R. (1990, December 7). D1 and D2 dopamine receptor-regulated gene expression of striatonigral and striatopallidal neurons. *Science*, *250*, 1429–1432.
- Gerfen, C. R., Staines, W. A., Arbuthnott, G. W., & Fibiger, H. C. (1982). Crossed connections of the substantia nigra in the rat. *Journal of Comparative Neurology*, *207*, 283–303.
- Gerstner, W., & Kistler, W. M. (2002). Mathematical formulations of Hebbian learning. *Biological Cybernetics*, *87*, 404–415.
- Gillies, A., & Arbuthnott, G. W. (2000). Computational models of the basal ganglia. *Movement Disorders*, *15*, 762–770.
- Glass, M., Dragunow, M., & Faull, R. L. (2000). The pattern of neurodegeneration in Huntington's disease: A comparative study of cannabinoid, dopamine, adenosine and GABA(A) receptor alterations in the human basal ganglia in Huntington's disease. *Neuroscience*, *97*, 505–519.
- Gluck, M. A., & Bower, G. H. (1988). From conditioning to category learning: An adaptive network model. *Journal of Experimental Psychology: General*, *117*, 227–247.
- 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*, 408–418.
- Goel, V., Buchel, C., Frith, C., & Dolan, R. J. (2000). Dissociation of mechanisms underlying syllogistic reasoning. *NeuroImage*, *12*, 504–514.
- Granger, R. (2006). Engines of the brain: The computational instruction set of human cognition. *AI Magazine*, *27*, 15–31.
- Graveland, G. A., Williams, R. S., & DiFiglia, M. (1985). A Golgi study of the human neostriatum: Neurons and afferent fibers. *Journal of Comparative Neurology*, *234*, 317–333.
- Graybiel, A. M. (2000). The basal ganglia. *Current Biology*, *10*, R509–R511.
- Graybiel, A. M. (2005). The basal ganglia: Learning new tricks and loving it. *Current Opinion in Neurobiology*, *15*, 638–644.
- Graybiel, A. M., & Ragsdale, C. W., Jr. (1978). Histochemically distinct compartments in the striatum of human, monkeys, and cat demonstrated by acetylthiocholinesterase staining. *Proceedings of the National Academy of Sciences, USA*, *75*, 5723–5726.
- Gurney, K., Prescott, T. J., & Redgrave, P. (2001). A computational model of action selection in the basal ganglia: I. A new functional anatomy. *Biological Cybernetics*, *84*, 401–410.
- Haber, S. N. (2003). The primate basal ganglia: Parallel and integrative networks. *Journal of Chemical Neuroanatomy*, *26*, 317–330.
- Haber, S. N., Fudge, J. L., & McFarland, N. R. (2000). Striatonigrostriatal pathways in primates form an ascending spiral from the shell to the dorsolateral striatum. *Journal of Neuroscience*, *20*, 2369–2382.
- Hajos, M., & Greenfield, S. A. (1994). Synaptic connections between pars compacta and pars reticulata neurones: Electrophysiological evidence for functional modules within the substantia nigra. *Brain Research*, *660*, 216–224.
- Hattori, T., Fibiger, H. C., & McGeer, P. L. (1975). Demonstration of a pallido-nigral projection innervating dopaminergic neurons. *Journal of Comparative Neurology*, *162*, 487–504.
- Hawkins, J., & Blakeslee, S. (2004). *On intelligence*. New York, NY: Times Books.
- Hazeltine, E., Teague, D., & Ivry, R. B. (2002). Simultaneous dual-task performance reveals parallel response selection after practice. *Journal of Experimental Psychology: Human Perception and Performance*, *28*, 527–545.
- Hikosaka, O., Rand, M. K., Miyachi, S., & Miyashita, K. (1995). Learning of sequential movements in the monkey: Process of learning and retention of memory. *Journal of Neurophysiology*, *74*, 1652–1661.
- Hill, N. M., & Schneider, W. (2006). Brain changes in the development of expertise: Neurological evidence on skill-based adaptations. In K. A. Ericsson, N. Charness, P. Feltovich, & R. Hoffman (Eds.), *Cambridge handbook of expertise and expert performance* (pp. 653–682). New York, NY: Cambridge University Press.
- Hoover, J. E., & Strick, P. L. (1993, February 5). Multiple output channels in the basal ganglia. *Science*, *259*, 819–821.
- Houk, J. C. (2005). Agents of the mind. *Biological Cybernetics*, *92*, 427–437.
- Houk, J. C., Adams, J. L., & Barto, A. G. (1995). A model of how the basal ganglia generate and use signals that predict reinforcement. In J. C. Houk, J. L. Davis, & D. G. Beiser (Eds.), *Models of information processing in the basal ganglia* (pp. 249–270). Cambridge, MA: MIT Press.
- Hubert, V., Beaunieux, H., Chetelat, G., Platel, H., Landeau, B., Danion, J. M., . . . Eustache, F. (2007). The dynamic network subserving the three phases of cognitive procedural learning. *Human Brain Mapping*, *28*, 1415–1429.
- Humphries, M. D., Stewart, R. D., & Gurney, K. N. (2006). A physiologically plausible model of action selection and oscillatory activity in the basal ganglia. *Journal of Neuroscience*, *26*, 12921–12942.
- Jackson, G. M., Jackson, S. R., Harrison, J., Henderson, L., & Kennard, C. (1995). Serial reaction time learning and Parkinson's disease: Evidence for a procedural learning deficit. *Neuropsychologia*, *33*, 577–593.
- Jaeger, D., Kita, H., & Wilson, C. J. (1994). Surround inhibition among projection neurons is weak or nonexistent in the rat neostriatum. *Journal of Neurophysiology*, *72*, 2555–2558.
- Jankovic, J. (2008). Parkinson's disease: Clinical features and diagnosis. *Journal of Neurology, Neurosurgery & Psychiatry*, *79*, 368–376.
- Joel, D., Niv, Y., & Ruppel, E. (2002). Actor-critic models of the basal ganglia: New anatomical and computational perspectives. *Neural Networks*, *15*, 535–547.
- Joel, D., & Weiner, I. (1994). The organization of the basal ganglia-thalamocortical circuits: Open interconnected rather than closed segregated. *Neuroscience*, *63*, 363–379.
- Joel, D., & Weiner, I. (2000). The connections of the dopaminergic system with the striatum in rats and primates: An analysis with respect to the functional and compartmental organization of the striatum. *Neuroscience*, *96*, 451–474.
- Jog, M. S., Kubota, Y., Connolly, C. I., Hillegeart, V., & Graybiel, A. M. (1999, November 26). Building neural representations of habits. *Science*, *286*, 1745–1749.
- Juraska, J. M., Wilson, C. J., & Groves, P. M. (1970). The substantia nigra of the rat: A Golgi study. *Journal of Comparative Neurology*, *172*, 585–600.
- Just, M. A., & Carpenter, P. A. (1992). A capacity theory of comprehension: Individual differences in working memory. *Psychological Review*, *99*, 122–149.
- Just, M. A., & Varma, S. (2007). The organization of thinking: What functional brain imaging reveals about the neuroarchitecture of complex cognition. *Cognitive, Affective, & Behavioral Neuroscience*, *7*, 153–191.
- Kawaguchi, Y. (1997). Neostriatal cell subtypes and their functional roles. *Neurosciences Research*, *27*, 1–8.
- Kelly, R. M., & Strick, P. L. (2004). Macro-architecture of basal ganglia loops with the cerebral cortex: Use of rabies virus to reveal multisynaptic circuits. *Progress in Brain Research*, *143*, 449–459.
- Kemp, J. M., & Powell, T. P. (1970). The cortico-striate projection in the monkey. *Brain*, *93*, 525–546.
- 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*, 462–470.

- Kita, H., & Kitai, S. T. (1988). Glutamate decarboxylase immunoreactive neurons in rat neostriatum: Their morphological types and populations. *Brain Research*, *447*, 346–352.
- Knowlton, B. J., Mangels, J. A., & Squire, L. R. (1996, September 6). A neostriatal habit learning system in humans. *Science*, *273*, 1399–1402.
- Knowlton, B. J., Squire, L. R., & Gluck, M. A. (1994). Probabilistic classification learning in amnesia. *Learning & Memory*, *1*, 106–120.
- Koos, T., & Tepper, J. M. (1999). Inhibitory control of neostriatal projection neurons by GABAergic interneurons. *Nature Neuroscience*, *2*, 467–472.
- Koos, T., & Tepper, J. M. (2002). Dual cholinergic control of fast-spiking interneurons in the neostriatum. *Journal of Neuroscience*, *22*, 529–535.
- Koos, T., Tepper, J. M., & Wilson, C. J. (2004). Comparison of IPSCs evoked by spiny and fast-spiking neurons in the neostriatum. *Journal of Neuroscience*, *24*, 7916–7922.
- Kötter, R., & Wickersham, J. (1995). Interactions of glutamate and dopamine in a computational model of the striatum. *Journal of Computational Neuroscience*, *2–3*, 195–214.
- Laird, J. E., Newell, A., & Rosenbloom, P. S. (1987). SOAR: An architecture for general intelligence. *Artificial Intelligence*, *33*, 1–64.
- Laird, J. E., Rosenbloom, P. S., & Newell, A. (1986). Chunking in SOAR: The anatomy of a general learning mechanism. *Machine Learning*, *1*, 11–46.
- Lawrence, A. D., Sahakian, B. J., & Robbins, T. W. (1998). Cognitive functions and corticostriatal circuits: Insights from Huntington's disease. *Trends in Cognitive Sciences*, *2*, 379–388.
- Lebiere, C., & Anderson, J. R. (1993). A connectionist implementation of the ACT-R production system. In *Proceedings of the Fifteenth Annual Conference of the Cognitive Science Society* (pp. 635–640). Mahwah, NJ: Erlbaum.
- Levin, B. E., Llabre, M. M., & Weiner, W. J. (1989). Cognitive impairments associated with early Parkinson's disease. *Neurology*, *39*, 557–561.
- Lewis, S. J., Dove, A., Robbins, T. W., Barker, R. A., & Owen, A. M. (2004). Striatal contributions to working memory: A functional magnetic resonance imaging study in humans. *European Journal of Neuroscience*, *19*, 755–760.
- Mallet, N., Le Moine, C., Charpier, S., & Gonon, F. (2005). Feedforward inhibition of projection neurons by fast-spiking GABA interneurons in the rat striatum in vivo. *Journal of Neuroscience*, *25*, 3857–3869.
- McFarland, N. R., & Haber, S. N. (2002). Thalamic relay nuclei of the basal ganglia form both reciprocal and nonreciprocal cortical connections, linking multiple frontal cortical areas. *Journal of Neuroscience*, *22*, 8117–8132.
- McNab, F., & Klingberg, T. (2008). Prefrontal cortex and basal ganglia control access to working memory. *Nature Neuroscience*, *11*, 103–107.
- Meyer, D. E., & Kieras, D. E. (1997a). A computational theory of executive cognitive processes and multiple-task performance: 1. Basic mechanisms. *Psychological Review*, *104*, 3–65.
- Meyer, D. E., & Kieras, D. E. (1997b). A computational theory of executive cognitive processes and multiple-task performance: 2. Accounts of psychological refractory-period phenomena. *Psychological Review*, *104*, 749–791.
- Middleton, F. A., & Strick, P. L. (1994, October 21). Anatomical evidence for cerebellar and basal ganglia involvement in higher cognitive function. *Science*, *266*, 458–461.
- Middleton, F. A., & Strick, P. L. (1996). The temporal lobe is a target of output from the basal ganglia. *Proceedings of the National Academy of Sciences, USA*, *93*, 8683–8687.
- Middleton, F. A., & Strick, P. L. (2000). Basal ganglia output and cognition: Evidence from anatomical, behavioral, and clinical studies. *Brain and Cognition*, *42*, 183–200.
- Miller, J. D., Sanghera, M. K., & German, D. C. (1981). Mesencephalic dopaminergic unit activity in the behaviorally conditioned rat. *Life Sciences*, *29*, 1255–1263.
- Mink, J. W., & Thach, W. T. (1993). Basal ganglia intrinsic circuits and their role in behavior. *Current Opinion in Neurobiology*, *3*, 950–957.
- Mishkin, M., & Petri, H. L. (1984). Memories and habits: Some implications for the analysis of learning and retention. In L. R. Squire & N. Butters (Eds.), *Neuropsychology of memory* (pp. 287–296). New York, NY: Guilford Press.
- Miyachi, S., Hikosaka, O., Miyashita, K., Karadi, Z., & Rand, M. K. (1997). Differential roles of monkey striatum in learning of sequential hand movement. *Experimental Brain Research*, *115*, 1–5.
- Moazami-Goudarzi, M., Sarnthein, J., Michels, L., Moukhtieva, R., & Jeanmonod, D. (2008). Enhanced frontal low and high frequency power and synchronization in the resting EEG of Parkinsonian patients. *NeuroImage*, *41*, 985–997.
- Monchi, O., Taylor, J. G., & Dagher, A. (2000). A neural model of working memory processes in normal subjects, Parkinson's disease and schizophrenia for fMRI design and predictions. *Neural Networks*, *13*, 953–973.
- Montague, P. R., Hyman, S. E., & Cohen, J. D. (2004, October 14). Computational roles for dopamine in behavioural control. *Nature*, *431*, 760–767.
- Montague, P. R., King-Casas, B., & Cohen, J. D. (2006). Imaging valuation models in human choice. *Annual Review of Neuroscience*, *29*, 417–448.
- Morris, G., Arkadir, D., Nevet, A., Vaadia, E., & Bergman, H. (2004). Coincident but distinct messages of midbrain dopamine and striatal tonically active neurons. *Neuron*, *43*, 133–143.
- Muslimovic, D., Post, B., Speelman, J. D., & Schmand, B. (2005). Cognitive profile of patients with newly diagnosed Parkinson disease. *Neurology*, *65*, 1239–1245.
- Nakanishi, H., Kita, H., & Kitai, S. T. (1987). Intracellular study of rat substantia nigra pars reticulata neurons in an in vitro slice preparation: Electrical membrane properties and response characteristics to subthalamic stimulation. *Brain Research*, *437*, 45–55.
- Nambu, A., Tokuno, H., Hamada, I., Kita, H., Imanishi, M., Akazawa, T., . . . Hasegawa, N. (2000). Excitatory cortical inputs to pallidal neurons via the subthalamic nucleus in the monkey. *Journal of Neurophysiology*, *84*, 289–300.
- Nambu, A., Tokuno, H., & Takada, M. (2002). Functional significance of the cortico-subthalamo-pallidal "hyperdirect" pathway. *Neuroscience Research*, *43*, 111–117.
- Newell, A. N. (1973). Production systems: Models of control structure. In W. G. Chase (Ed.), *Visual information processing* (pp. 526–547). New York, NY: Academic Press.
- Newell, A. N. (1990). *Unified theories of cognition*. Cambridge, MA: Harvard University Press.
- Nicola, S. M., Surmeier, J., & Malenka, R. C. (2000). Dopaminergic modulation of neuronal excitability in the striatum and nucleus accumbens. *Annual Review of Neuroscience*, *23*, 185–215.
- Nissen, M. J., & Bullemer, P. (1987). Attention requirements of learning: Evidence from performance measures. *Cognitive Psychology*, *19*, 1–32.
- O'Doherty, J. P. (2004). Reward representations and reward-related learning in the human brain: Insights from neuroimaging. *Current Opinion in Neurobiology*, *14*, 769–776.
- O'Doherty, J. P., Deichmann, R., Critchley, H. D., & Dolan, R. J. (2002). Neural responses during anticipation of a primary taste reward. *Neuron*, *33*, 815–826.
- O'Reilly, R. C. (1996). Biologically plausible error-driven learning using local activation differences: The generalized recirculation algorithm. *Neural Computation*, *8*, 895–938.
- O'Reilly, R. C., Braver, T. S., & Cohen, J. D. (1999). A biologically-based computational model of working memory. In A. Miyake & P. Shah (Eds.), *Models of working memory: Mechanisms of active maintenance and executive control* (pp. 375–411). New York, NY: Cambridge University Press.

- 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 Computation*, *18*, 283–328.
- O'Reilly, R. C., & Munakata, Y. (2000). *Computational explorations in cognitive neuroscience*. Cambridge, MA: MIT Press.
- Owen, A. M. (2004). Cognitive dysfunction in Parkinson's disease: The role of frontostriatal circuitry. *The Neuroscientist*, *10*, 525–537.
- Owen, A. M., Doyon, J., Petrides, M., & Evans, A. C. (1996). Planning and spatial working memory: A positron emission tomography study in humans. *European Journal of Neuroscience*, *8*, 353–364.
- Owen, A. M., Iddon, J. L., Hodges, J. R., Summers, B. A., & Robbins, T. W. (1997). Spatial and nonspatial working memory at different stages of Parkinson's disease. *Neuropsychologia*, *35*, 519–532.
- 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*, 1465–1472.
- Packard, M. G., & Knowlton, B. J. (2002). Learning and memory functions of the basal ganglia. *Annual Review of Neuroscience*, *25*, 563–593.
- Packard, M. G., & McGaugh, J. L. (1992). Double dissociation of fornix and caudate nucleus lesions on acquisition of two water maze tasks: Further evidence for multiple memory systems. *Behavioral Neuroscience*, *106*, 439–446.
- 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 and Memory*, *65*, 65–72.
- Pakhotin, P., & Bracci, E. (2007). Cholinergic interneurons control the excitatory input to the striatum. *Journal of Neuroscience*, *27*, 391–400.
- Parent, A., & Hazrati, L. N. (1995). Functional anatomy of the basal ganglia: II. The place of subthalamic nucleus and external pallidum in basal ganglia circuitry. *Brain Research Reviews*, *20*, 128–154.
- Parthasarathy, H. B., Schall, J. D., & Graybiel, A. M. (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.
- Pasupathy, A., & Miller, E. K. (2005, February 24). Different time courses of learning-related activity in the prefrontal cortex and striatum. *Nature*, *433*, 873–876.
- Paulsen, J. S., Butters, N., Salmon, D. P., Heindel, W. C., & Swenson, M. R. (1993). Prism adaptation in Alzheimer's and Huntington's disease. *Neuropsychology*, *7*, 73–81.
- Peigneux, P., Maquet, P., Meulemans, T., Destrebecqz, A., Laureys, S., Degueldre, C., ... Cleeremans, A. (2000). Striatum forever, despite sequence learning variability: A random effect analysis of PET data. *Human Brain Mapping*, *10*, 179–194.
- Penney, J. B., Jr., & Young, A. B. (1986). Striatal inhomogeneities and basal ganglia function. *Movement Disorders*, *1*, 3–15.
- Plate, T. A. (1995). Holographic reduced representations. *IEEE Transactions on Neural Networks*, *6*, 623–641.
- Plenz, D. (2003). When inhibition goes incognito: Feedback interaction between spiny projection neurons in striatal function. *Trends in Neurosciences*, *26*, 436–443.
- Plenz, D., & Kital, S. T. (1999, August 12). A basal ganglia pacemaker formed by the subthalamic nucleus and external globus pallidus. *Nature*, *400*, 677–682.
- Prat, C. S., & Just, M. A. (2010). *Exploring the cortical dynamics underpinning individual differences in sentence comprehension*. Manuscript submitted for publication.
- Prat, C. S., Keller, T. A., & Just, M. A. (2007). Individual differences in sentence comprehension: A functional magnetic resonance imaging investigation of syntactic and lexical processing demands. *Journal of Cognitive Neuroscience*, *19*, 1950–1963.
- Preuschhoff, K., Bossaerts, P., & Quartz, S. R. (2006). Neural differentiation of expected reward and risk in human subcortical structures. *Neuron*, *51*, 381–390.
- Qin, Y., Sohn, M.-H., Anderson, J. R., Stenger, V. A., Fissell, K., Goode, A., & Carter, C. S. (2003). Predicting the practice effects on the blood oxygenation level-dependent (BOLD) function of fMRI in a symbolic manipulation task. *Proceedings of the National Academy of Sciences, USA*, *100*, 4951–4956.
- Raichle, M. E., Fiez, J. A., Videen, T. O., MacLeod, A. M., Pardo, J. V., Fox, P. T., & Petersen, S. E. (1994). Practice-related changes in human brain functional anatomy during nonmotor learning. *Cerebral Cortex*, *4*, 8–26.
- Rauch, S. L., Whalen, P. J., Savage, C. R., Curran, T., Kendrick, A., Brown, H. D., ... Rosen, B. R. (1997). Striatal recruitment during an implicit sequence learning task as measured by functional magnetic resonance imaging. *Human Brain Mapping*, *5*, 124–132.
- Ravizza, S. M., & Ivry, R. B. (2001). Comparison of the basal ganglia and cerebellum in shifting attention. *Journal of Cognitive Neuroscience*, *13*, 285–297.
- Reder, L. M., Park, H., & Kieffaber, P. D. (2009). Memory systems do not divide on consciousness: Reinterpreting memory in terms of activation and binding. *Psychological Bulletin*, *135*, 23–49.
- Redgrave, P., & Gurney, K. (2006). The short-latency dopamine signal: A role in discovering novel actions? *Nature Reviews Neuroscience*, *7*, 967–975.
- Redgrave, P., Prescott, T. J., & Gurney, K. (1999). The basal ganglia: A vertebrate solution to the selection problem? *Neuroscience*, *89*, 1009–1023.
- Reiner, A., Albin, R. L., Anderson, K. D., D'Amato, C. J., Penney, J. B., & Young, A. B. (1988). Differential loss of striatal projection neurons in Huntington disease. *Proceedings of the National Academy of Sciences, USA*, *85*, 5733–5737.
- Reynolds, J. N., Hyland, B. I., & Wickens, J. R. (2004). Modulation of an afterhyperpolarization by the substantia nigra induces pauses in the tonic firing of striatal cholinergic interneurons. *Journal of Neuroscience*, *24*, 9870–9877.
- Reynolds, J. N., & Wickens, J. R. (2004). The corticostriatal input to giant aspiny interneurons in the rat: A candidate pathway for synchronising the response to reward-related cues. *Brain Research*, *1011*, 115–128.
- Rolls, E. T., & Treves, A. (1998). *Neural networks and brain function*. New York, NY: Oxford University Press.
- Rouzaire-Dubois, B., & Scarnati, E. (1987). Pharmacological study of the cortical-induced excitation of subthalamic nucleus neurons in the rat: Evidence for amino acids as putative neurotransmitters. *Neuroscience*, *21*, 429–440.
- Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986, October 9). Learning representations by back-propagating errors. *Nature*, *323*, 533–536.
- Rymar, V. V., Sasseville, R., Luk, K. C., & Sadikot, A. F. (2004). Neurogenesis and stereological morphometry of calretinin-immunoreactive GABAergic interneurons of the neostriatum. *Journal of Comparative Neurology*, *469*, 325–339.
- Rypma, B., Prabhakaran, V., Desmond, J. E., Glover, G. H., & Gabrieli, J. D. E. (1999). Load-dependent roles of frontal brain regions in the maintenance of working memory. *NeuroImage*, *9*, 216–226.
- Schultz, W. (1998). Predictive reward signal of dopamine neurons. *Journal of Neurophysiology*, *80*, 1–27.
- Schultz, W. (2002). Getting formal with dopamine and reward. *Neuron*, *36*, 241–263.
- 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., Dayan, P., & Montague, P. R. (1997, March 14). A neural substrate of prediction and reward. *Science*, *275*, 1593–1599.
- Schumacher, E. H., Seymour, T. L., Glass, J. M., Fencsik, D. E., Lauber, E. J., Kieras, D. E., & Meyer, D. E. (2001). Virtually perfect time

- sharing in dual-task performance: Uncorking the central cognitive bottleneck. *Psychological Science*, 12, 101–108.
- Seger, C. A., & Cincotta, C. M. (2005). The roles of the caudate nucleus in human classification learning. *Journal of Neuroscience*, 25, 2941–2951.
- Shallice, T. (1982). Specific impairments of planning. *Philosophical Transactions of the Royal Society of London*, 298, 199–209.
- Shanks, D. R., & St. John, M. F. (1994). Characteristics of dissociable human learning systems. *Behavioral and Brain Sciences*, 17, 367–447.
- Shiffrin, R. M., & Schneider, W. (1977). Controlled and automatic human information processing: II. Perceptual learning, automatic attending and a general theory. *Psychological Review*, 84, 127–190.
- Shink, E., Bevan, M. D., Bolam, J. P., & Smith, Y. (1996). The subthalamic nucleus and the external pallidum: Two tightly interconnected structures that control the output of the basal ganglia in the monkey. *Neuroscience*, 73, 335–357.
- Smith, A. D., & Bolam, J. P. (1990). The neural network of the basal ganglia as revealed by the study of synaptic connections of identified neurons. *Trends in Neurosciences*, 13, 259–265.
- Smith, Y., Raju, D. V., Pare, J. F., & Sidibe, M. (2004). The thalamostriatal system: A highly specific network of the basal ganglia circuitry. *Trends in Neurosciences*, 27, 520–527.
- Smolensky, P. (1990). Tensor product variable binding and the representation of symbolic structures in connectionist systems. *Artificial Intelligence*, 46, 159–216.
- Sohn, M.-H., Goode, A., Stenger, V. A., Carter, C. S., & Anderson, J. R. (2003). Competition and representation during memory retrieval: Roles of the prefrontal cortex and the posterior parietal cortex. *Proceedings of the National Academy of Sciences, USA*, 100, 7412–7417.
- Somogyi, P., Bolam, J. P., & Smith, A. D. (1981). Monosynaptic cortical input and local axon collaterals of identified striatonigral neurons: A light and electron microscopic study using the Golgi-peroxidase transport-degeneration procedure. *Journal of Comparative Neurology*, 195, 567–584.
- Sprengelmeyer, R., Canavan, A. G., Lange, H. W., & Homberg, V. (1995). Associative learning in degenerative neostriatal disorders: Contrasts in explicit and implicit remembering between Parkinson's and Huntington's diseases. *Movement Disorders*, 10, 51–65.
- Squire, L. R. (1992). Declarative and nondeclarative memory: Multiple brain systems supporting learning and memory. *Journal of Cognitive Neuroscience*, 4, 232–243.
- Squire, L. R., & Zola, S. M. (1996). Structure and function of declarative and nondeclarative memory systems. *Proceedings of the National Academy of Sciences, USA*, 93, 13515–13522.
- Stewart, T. C., & Eliasmith, C. (2008). Building production systems with realistic spiking neurons. In V. Sloutsky, B. Love, & K. McRae (Eds.), *Proceedings of the 30th annual meeting of the Cognitive Science Society*. Austin, TX: Cognitive Science Society.
- Stocco, A., & Anderson, J. R. (2008). Endogenous control and task representation: An fMRI study of algebraic problem solving. *Journal of Cognitive Neuroscience*, 20, 1300–1314.
- Stoffers, D., Bosboom, J. L., Deijen, J. B., Wolters, E., Stam, C. J., & Berendse, H. W. (2008). Increased cortico-cortical functional connectivity in early-stage Parkinson's disease: An MEG study. *NeuroImage*, 41, 212–222.
- Suri, R. E.,argas, J., & Arbib, M. A. (2001). Modeling functions of striatal dopamine modulation in learning and planning. *Neuroscience*, 103, 65–85.
- Suri, R. E., & Schultz, W. (1998). Learning of sequential movements by neural network model with dopamine-like reinforcement signal. *Experimental Brain Research*, 121, 350–354.
- Suri, R. E., & Schultz, W. (1999). A neural network model with dopamine-like reinforcement signal that learns a spatial delayed response task. *Neuroscience*, 91, 871–890.
- Sutton, R. S. (1988). Learning to predict by the methods of temporal differences. *Machine Learning*, 3, 9–44.
- Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction*. Cambridge, MA: MIT Press.
- Taatgen, N. A. (2005). Modeling parallelization and flexibility improvements in skill acquisition: From dual tasks to complex dynamic skills. *Cognitive Science*, 29, 421–455.
- Taatgen, N. A., & Lee, F. J. (2003). Production compilation: A simple mechanism to model complex skill acquisition. *Human Factors*, 45, 61–76.
- Teicher, M. H., Anderson, C. M., Polcari, A., Glod, C. A., Maas, L. C., & Renshaw, P. F. (2000). Functional deficits in basal ganglia of children with attention-deficit/hyperactivity disorder shown with functional magnetic resonance imaging relaxometry. *Nature Medicine*, 6, 470–473.
- Teichmann, M., Dupoux, E., Kouider, S., & Bachoud-Levi, A. C. (2006). The role of the striatum in processing language rules: Evidence from word perception in Huntington's disease. *Journal of Cognitive Neuroscience*, 18, 1555–1569.
- Teng, E., Stefanacci, L., Squire, L. R., & Zola, S. M. (2000). Contrasting effects on discrimination learning after hippocampal lesions and conjoint hippocampal–caudate lesions in monkeys. *Journal of Neuroscience*, 20, 3853–3863.
- Tepper, J. M., & Bolam, J. P. (2004). Functional diversity and specificity of neostriatal interneurons. *Current Opinion in Neurobiology*, 14, 685–692.
- Tepper, J. M., Martin, L. P., & Anderson, D. R. (1995). GABA_A receptor-mediated inhibition of rat substantia nigra dopaminergic neurons by pars reticulata projection neurons. *Journal of Neuroscience*, 15, 3092–3103.
- Terman, D., Rubin, J. E., Yew, A. C., & Wilson, C. J. (2002). Activity patterns in a model for the subthalamopallidal network of the basal ganglia. *Journal of Neuroscience*, 22, 2963–2976.
- Thompson-Schill, S. L., D'Esposito, M., & Kan, I. P. (1999). Effects of repetition and competition on activity in left prefrontal cortex during word generation. *Neuron*, 23, 513–522.
- Tobler, P. N., Fiorillo, C. D., & Schultz, W. (2005, March 11). Adaptive coding of reward value by dopamine neurons. *Science*, 307, 1642–1645.
- Tom, S. M., Fox, C. R., Trepel, C., & Poldrack, R. A. (2007, January 26). The neural basis of loss aversion in decision-making under risk. *Science*, 315, 515–518.
- Touretzky, D., & Hinton, G. H. (1988). A distributed connectionist production system. *Cognitive Science*, 12, 423–466.
- Tunstall, M. J., Oorschot, D. E., Kean, A., & Wickens, J. R. (2002). Inhibitory interactions between spiny projection neurons in the rat striatum. *Journal of Neurophysiology*, 88, 1263–1269.
- Ullman, M. T., Corkin, S., Coppola, M., Hickok, G., Growdon, J. H., Koroshetz, W. J., & Pinker, S. (1997). A neural dissociation within language: Evidence that the mental dictionary is part of declarative memory, and that grammatical rules are processed by the procedural system. *Journal of Cognitive Neuroscience*, 9, 266–276.
- Walker, F. O. (2007). Huntington's disease. *Lancet*, 369, 218–228.
- Wickens, J. R., Alexander, M. E., & Miller, R. (1991). Two dynamic modes of striatal function under dopaminergic–cholinergic control: Simulation and analysis of a model. *Synapse*, 8, 1–12.
- Wickens, J. R., & Arbutnot, G. W. (1993). The corticostriatal system on computer simulation: An intermediate mechanism for sequencing of actions. *Progress in Brain Research*, 99, 325–339.
- Wickens, J. R., Begg, A. J., & Arbutnot, G. W. (1996). Dopamine reverses the depression of rat corticostriatal synapses which normally follows high-frequency stimulation of cortex in vitro. *Neuroscience*, 70, 1–5.
- Willingham, D. B., & Koroshetz, W. J. (1993). Evidence for dissociable motor skills in Huntington's disease patients. *Psychobiology*, 21, 173–182.
- Wilson, C. J. (1993). The generation of natural firing patterns in neostriatal neurons. *Progress in Brain Research*, 99, 277–297.
- Wilson, C. J., & Groves, P. M. (1980). Fine structure and synaptic con-

- nections of the common spiny neuron of the rat neostriatum: A study employing intracellular inject of horseradish peroxidase. *Journal of Comparative Neurology*, 194, 599–615.
- Wilson, C. J., & Groves, P. M. (1981). Spontaneous firing patterns of identified spiny neurons in the rat neostriatum. *Brain Research*, 220, 67–80.
- Xie, X., & Seung, S. (2003). Equivalence of backpropagation and contrastive Hebbian learning in a layered network. *Neural Computation*, 15, 441–454.
- Yelnik, J., Francois, C., Percheron, G., & Tande, D. (1991). Morphological taxonomy of the neurons of the primate striatum. *Journal of Comparative Neurology*, 313, 273–294.
- Zhang, Y., Bertolino, A., Fazio, L., Blasi, G., Rampino, A., Romano, R., . . . Sadé, W. (2007). Polymorphisms in human dopamine D2 receptor gene affect gene expression, splicing, and neuronal activity during working memory. *Proceedings of the National Academy of Sciences, USA*, 104, 20552–20557.
- Zheng, T., & Wilson, C. J. (2002). Corticostriatal combinatorics: The implications of corticostriatal axonal arborizations. *Journal of Neurophysiology*, 87, 1007–1017.

Appendix A

Model Specification

This appendix gives an overview of the model's different types of units, their connections, and functions.

Model Neurons

In the model, neurons were implemented as simple computational units that apply an activation function f over an input value η to yield an activation value, denoted by x . The input value η is simply the sum of all the activations coming from the projecting neurons, weighted by the corresponding synaptic strengths:

$$\eta = \sum_j w_j x_j,$$

where w_j is the value (or *synaptic weight*) of the synapse from neuron j , and x_j is the activation of neuron j . This is perhaps the simplest and most common representation for artificial neurons, and it is widely adopted in many biological models (see O'Reilly & Munakata, 2000; Rolls & Treves, 1998).

The activation value x is obtained from the net input η by applying the activation function f :

$$x = f(\eta - \theta),$$

where θ is the neuron's threshold, which can be thought of an initial resistance of every neuron to be excited. A negative threshold (so that the quantity $\eta - \theta$ is positive in absence of direct stimulation) can be used to model neurons with high baseline activities, or to compensate for the effects of convergent inhibitory projections. The activation value x is supposed to be the computational counterpart of a neuron's firing rate. Note that a neuron's dynamic is completely characterized by its activation function and threshold.

With the exception of striatal interneurons (discussed below), all the neurons in the model use the hyperbolic tangent as their activation function:

$$x = \tanh(\gamma[\eta - \theta]_+),$$

where γ is the gain parameter that determines the curves' steepness, and the $[x]_+$ notation indicates that negative values of x are treated as zeroes. This ensures that the output of the function is in the range $[0, 1]$. Together with the sigmoid function, the hyperbolic tangent is among the simplest formulae that fit the change of spiking rates following changes in membrane potential in biological neurons; the curve also closely mimics the variation of spike rates to a change in the membrane potentials in biological neurons (see O'Reilly & Munakata, 2000). Table A1 details the values of γ and θ for each type of neuron in the model. Figure A1 gives a visual rendition of the corresponding activation curves.

Special Activation Function for Striatal Interneurons

Striatal interneurons exhibit special dynamics. They are tonically active and exert inhibitory pressure on striatal projection neurons, unless cortical activation reduces their firing rates. This behavior is likely produced by the interaction between cholinergic and GABAergic interneurons (e.g., Tepper & Bolam, 2004; see Figure 2). To account for this behavior, the only type of striatal interneurons in our model were provided with a special activation function, consisting of a sigmoid function with positive exponent. This function is monotonically decreasing, so that increased cortical inputs decrease the activity of interneurons. The function and its parameters are reported in Table A1 and visually depicted in Figure A1.

Baseline Activation Values

Each neuron also has an associated quantity called *baseline activation* value, which is indicated as $\langle x \rangle$. The baseline can be interpreted as the neuron's tonic activity. The baseline value provides a simple means to measure how much of a certain neuron's activation is due to its current inputs. It plays an important role in the Hebbian learning rule that is used in the model:

$$\Delta w_{ij} \approx r(x_i - \langle x_i \rangle)(x_j - \langle x_j \rangle).$$

Table A1
Summary of the Equations and Parameters That Govern Each Type of Neuron's Response

Neuron type	Activation function	Gain parameter	Threshold
Striatal interneuron	$x = \frac{1}{1 + e^{\gamma(\eta - \theta)}}$ $x = \tanh[\gamma(\eta - \theta)_+]$	$\gamma = 8.0$	$\theta = 1/2$
Striatal projection neuron (SN/SP)		$\gamma = 2.0$	$\theta_p \approx \sum_i w_{p,i} x_i^{*a}$
Thalamic neuron		$\gamma = 2.0$	
SNr/GPi neurons		$\gamma = 2.0$	$\theta = -2$
STN neurons		$\gamma = 3.0$	$\theta = -1$
GPe neuron		$\gamma = 3.5$	$\theta = -1$
SNc interneuron		$\gamma = 2.0$	$\theta = -1$
SNc dopamine neuron	$\gamma = 2.0$	$\theta = -1$	

Note. SN = striatonigral; SP = striatopallidal; SNr = pars reticulata of the substantia nigra; GPi = internal part of the globus pallidus; STN = subthalamic nucleus; GPe = external part of the substantia nigra; SNc = pars compacta of the substantia nigra.

^aStriatal projection neurons' thresholds were dynamically recalculated to match interneuron inhibition when cortical representations were transferred.

In this rule, subtracting the baseline from the activation prevents two neurons from becoming strongly associated when their activity is due to their "usual," tonic condition. It also makes it possible for a synapse to lose strength, whenever activation in one neuron is coupled with a decrease of activation in the other. The baselines for each different neuron type were calculated as their activation values in absence of stimulation, that is, when $\eta = 0$.

"Up" and "Down" States in Striatal Projection Neurons

Projection neurons in the striatum have special dynamics. They cannot be excited while they are in the "down" state. Cortical activity puts them in an "up" state; when in "up" state, an increase in excitatory input or a decrease of inhibition triggers a response (Bolam et al., 2000; Wilson, 1993). A realistically complex model

of this behavior was beyond the scope of our research. A simple approximation, however, consists of using neurons with a *dynamic threshold*. A threshold value θ is said to be dynamic when it is allowed to change over time. Some learning rules that have found biological support, such as the BCM rule (Bienenstock, Cooper, & Munro, 1982) make use of dynamic thresholds. In the model, the dynamic threshold θ_p for a projection neuron p approximates the expected input from striatal interneurons when cortical patterns are being gated. That is, the threshold is adapted to match the amount of inhibition that a projection neuron receives from interneurons when the projection neuron is nonetheless firing:

$$\theta_p \approx \sum_i w_{p,i} x_i^*$$

where x_i^* indicates the average activation of interneuron i when cortical signals are allowed to pass. This value depends on the

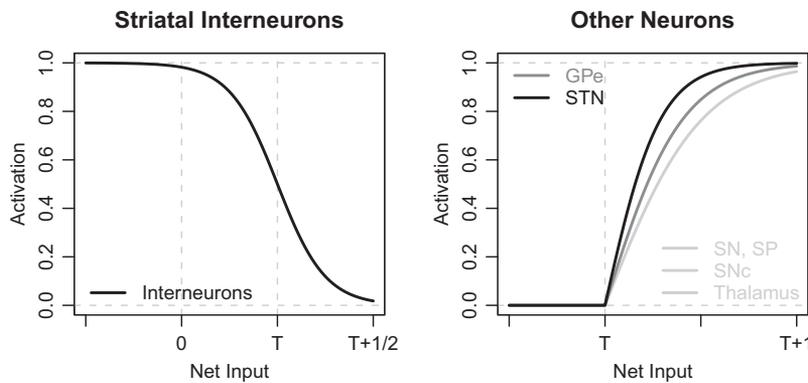


Figure A1. A visual rendition of the different activation functions used in the model. Left: The monotonically decreasing function used to simulate the firing rate patterns of striatal interneurons. Right: The monotonically increasing functions used to simulate the firing patterns of all the other neurons. GPe = external part of the globus pallidus; STN = subthalamic nucleus; SN = striatonigral; SP = striatopallidal; SNc = pars compacta of the substantia nigra.

(Appendices continue)

routing patterns encoded in the model, and was therefore calculated separately for each set of simulations. Note that the values of θ_p are dynamic because they depend on the strength of synapses $w_{p,i}$. Therefore, they are recalculated every time the synaptic weights are changed by Hebbian learning.

A model projection neuron's activation remains constant and equal to zero until the sum of all its inputs stays below θ_p . This corresponds to the "down" state. When the interneuron inhibition matches the threshold, the neuron reaches the "up" state, and any additional input, from either the cortex or interneurons, increases its activation.

More on Synapses

Inhibitory synapses were encoded as negative weights, and excitatory synapses were encoded as positive weights. Although the value of the synaptic change was left free to change according to Hebbian learning, no synapse could ever change *sign*. That is, negative synapses could not rise above zero, and positive synapses could not decrease below zero. This reflects the biological fact that inhibitory synapses cannot turn excitatory, and vice versa. The only exception to this rule consists of the synapses between cortical neurons and striatal interneurons. The reason for this exception is that striatal interneurons represent the net contribution of GABAergic and cholinergic interneurons.

Receptive Fields and Representation Compression

Nuclei in the basal ganglia have increasingly smaller sizes, which suggests a progressive *funneling* of information (Alexander et al., 1986). This is an important characteristic of the basal ganglia

physiology and needs to be addressed in a realistic model of the circuit. The easiest way to model this compression of information is to arrange the synaptic inputs so that a neuron from a smaller region receives inputs from many neurons that occupy the same position in a larger, input nucleus. Let us suppose that the projecting region has m neuron, and its target region contains n neurons (with $n < m$). If we indicate with j a neuron in the projecting region, and with i a neuron in the target region, then the synaptic weight $w_{i,j}$ is given by

$$w_{i,j} = G[i - j \times (n/m), \sigma],$$

where $G(x, \sigma)$ is a Gaussian (normal) function with mean 0 and standard deviation σ . In the expression, the term n/m is used to express the position of the neuron j within a range between 0 and n . This way, the relative positions of neurons i and j in the two regions can be compared. The term $i - j \times (n/m)$ can be read as the difference between the two relative positions. When this difference is zero, j is at the center of i 's receptive field.

Figure A2 illustrates the shape of such receptive field in the case of projections from the cortex ($m = 100$) to the striatum ($n = 10$); this m/n ratio is actually close to the real ratio of cortical projection neurons to striatal projection neurons, as estimated by Zheng & Wilson, 2002). In the model, similar functions are used to model connections between all nuclei, which usually differ in size. Notice that synaptic weights depend only on m and n and the free parameter σ . In the model, $\sigma = 1/2$ across all projections. The only exception was the striatal projections from interneurons to output neurons, where $\sigma = n/2$. This created an almost uniform inhibitory pressure.

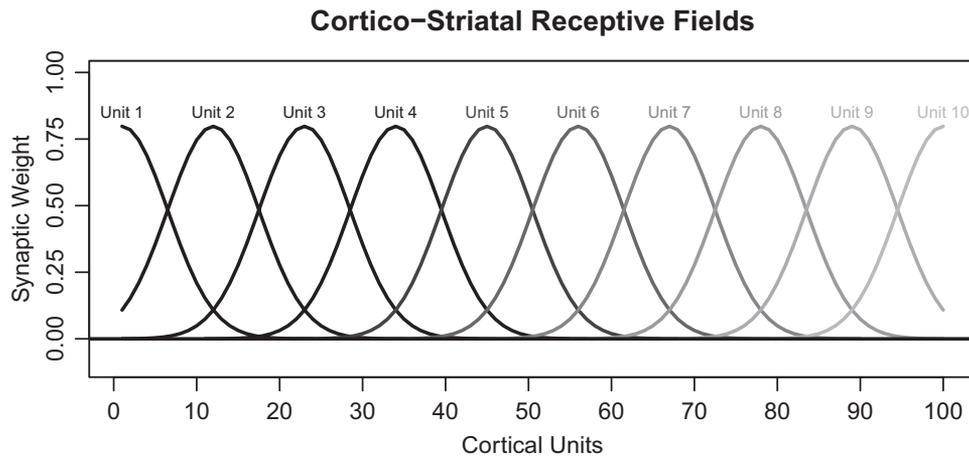


Figure A2. A visual rendition of the Gaussian striatal receptive fields used in the model. In this figure, 10 striatal units receive inputs from 100 cortical units. Their receptive fields are shown as bell curves of different shades of gray. They are shaped in such a way that each striatal unit is maximally sensitive to those cortical neurons that occupy a similar position in the cortical regions. This way, cortical topology is maintained within a striatal subdivision.

Appendix B

Learning Algorithm

This article focuses on habit learning, that is, practice-related changes occurring by repeated execution of routing operations. For this form of learning to occur, the model needs to have already learned a preliminary number of routing operations. This appendix describes the supervised learning procedure that was used to learn these initial operations. The preliminary learning of routing operations occurs in two sites: (a) the projections from cortical neurons to striatal interneurons, and (b) the inhibitory projections between striatal interneurons and striatal projection neurons. These projections occur within a simple three-layer network, which is illustrated in Figure B1.

Several algorithms exist to develop learning in a three-layered network. Perhaps the most famous is backpropagation, which consists of propagating backward the vector of differences between each output layer neuron’s target and actual value (Rumelhart, Hinton, & Williams, 1986). This algorithm, however, is implausible, as it depends on computations that are nonlocal and require signals to propagate back from the dendrite through the axon. The same computations can be achieved by a Hebbian-like rule known as contrastive Hebbian learning (CHL; Dayan & Abbott, 2001; Rolls & Treves, 1998). The CHL rule is defined as follows:

$$\Delta w_{i,j} = r(x_j t_i - x_j x_i) = r x_j (t_i - x_i), \quad (B1)$$

where x_i represents the actual activation of neuron i and t_i represents its desired (target) value. This rule is biologically admissible in that it requires only local computations between pairs of neurons. Variations of the CHL algorithm have also been used as the basis for reward-based learning in other models of the basal ganglia (e.g., Frank et al., 2001; O’Reilly & Frank, 2006). It can be

seen that the rule consists of two Hebbian steps: an anti-Hebbian update between the presynaptic activation x_j and postsynaptic activation x_i , and a Hebbian update between the presynaptic activation x_j and the postsynaptic target t_i .

Although Equation B1 can be applied only on two-layer networks, the CHL algorithm can be generalized to multilayer network, where it performs comparably to the backpropagation algorithm (O’Reilly, 1996; Xie & Seung, 2003). However, the generalization requires bidirectional connections between the consecutive layers of a network; certain variants (e.g., Xie & Seung, 2003) even require these connections to be symmetrical. These requirements do not hold in our three-layered cortico–striatal network.

To perform the initial training of the network, the CHL algorithm was broken in two steps and was performed separately for the cortico–striatal and the striato–striatal projections (see Figure B1). For the cortico–striatal projections, the exact procedure was the following:

1. An initial random state was generated for the striatal interneurons by setting a number of k different striatal interneurons to zero and letting all the others go to their baseline levels. This was designated as the desired state t for striatal interneurons.
2. The activation values of cortical neurons were then clamped to the cortical pattern the model should respond to.
3. The activation values x_i of striatal interneurons are calculated, and the anti-Hebbian term $-x_j x_i$ is calculated.
4. The activation values of interneurons are clamped to the desired state t_i .
5. The Hebbian term $x_j t_i$ is calculated, and each cortico–striatal synapse is then updated according to Equation B1.

For the striato–striatal projection, the procedure was the following:

1. The desired target state t of striatal projection neurons was generated by setting all their activation values to a negative number. If a neuron belongs to an ensemble corresponding to a desired source–destination binding, its value is set to a small positive number.
2. The activation values of striatal interneurons are clamped to their target state, as calculated in the previous procedure.
3. The activation values of striatal projection neuron are calculated, and the anti-Hebbian term $-x_j x_i$ is calculated.

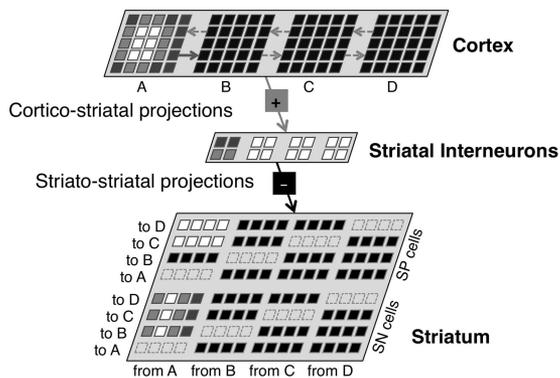


Figure B1. The three-layer network that was trained for the initial learning of routing operations. SP = striatopallidal; SN = striatonigral.

(Appendices continue)

4. The activation value of projection neurons are clamped to their desired state.
5. The Hebbian term $-x_j t_i$ is calculated, and each cortico-striatal synapse is then updated according to Equation B1.

This process was repeated until the network error E between the desired and actual striatal outputs was smaller than a criterion value: $E < .001$. The error was calculated as:

$$E = \sum_i (x_i - t_i)^2.$$

Contrastive Hebbian Learning and Reward

Although the CHL is plausible in that it requires only local computations, it still is a supervised learning procedure because it requires a detailed representation of the desired output for each neuron (the value t_i). In contrast, the simple error signal conveyed by dopamine neurons is thought to reflect a basic scalar value (e.g., Schultz, 1998, 2002). However, this rule can be thought of as approximation of unsupervised simple error-driven learning, which can be triggered by dopamine release. To see how this is possible, we need to rewrite Equation B1 in this form:

$$\Delta w_{ij} = r x_j t_i - r x_j x_i. \quad (\text{B2})$$

Equation B2 makes it clear that CHL is just the combination of two standard Hebbian updates: a Hebbian increment when the output of striatal interneurons equals t_i and an anti-Hebbian update when it equals any other possible value x_i .

As outlined in the article, Hebbian learning is actually a three-term interaction between these neurons and dopamine release, which we can simply write as an additional term d .

$$\Delta w_{ij} = d_{x=t} r x_j x_i + d_{x \neq t} r x_j x_i, \quad (\text{B3})$$

where $d_{x=t}$ is the dopamine release when the value is the target value, and $d_{x \neq t}$ is the dopamine when the interneuron activation is off the target. It is easy to see that this equation reduces to CHL when $d_{x=t} = -d_{x \neq t}$. But this requirement simply means that dopamine response needs to have the opposite sign when the interneuron response is off target. That is, dopamine release needs to be reduced when the striatum makes an error. This reduction is entirely compatible with the dopamine response to reward prediction errors (see Schultz, 1998, 2002, for an extensive review).

Contrastive Hebbian Learning and SP Neurons

The CHL algorithm still requires an exact representation of the correct destination in striatopallidal projection (SP) neurons. To understand how this could possibly happen without supervised training, we can express Equation B3 in probabilistic form:

$$\Delta w_{ij} \approx P(x = t) x_j x_i - P(x \neq t) x_j x_i.$$

Assuming that there is sufficient initial noise in the neuron response to explore among different possible patterns, the weight update is eventually dominated by the probability of an ensemble to produce a reward or a punishment. In particular, SP neurons that encode for the desired destination are more likely to incur in a negative learning signal than other destinations (which, as a simplification, we can imagine as having no effect). In the long run, therefore, their synaptic weights with the striatal interneurons will fall to zero, and they will simply remain in the “down” state (i.e., inactive) when the corresponding striatonigral projection neuron (SN) neurons are activated.

Received September 5, 2008

Revision received December 28, 2009

Accepted January 1, 2010 ■