Available online at www.sciencedirect.com
ScienceDirect
Neural Networks 19 (2006) 1120–1136
Neural Networks
www.elsevier.com/locate/neunet

2006 Special Issue

Hold your horses: A dynamic computational role for the subthalamic nucleus in decision making

Michael J. Frank*,[1]

Department of Psychology, Program in Neuroscience, University of Arizona, 1503 E University Blvd, Tucson, AZ 85721, United States

Received 31 October 2005; accepted 30 March 2006

Abstract

The basal ganglia (BG) coordinate decision making processes by facilitating adaptive frontal motor commands while suppressing others. In previous work, neural network simulations accounted for response selection deficits associated with BG dopamine depletion in Parkinson's disease. Novel predictions from this model have been subsequently confirmed in Parkinson patients and in healthy participants under pharmacological challenge. Nevertheless, one clear limitation of that model is in its omission of the subthalamic nucleus (STN), a key BG structure that participates in both motor and cognitive processes. The present model incorporates the STN and shows that by modulating *when* a response is executed, the STN reduces premature responding and therefore has substantial effects on *which* response is ultimately selected, particularly when there are multiple competing responses. Increased cortical response conflict leads to dynamic adjustments in response thresholds via cortico-subthalamic-pallidal pathways. The model accurately captures the dynamics of activity in various BG areas during response selection. Simulated dopamine depletion results in emergent oscillatory activity in BG structures, which has been linked with Parkinson's tremor. Finally, the model accounts for the beneficial effects of STN lesions on these oscillations, but suggests that this benefit may come at the expense of impaired decision making.

Keywords: Basal ganglia; Decision making; Subthalamic nucleus; Neural network model; Parkinson's disease; Reinforcement learning

1. Introduction

Deciphering the mechanisms by which the brain supports response selection, a central process in decision making, is an important challenge for both the artificial intelligence and cognitive neuroscience communities. Based on a wealth of data, the basal ganglia (BG) are thought to play a principal role in these processes. In the context of motor control, various authors have suggested that the role of the BG is to selectively facilitate the execution of a single adaptive motor command, while suppressing all others (Basso & Wurtz, 2002; Brown, Bullock, & Grossberg, 2004; Frank, 2005a; Gurney, Prescott, & Redgrave, 2001; Hikosaka, 1994; Jiang, Stein, & McHaffie, 2003; Mink, 1996; Redgrave, Prescott, & Gurney, 1999). Interestingly, circuits linking the BG with more cognitive areas of frontal cortex (e.g., prefrontal) are strikingly similar to those observed in the motor domain (Alexander, DeLong, & Strick, 1986), raising the possibility that the BG participate in cognitive decision making in an analogous fashion to their role in motor control (Beiser & Houk, 1998; Frank, 2005a; Frank & Claus, 2006; Frank, Loughry, & O'Reilly, 2001; Middleton & Strick, 2000, 2002). Studies with Parkinson's patients, who have severely depleted levels of dopamine (DA) in the BG (Kish, Shannak, & Hornykiewicz, 1988), have provided insights into the functional roles of the BG/DA system in both motor and higher level cognitive processes (Cools, 2005; Frank, 2005a; Shohamy, Myers, Grossman, Sage, & Gluck, 2005). Of particular recent interest is the finding that deep brain stimulation in the subthalamic nucleus (STN) dramatically improves Parkinson motor symptoms, with both reported enhancements and impairments in cognition (Karachi et al., 2004; Witt et al., 2004). Because the BG consists of a complex network of dynamically interacting brain areas, a mechanistic understanding of exactly how the STN participates in response selection and decision making is difficult to develop with traditional box and arrow models. Computational models

* Tel.: +1 520 626 4787; fax: +1 520 621 9306.
E-mail address: mfrank@u.arizona.edu.
URL: http://www.u.arizona.edu/~mfrank/.

[1] Portions of this paper were previously presented in conference format at the International Workshop on Models of Natural Action Selection (Frank, 2005b).

0893-6080/$ - see front matter © 2006 Elsevier Ltd. All rights reserved.
doi:10.1016/j.neunet.2006.03.006

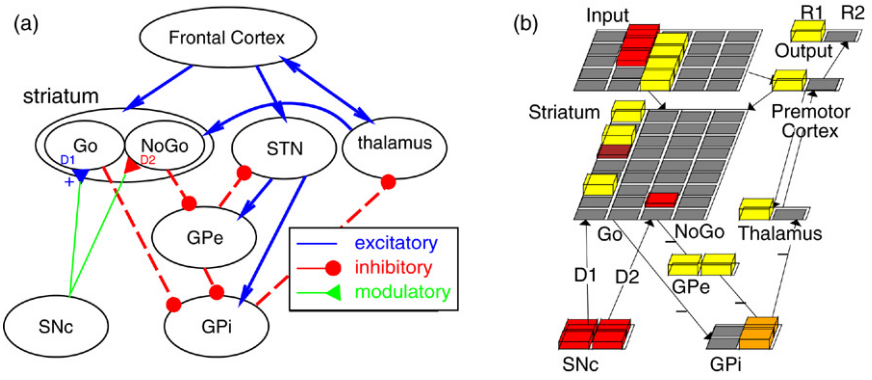

Fig. 1. (a) The striato-cortical loops, including the direct ("Go") and indirect ("NoGo") pathways of the basal ganglia. The Go cells disinhibit the thalamus via GPi, thereby facilitating the execution of an action represented in cortex. The NoGo cells have an opposing effect by increasing inhibition of the thalamus, suppressing actions from getting executed. Dopamine from the SNc projects to the dorsal striatum, causing excitation of Go cells via D1 receptors, and inhibition of NoGo via D2 receptors. GPi: internal segment of globus pallidus; GPe: external segment of globus pallidus; SNc: substantia nigra pars compacta; STN: subthalamic nucleus. (b) The Frank (2005a, 2005b) neural network model of this circuit (squares represent units, with height reflecting neural activity). The Premotor cortex selects an output response via direct projections from the sensory input, and is modulated by the BG projections from thalamus. Go units are in the left half of the striatum layer; NoGo units are in the right half, with separate columns for the two responses (R1 and R2). In the case shown, striatum Go is stronger than NoGo for R1, inhibiting GPi, disinhibiting thalamus, and facilitating R1 execution in cortex. A tonic level of dopamine is shown in SNc; a burst or dip ensues in a subsequent error feedback phase (not shown), driving Go/NoGo learning. The contributions of the STN were omitted from this model, but are explored in the current simulations (Fig. 2).

that explore the dynamics of BG network activity are therefore useful tools for providing insight into these issues, and in turn, how they affect individuals with Parkinson's disease and related disorders.

In this paper, I review converging evidence for a mechanistic, functional account of how interacting areas within the BG-frontal system learn to select adaptive responses and participate in cognitive decision making, as informed by prior computational simulations. I then present a neural network model that explores the unique contribution of the STN within the overall BG circuitry. The simulations reveal that the STN can dynamically control the threshold for executing a response, and that this function is adaptively modulated by the degree to which multiple competing responses are activated, as in difficult decisions. It is concluded that the STN may be essential to allow all information to be integrated before making decisions, and thereby prevents impulsive or premature responding during high-conflict decision trials. Furthermore, analysis of the dynamics of activity within various BG areas during response selection in intact and simulated Parkinson states demonstrates a striking relationship to the same patterns observed physiologically, providing support for the model's biological plausibility and further insight into the neural processes underlying response selection.

2. Overall BG network functionality

The "standard model" proposes that two BG pathways independently act to selectively facilitate the execution of the most appropriate cortical motor command, while suppressing competing commands (Albin, Young, & Penney, 1989; Mink, 1996). Two main projection pathways from the striatum go through different BG nuclei on the way to thalamus and up to cortex (Fig. 1(a)). Activity in the direct pathway sends a "Go" signal to facilitate the execution of a response considered in cortex, whereas activity in the indirect pathway sends a "NoGo" signal to suppress competing responses (Alexander & Crutcher, 1990a; Gerfen & Wilson, 1996). More specifically, striatal Go cells directly project to and inhibit the *internal* segment of the globus pallidus (GPi), which in turn disinhibits the thalamus, ultimately facilitating the execution of cortical motor responses. Conversely, striatal NoGo cells project to and inhibit the *external* segment of the globus pallidus (GPe), releasing the tonic inhibition of GPe onto GPi, and therefore having an opposing effect on motor activity. Dopamine modulates the relative balance of these pathways by exciting synaptically-driven activity in Go cells via D1 receptors, while inhibiting NoGo activity via D2 receptors (Aubert, Ghorayeb, Normand, & Bloch, 2000; Brown et al., 2004; Frank, 2005a; Gerfen, 1992; Hernandez-Lopez, Bargas, Surmeier, Reyes, & Galarraga, 1997; Hernandez-Lopez et al., 2000; Joel & Weiner, 1999). Physiological evidence for this model comes from studies showing opposite effects of D1 and D2 agents on activity within the two types of cells and BG output nuclei (Boraud, Bezard, Bioulac, & Gross, 2002; Gerfen, 2000; Gerfen, Keefe, & Gauda, 1995; Robertson, Vincent, & Fibiger, 1992; Salin, Hajji, & Kerkerian-Le Goff, 1996). Moreover, the general aspects of this model have been successfully leveraged to explain various motor deficits observed in patients with BG dysfunction (e.g., Albin et al. (1989)).

Recently, several researchers have pointed out that the simplest version of the standard BG models is inadequate, and that a more advanced dynamic conceptualization of BG function is required, motivating the use of computational modeling (Bar-Gad, Morris, & Bergman, 2003; Brown et al., 2004; Frank, 2005a; Gurney et al., 2001). Others question the most basic assumptions of the standard model, suggesting that the segregation of the "direct" and "indirect" BG pathways is not as clear as once thought (Kawaguchi, Wilson, & Emson, 1990; Levesque & Parent, 2005; Wu, Richard, & Parent,

2000). These studies found that rather than the striatum having separate populations projecting to GPi and GPe, virtually all striatal cells projected to GPe, while there was still a subpopulation that *also* projected to GPi. While these data certainly challenge the simplest version of the direct/indirect model, one need not "throw the baby out with the bathwater" and reject the standard model altogether, especially given its contribution to many theoretical and practical advances. First, putative NoGo cells remain evident in the above studies, in that many cells only projected to GPe and not GPi. Second, the possibility that Go cells projecting to GPi also project to GPe may signify that rather than computing raw Go signals, the BG output may compute the temporal derivative of these signals. According to this scheme, a striatal Go signal would first inhibit the GPi and disinhibit the thalamus, thereby facilitating the cortical response as proposed by the standard model. Concurrently, the same striatal activity would inhibit GPe, opposing the initial facilitation via GPe–GPi disinhibition. Importantly, this opposing signal would be temporally delayed relative to the initial facilitation (given the extra synapse and slower time constant associated with disinhibition). The net result would be that Go signals lead to short-lasting GPi pauses and associated thalamic bursting activity. This overall functionality may be useful for rapid sequencing of several sub-motor commands, in which an individual command should be facilitated and then immediately suppressed in favor of the subsequent command. Consistent with this idea, motor cortical neurons that send their main axons to the pyramidal tract (and therefore directly involved in movement) also send collaterals only to NoGo-type striatal cells (Lei, Jiao, Del Mar, & Reiner, 2004), which would then act to suppress/terminate the movement. Overall, while some of the anatomical details are undoubtedly missing from any computational model of the BG, this simplification enables analysis of the *dynamics* of activity among multiple BG regions which is not feasible with detailed but static anatomical diagrams.

2.1. A model of reinforcement learning and decision making in Parkinson's disease

Studies with patients with Parkinson's disease (PD) provide insight into the response selection/decision making functions of the basal ganglia. PD is characterized by death of midbrain dopamine cells projecting to the BG (Kish et al., 1988), and associated motor symptoms including tremor, rigidity, and slowness of movement (McAuley, 2003). PD patients also have a variety of cognitive deficits, ranging from procedural learning to working memory, decision making and attention (Ashby, Noble, Ell, Filoteo, & Waldron, 2003; Brown & Marsden, 1988; Cools, Barker, Sahakian, & Robbins, 2001b, 2003; Frank, 2005a; Frank, Seeberger, & O'Reilly, 2004; Gotham, Brown, & Marsden, 1988; Jackson, Jackson, Harrison, Henderson, & Kennard, 1995; Knowlton, Mangels, & Squire, 1996; Maddox & Filoteo, 2001; Rogers et al., 1998; Shohamy et al., 2005). Given the proposed BG role in selecting among various competing low-level motor responses by modulating frontal motor activity, and the parallel circuits linking the BG with more frontal cognitive areas (Alexander et al., 1986; Middleton & Strick, 2000, 2002), it is natural to extend this action selection functionality to include higher-level cognitive decisions (Frank, 2005a; Frank & Claus, 2006; Frank et al., 2001; Houk, 2005). While others explore BG mechanisms that lead to adaptive selection of a salient response in the face of less salient competitors (e.g., Gurney et al. (2001) and Gurney, Humphries, Wood, Prescott, and Redgrave (2004)), a complementary question is how the BG *learn* which action to select. This question is relevant because the most salient input may not always be the best choice.

Previous computational modeling of the basal ganglia/dopamine system provided an explicit formulation that attempts to address this question and ties together various cognitive deficits in Parkinson's disease (PD) (Frank, 2005a). Specifically, the model (Fig. 1(b)) posits that phasic changes in DA during error feedback are critical for modulating Go/NoGo representations in the BG that facilitate or suppress the execution of motor commands. The main assumption was that during positive and negative feedback (e.g., when participants are told that their responses were correct or incorrect), bursts and dips of DA occur that drive learning about the response just executed. This assumption was motivated by a large amount of evidence for bursts and dips of DA during rewards or their absence in animals, respectively (Bayer & Glimcher, 2005; Pan, Schmidt, Wickens, & Hyland, 2005; Satoh, Nakai, Sato, & Kimura, 2003; Schultz, 2002; Schultz, Dayan, & Montague, 1997). These DA changes have also been inferred to occur in humans receiving positive and negative feedback in cognitive tasks (Delgado, Locke, Stenger, & Fiez, 2003; Frank, Woroch, & Curran, 2005; Holroyd & Coles, 2002). Moreover, these phasic changes in DA modulate neuronal excitability, and may therefore act to reinforce the efficacy of recently active synapses (e.g., Hebb (1949); see Mahon, Casassus, Mulle, and Charpier (2003) for evidence of Hebbian learning in striatum), leading to the learning of rewarding behaviors. In the model, "correct" responses are followed by transient increases in simulated DA that enhance synaptically-driven activity in the direct/Go pathway via simulated D1 receptors, while concurrently suppressing the indirect/NoGo pathway via simulated D2 receptors (for detailed neurobiological support, see Brown et al. (2004), Frank (2005a) and Frank and O'Reilly (2006)). This drives Go learning, enabling the model to select responses that on average result in positive feedback. Conversely, phasic dips in DA following incorrect responses release NoGo neurons from the suppressive influence of DA, allowing them to be further excited by corticostriatal glutamate, and driving NoGo learning.[2] Without ever having access to a supervised training signal as to which response *should* have been selected, over the course of training intact networks nevertheless learned how to respond in probabilistic classification tasks, similarly to healthy participants. When 75% of units in the SNc DA layer of the model were lesioned to simulate the approximate amount of damage in PD patients, the model was impaired similarly to patients.

[2] See Frank and O'Reilly (2006) for more biological justification, including discussion on how DA dips can be effective learning signals despite the already low tonic firing rates of DA neurons.

The details of the BG model are described in Frank (2005a, 2005b). In brief, the premotor cortex represents and "considers" two possible responses (R1 and R2) for each input stimulus. The BG system modulates which one of these responses is facilitated and which is suppressed by signaling Go or NoGo to each of the responses. The four columns of units in the striatum represent, from left to right, Go-R1, Go-R2, NoGo-R1 and NoGo-R2. In the absence of synaptic input, GPi and GPe units are tonically active. Go and NoGo representations for each response compete at the level of GPi, such that stronger Go representations lead to disinhibition of the corresponding column of the thalamus, which in turn amplifies and facilitates the execution of that response in premotor cortex. Concurrently, the alternative response is suppressed.

Striatal Go/NoGo representations are learned via phasic changes in simulated DA firing in the SNc layer during positive and negative reinforcement. After correct responses, increases in DA firing excite Go units for the just-selected response, while suppressing NoGo units, via simulated D1 and D2 receptors. Conversely, decreases in DA after incorrect responses, together with corticostriatal glutamate release, results in increased NoGo activity for that response. This DA modulation of Go/NoGo activity drives learning as described above. This overall functionality is consistent with recent observations that separate striatal populations encode both positive and negative action values (Samejima, Ueda, Doya, and Kimura (2005)), potentially corresponding to Go and NoGo cells, respectively.

As DA bursts and dips reinforce Go and NoGo representations in the BG, our model showed that the most adaptive (i.e., rewarding) responses represented in premotor areas will tend to get facilitated while less adaptive ones are suppressed. Further, as the BG learns to facilitate adaptive responses, the associated premotor representations become enhanced (via Hebbian learning). In this way, DA reward processes within the BG may ingrain prepotent motor "habits" in frontal cortical areas (Frank, 2005a; Frank & Claus, 2006). Once these habits are ingrained, there is less need for selective facilitation by the BG. This is consistent with observations that dopaminergic integrity within the BG is much more critical for the acquisition rather than the execution of instrumental responses (Choi, Balsam, & Horvitz, 2005; Parkinson et al., 2002; Smith-Roe & Kelley, 2000), and with recent physiological observations that learning-related activity is initially seen in the BG, and is only observed later in frontal cortex (Delgado, Miller, Inati, & Phelps, 2005; Pasupathy & Miller, 2005).

2.2. Modeling dopaminergic medication effects on cognitive function in PD

The same model was used to explain certain negative effects of dopaminergic medication on cognition in PD (Frank, 2005a). While medication improves cognitive performance in some attentional tasks (Cools, Barker, Sahakian, & Robbins, 2001a; Shohamy et al., 2005; Swainson et al., 2000), it actually impairs performance in probabilistic reversal learning (Cools, 2005; Cools et al., 2001a; Swainson et al., 2000), that is when having to make decisions that require learning to overriding previously adaptive responses in favor of those that were less adaptive.

In order to simulate medication effects, it was hypothesized that medication increases the tonic level of DA, but that this interferes with the natural biological system's ability to dynamically regulate phasic DA changes. Specifically, phasic DA dips during negative feedback may be partially blocked by DA agonists (or increases in tonic DA by L-Dopa; Pothos, Davila, and Sulzer (1998)) that continue to bind to receptors. When this was simulated in the model, selective deficits were observed during probabilistic reversal, despite equivalent performance in the acquisition phase (Frank, 2005a), mirroring the results found in medicated patients. Because increased tonic levels of DA suppressed the indirect/NoGo pathway, networks were unable to learn "NoGo" to override the prepotent response learned in the acquisition stage. This account is consistent with similar reversal deficits observed in healthy participants administered an acute dose of bromocriptine, a D2 agonist (Mehta, Swainson, Ogilvie, Sahakian, & Robbins, 2001), and with several other learning deficits induced by DA medications that are consistent with NoGo impairments (Bokura, Yamaguchi, & Kobayashi, 2005; Charbonneau, Riopelle, & Beninger, 1996; Cools et al., 2003; Czernecki et al., 2002; Frank et al., 2004; Ridley, Haystead, & Baker, 1981; Shohamy, Myers, Geghman, Sage, & Gluck, 2006; Smith, Neill, & Costall, 1999).

2.3. Empirical tests of the model

Recently, we have tested various aspects of the hypothesized roles of the basal ganglia/dopamine system across both multiple cognitive processes. First, we demonstrated support for a central prediction of the Frank (2005a) model regarding BG dopamine involvement in "Go" and "NoGo" cognitive reinforcement learning. We tested Parkinson patients on and off medication, along with healthy senior control participants (Frank et al., 2004). We predicted that decreased levels of dopamine in Parkinson's disease would lead to spared NoGo learning, but impaired Go learning (which depends on DA bursts). We further predicted that dopaminergic medication should alleviate the Go learning deficit, but would block the effects of dopamine dips needed to support NoGo learning. Results were consistent with these predictions. In a probabilistic learning task, all patients and aged-matched controls learned to make choices that were more likely to result in positive rather than negative reinforcement. The difference was in their learning biases: patients taking their regular dose of dopaminergic medication implicitly learned more about the positive outcomes of their decisions (i.e., they were better at Go learning), whereas those who had abstained from taking medication implicitly learned to avoid negative outcomes (better NoGo learning). Age-matched controls did not differ in their tendency to learn more from the positive/negative outcomes of their decisions. We have also found the same pattern in young healthy participants administered dopamine D2 receptors agonists and antagonists, which at low doses modulate striatal dopamine release (Frank & O'Reilly, 2006). Again, dopamine increases improved Go learning and impaired NoGo learning, while decreases had the opposite effect.

The same BG modeling framework accurately predicted the pattern of event-related potentials recorded from healthy participants who were biased to learn more from either positive or negative reinforcement (Frank et al., 2005), as well as a counter-intuitive improvement in BG/DA-dependent choices when hippocampal explicit memory systems were taken offline by the drug midazolam (Frank, O'Reilly, & Curran, 2006). Finally, in the D2 drug study mentioned above, the same BG/DA effects extended to higher level working memory tasks that required paying attention to task-relevant (i.e., positively valenced) information while ignoring distracting (negative) information (Frank & O'Reilly, 2006), consistent with predictions from extended BG models that include interactions with prefrontal cortex in working memory and attention (Frank & Claus, 2006; Frank et al., 2001; O'Reilly & Frank, 2006).

3. Integrating contributions of the subthalamic nucleus in the model

Despite its success in capturing dopamine-driven individual differences in learning and attentional processes, the above model falls short in its ability to provide insight into BG dynamics that depend on the subthalamic nucleus (STN). The model was designed to simulate how the BG can learn to selectively facilitate (Go) one response while selectively suppressing (NoGo) another. Because the projections from the STN to BG nuclei (GPe and GPi) are diffuse (Mink, 1996; Parent & Hazrati, 1995), it may not be well suited to provide selective (focused) modulation of specific responses, and was therefore omitted from the model. Instead the model simulated the focused projections from striatum to GPi and GPe, as well as the focused projections from GPe to GPi, to demonstrate how direct and indirect pathways may compete with one another at the level of each response, but may act in parallel to facilitate and suppress alternative responses (see Frank (2005a, 2005b) for details and discussion).

Nevertheless, there is substantial evidence that the STN is critically involved in both motor control and cognitive processes (Baunez et al., 2001; Bergman, Wichmann, Karmon, & DeLong, 1994; Boraud et al., 2002; Karachi et al., 2004; Witt et al., 2004). Further, other computational models of action selection also implicate a key role of the STN (Brown et al., 2004; Gurney et al., 2001; Rubchinsky, Kopell, & Sigvardt, 2003). The present model explored the contributions of the STN within the computational framework of the previous model of cognitive reinforcement learning and decision making (Frank, 2005a), scaled up to include four competing responses. By virtue of its diffuse connectivity to BG nuclei, the STN may support more of a global modulatory signal on facilitation and suppression of *all* responses, rather than modulating the execution of any *particular* response. The simulations described below reveal that this global modulatory signal could not be replaced by a simple response threshold parameter, because its effects are dynamic as response selection processes evolve, and its efficacy depends on excitatory input from premotor cortex. Further, simulated dopamine depletion in

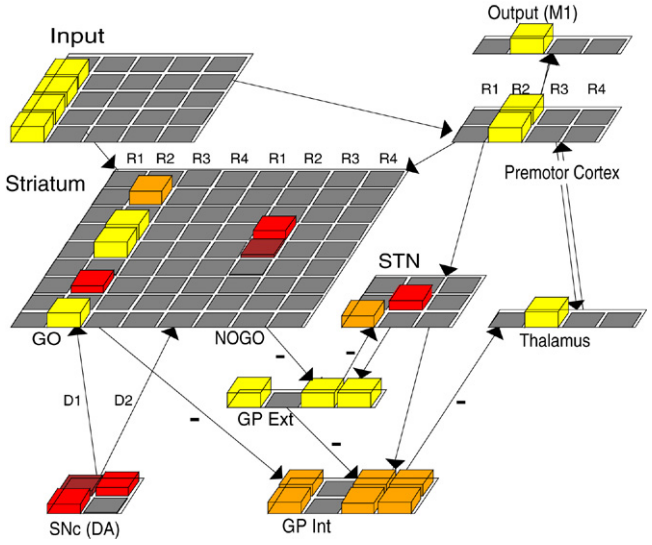

Fig. 2. The subthalamic nucleus is incorporated into a scaled-up model that includes four competing responses (R1–R4). The STN receives excitatory projections from pre/motor cortex in the "hyperdirect pathway" and excites both GPi and GPe; GPe provides inhibitory feedback on STN activity.

the augmented model results in emergent oscillations in the STN and BG output structures, which have been documented empirically and are thought to be the source of Parkinson's tremor. Finally, the simulations show that the STN may be critical for action selection processes to prevent premature responding, so that all potential responses are considered before facilitating the most appropriate one.

3.1. STN connectivity with other BG and cortical structures

The STN was included in the model in accordance with known constraints on its connectivity in BG circuitry, as depicted in Fig. 2. First, the STN forms part of the "hyperdirect" pathway, so named because cortical activity targets the STN, which directly excites GPi, bypassing the striatum altogether (Nambu et al., 2000). Thus initial activation of the STN by cortex leads to an initial excitatory drive on the already tonically active GPi, effectively making the latter structure more inhibitory on the thalamus, and therefore less likely to facilitate a response. Further, the STN gets increasingly excited with increasing cortical activity. Thus, if several competing responses are activated, the STN sends a stronger "Global NoGo" signal which allows the BG system to fully consider all possible options before sending a Go signal to facilitate the most adaptive one.

Second, the STN and GPe are reciprocally connected in a negative feedback loop, with the STN exciting the GPe and the GPe inhibiting the STN (Parent & Hazrati, 1995). As noted above, the connections from STN to GPe are diffuse, and therefore are not likely to be involved in suppressing a specific response. Of the STN neurons that project to GPe, the vast majority also project to GPi (Sato, Parent, Levesque, & Parent, 2000). In the model, each STN neuron receives projections from two randomly selected GPe neurons. This was motivated by data showing that multiple GPe neurons converge on a single

STN neuron (Karachi et al., 2004). In contrast, each GPe neuron receives diffuse projections from all STN neurons (but with randomly different synaptic weights). Please see Appendix A for additional model equations and parameters.

3.2. Simulated BG firing patterns during response selection

The firing patterns of simulated BG, thalamic, and cortical structures are shown in a representative response selection trial in Fig. 3(b). Upon presentation of a stimulus input, multiple competing responses are simultaneously but weakly activated in premotor cortex. Concurrently, response-specific striatal NoGo signals cause GPe activity to decrease. The combined effects of initial cortical activity and decreases in GPe activity produce an initial STN surge at approximately 20 cycles of network settling (in this particular trial). This STN activity is excitatory onto all GPi cells, preventing them from getting inhibited by early striatal Go signals that would otherwise facilitate response execution. However, STN activity also excites GPe neurons, which in turn reciprocally inhibit the initial STN activity surge, thereby removing the Global NoGo signal. At this point, a striatal Go signal for a particular response can then inhibit the corresponding GPi column, resulting in thalamic disinhibition and subsequent selection of that response in motor cortex. Because activity values are displayed in terms of average activity across each layer, the selection of a single motor response together with suppression of other responses results in a net decrease in average premotor cortex activity. Finally, in some trials, a late striatal NoGo signal causes GPe inhibition and a second surge in STN activity.

The above description of STN dynamics is consistent with data from physiological recordings showing an early discharge in STN cells during either response selection or direct cortical stimulation (Kolomiets et al., 2001; Magill, Sharott, Bevan, Brown, & Bolam, 2004; Nambu et al., 2000; Wichmann, Bergman, & DeLong, 1994; Fig. 3(a)). Moreover, this model is a formal implementation of existing theoretical constructs regarding the role of the STN in initial response suppression, followed by a direct pathway response facilitation, and then finally an indirect pathway response termination (Maurice, Deniau, Glowinski, & Thierry, 1998; Nambu, Kaneda, Tokuno, & Takada, 2002). This dynamic functionality of BG activity in response selection may have important implications for higher level decision making, as described below. But first, an obvious question is whether this model also accounts for patterns of activity in the dopamine-depleted state, for which there are abundant data.

3.3. Dopamine depletion is associated with subthalamic and pallidal oscillations

Dopaminergic depletion in Parkinson's disease is associated with changes in the firing patterns and activity levels in various BG nuclei (Boraud et al., 2002; Mink, 1996). Lowered dopamine levels result in excessive striatal NoGo (indirect pathway) activity, and concomitant decreases in GPe and increases in GPi activity (Boraud et al., 2002). Parkinsonism is also associated with increased STN activity, thought to arise from reduced GABAergic GPe input (DeLong, 1990; Miller & Delong, 1987). Perhaps most notably, DA depletion has been reliably associated with low-rate oscillatory bursting activity in STN, GPe and GPi, which is correlated with the development of Parkinson's tremor (Bergman et al., 1994, 1998; Levy, Hutchison, Lozano, & Dostrovsky, 2000; Raz, Vaadia, & Bergman, 2000). Finally, these oscillations and associated PD symptoms are eliminated in DA-depleted animals after they are given experimental STN lesions (Bergman, Wichmann, & DeLong, 1990; Ni, Bouali-Benazzouz, Gao, Benabid, & Benazzouz, 2000).[3]

Interestingly, when Parkinson's disease was simulated in the model, these effects emerged naturally (Figs. 3(d) and 4(a) for averaged activity across multiple trials). First, simulated DA depletion (setting tonic SNc firing rates to zero) led to increased striatal NoGo activity, as described previously (Frank, 2005a). Second, this led to increased overall STN and GPi activity, consistent with empirical recordings. Third, DA depletion led to emergent network oscillations between the STN, GPi and GPe layers, which have been linked to Parkinson's tremor as described above. These oscillations were more prominent when no motor response was selected (Fig. 3(d)), consistent with empirical observations that movements suppress STN oscillations (Amirnovin, Williams, Cosgrove, & Eskandar, 2004), and with the fact that tremor is usually seen in the resting state. Further, oscillations could be observed even when layer activity levels were averaged across multiple trials (Fig. 4(a)), suggesting that they are highly regular (for a given network configuration; random GPe–STN connectivity leads to variability across networks). Finally, simulated STN lesions (by removing the STN layer from processing) in DA-depleted networks normalized GPi activity and eliminated GPi/GPe oscillations (Fig. 4(b)). As mentioned above, this same pattern of results has been observed as a consequence of STN lesions in the dopamine-depleted animal (Ni et al., 2000). Similarly, because cortex provides the primary excitatory input onto STN, simulated cortical lesions also eliminated oscillations (data not shown), consistent with experimental data (Magill, Bolam, & Bevan, 2001). In sum, the close correspondence with various effects of DA manipulation on BG firing patterns supports the model's biological plausibility, particularly in light of the fact that it was not specifically designed to reproduce these physiological data. Next, the relevance of these patterns to response selection processes are considered.

3.4. The STN and action selection

If STN lesions improve Parkinson symptoms, it is natural to consider what deleterious effects they might have. In other words, what is the essential computational function of the STN in action selection/decision making? Some evidence comes from the animal literature showing that STN lesions impair response selection processes, and lead to premature responding when having to suppress competing responses (Baunez et al., 2001; Baunez & Robbins, 1997). This leaves open the possibility that the Global NoGo signal provided by the STN

[3] The therapeutic effects of human STN deep brain stimulation are thought to rely on similar mechanisms (Benazzouz & Hallett, 2000; Meissner et al., 2005).

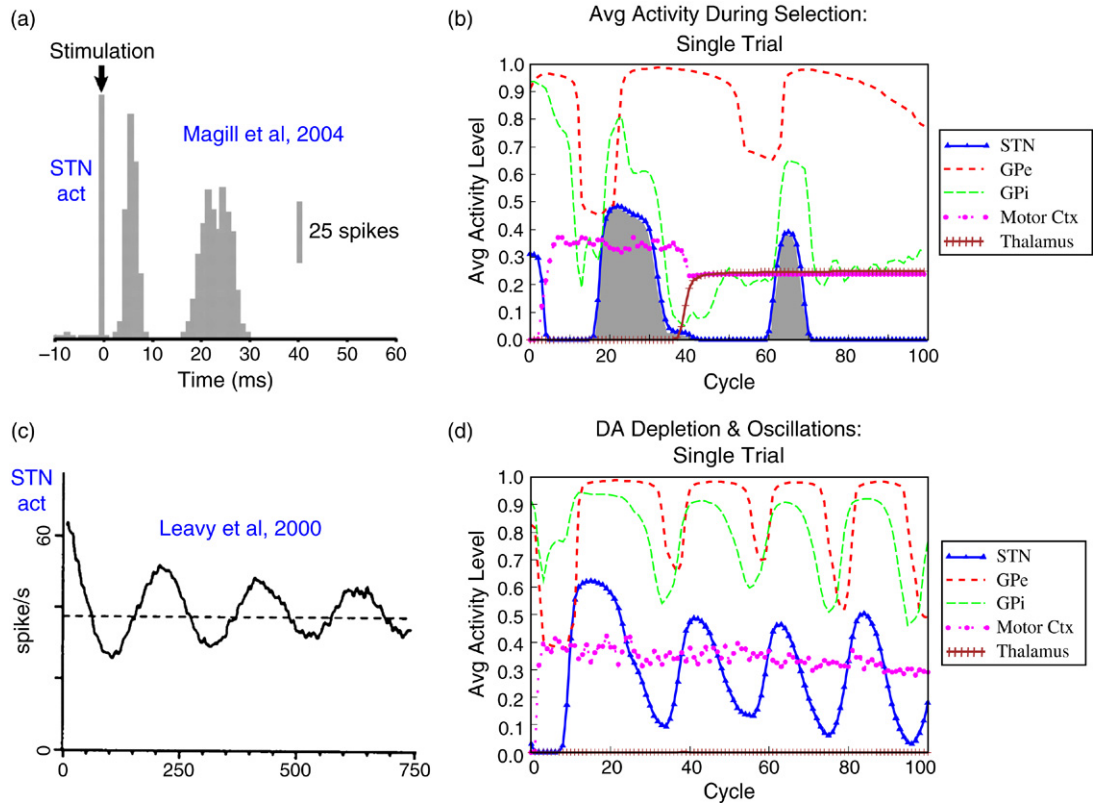

Fig. 3. (a) Physiological STN recordings following cortical stimulation, adapted from Magill et al. (2004), showing two bursts of activity. The same patterns are observed during natural response selection (Wichmann et al., 1994). (b) A single response selection trial in the model. Activity levels (normalized with respect to maximal firing rates) are averaged across units within each layer as a function of network settling cycles. STN activity is shaded in grey for comparison with (a). Initially, multiple simultaneously active and competing premotor cortex responses excite STN via the hyperdirect pathway (≈cycle 20). The resulting "Global NoGo" signal prevents premature responding by keeping GPi units tonically active. Sustained GPe activity subsequently inhibits STN (cycle 35), turning off the Global NoGo signal. Striatal Go signals then facilitate a response by inhibiting GPi and disinhibiting Thalamus (cycle 40). This is reflected in premotor cortex as an overall decrease in activity, due to suppression of the three alternative responses. Finally, striatal NoGo signals inhibit GPe, causing a second STN surge (cycle 60), thought to terminate the executed response. (c) Dopamine depletion in animals and humans leads to oscillatory activity in STN (shown here from Levy et al. (2000)) and GPe and GPi (not shown), which are associated with Parkinson's tremor. (d) Simulated dopamine depletion in the model leads to emergent network oscillations in STN, GPi and GPe.

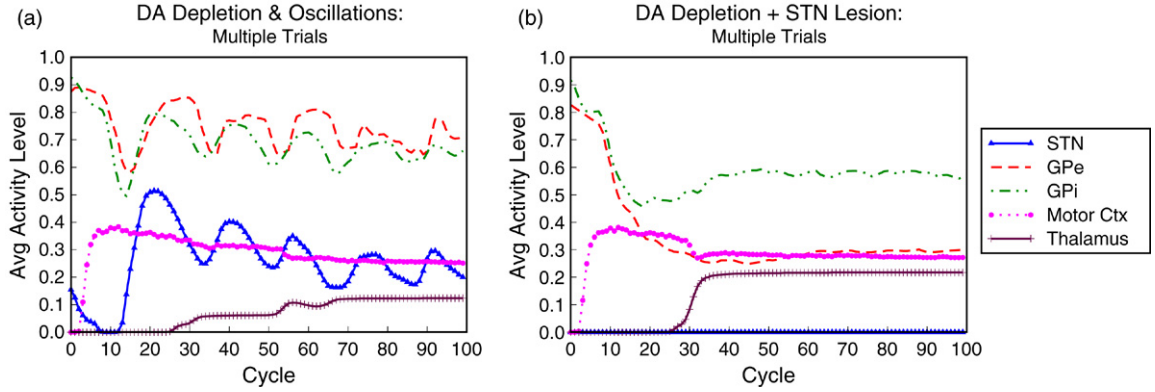

Fig. 4. Average unit activity during response selection as a function of network settling cycles. Data are averaged across units within each area, and across 100 trials. (a) Simulated Parkinsonism (DA depletion) led to oscillations in STN, GPe and GPi, as is observed in DA-depleted animals. Regular oscillations are observed despite averaging across multiple trials, but are dampened relative to those observed in any single trial (e.g., Fig. 3(d)). (b) STN lesions in DA-depleted networks eliminated the oscillations observed in GPe and GPi, and facilitated response selection, as has been observed in experimental animals (Bergman, Wichmann, & DeLong, 1990; Ni et al., 2000).

is adaptive and allows the animal (or the model) sufficient time to consider all possible responses before selecting the most adaptive one. This hypothesis is further supported by observations that low-amplitude STN stimulation decreases premature responding in rats (Desbonnet et al., 2004). The question is whether a formal simulation of STN involvement in

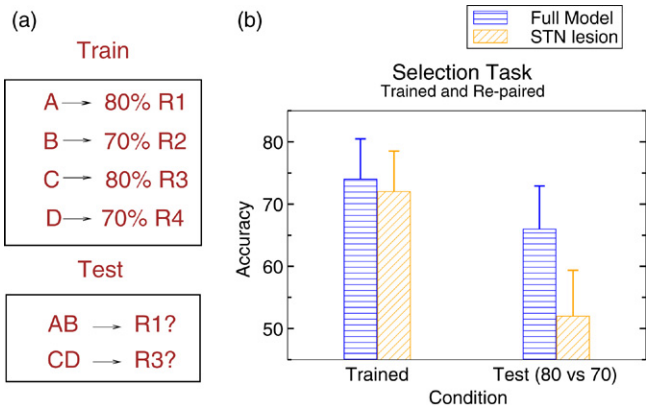

Fig. 5. (a) Response selection paradigm. Four cues are independently associated with one of four possible responses. Responses R1 and R3 are reinforced on 80% of trials in the presence of cues A and C, respectively. R2 and R4 are reinforced in 70% of trials to cues B and D. The test phase measures the network's ability to choose the 80% over the 70% response when presented with cues A and B or C and D together. (b) Both intact networks and those with STN lesions successfully learned to choose the appropriate response for each training cue. STN lesions selectively impaired selection among two competing responses, due to premature responding before being able to integrate over all possible responses. Error bars reflect standard error across 25 runs of the model with random initial synaptic weights.

BG dynamics can account for these data in a response selection paradigm.

To address this question, a reinforcement learning paradigm was simulated in which the model was presented with one of four cues, each represented by a column of simultaneously active units in the input layer. The network's task was to select one of four possible responses for each cue (Fig. 5(a)). "Feedback" is then provided to the network by either increasing or decreasing dopamine levels. The network learns based on the difference in Go/NoGo activity levels in the response selection and feedback phase, as detailed in Frank (2005a) and in Appendix A.

The stimulus-response mappings are probabilistic, such that the optimal response for some cues will lead to positive reinforcement (DA bursts) in 80% of trials; in the remaining 20% of trials some alternative response is reinforced. For all incorrect responses, DA dips are applied. Other cue-response mappings are less reliable, such that the optimal response is positively reinforced in only 70% of trials. Networks were trained with 15 epochs consisting of 10 trials of each stimulus cue. As in previous simulations, BG networks should be able to learn to select the response most associated with positive reinforcement based on Go/NoGo learning within the striatum (Frank, 2005a; Frank et al., 2004). But in these prior simulations, the STN was not incorporated and was therefore not critical for this learning to take place.

To determine whether the STN is beneficial for selecting among multiple competing responses, a test phase was administered. Two cues were presented in the input simultaneously, one of which had been associated with 80% positive reinforcement if responded to by one response, while the other had been associated with 70% positive reinforcement for an alternative response. Although the models had not been trained with these stimulus combinations, they should nevertheless be able to select the response that was most likely to result in positive reinforcement (i.e. the 80% response). However, premature responding could result in selection of the 70% reinforced response if its corresponding striatal Go signal happened to get active prior to that of the 80% response (due to noise in striatum or in the premotor representations themselves). This is precisely the kind of situation in which an initial STN Global NoGo signal may be useful, so that the network can integrate over multiple possible responses before selecting the most appropriate one.

Simulation results were consistent with this depiction (Fig. 5(b)). While there was no difference between networks in their ability to select the most adaptive response for each cue, models with STN lesions were impaired at making high-conflict decisions (e.g., choosing the best among two positively associated responses). This result is consistent with the notion that the STN is critical for preventing premature responding, as networks without the STN were equally likely to choose the 70% response as the 80% response.[4]

Further support for the above conclusion comes from analysis of model dynamics during the learning and testing stages of this task (Fig. 6, averaged across trials of each type). This analysis reveals that the strength of the initial STN Global NoGo signal is modulated by the degree of response conflict present in premotor cortex, such that if multiple competing responses are active, the network may take more time to select a given one. Recall that in training trials, only one cue was presented that had been most reliably associated with a single response. Thus a minimal amount of response conflict in premotor cortex led to a small initial STN Global NoGo signal (Fig. 6(a)). In these trials, virtually all responses were selected by cycle 50 (as evidenced by average Thalamic activity). In contrast, the STN surge occurred earlier and was larger in magnitude during high-conflict test trials (Fig. 6(b)), due to increased cortical synaptic activity (from prior cortical learning) and resulting in slower response execution (slower increase in average Thalamus activity). Notably, models with simulated STN lesions did not demonstrate this modulation of response time by degree of conflict; these models continued to select all responses by cycle 50 in both training and test trials (Fig. 6(c), (d)).

4. Discussion

This work presents a novel computational exploration of the subthalamic nucleus within the overall basal ganglia circuitry. The model integrates various neural and behavioral findings and provides insight into the STN role in response selection

[4] Similar patterns of results were observed in a network that was trained to select among two responses (instead of four; data not shown). Thus whereas previous models omitting the STN were capable of learning complex probabilistic tasks (Frank, 2005a) and produced the correct pattern of DA-dependent learning biases (in terms of Go versus NoGo striatal representations; Frank et al., 2004), the networks in those studies were not specifically tested in their ability to select among two responses that had been independently associated with similar reinforcement values. The present simulations reveal that the STN improves choice selection in these high-conflict decisions.

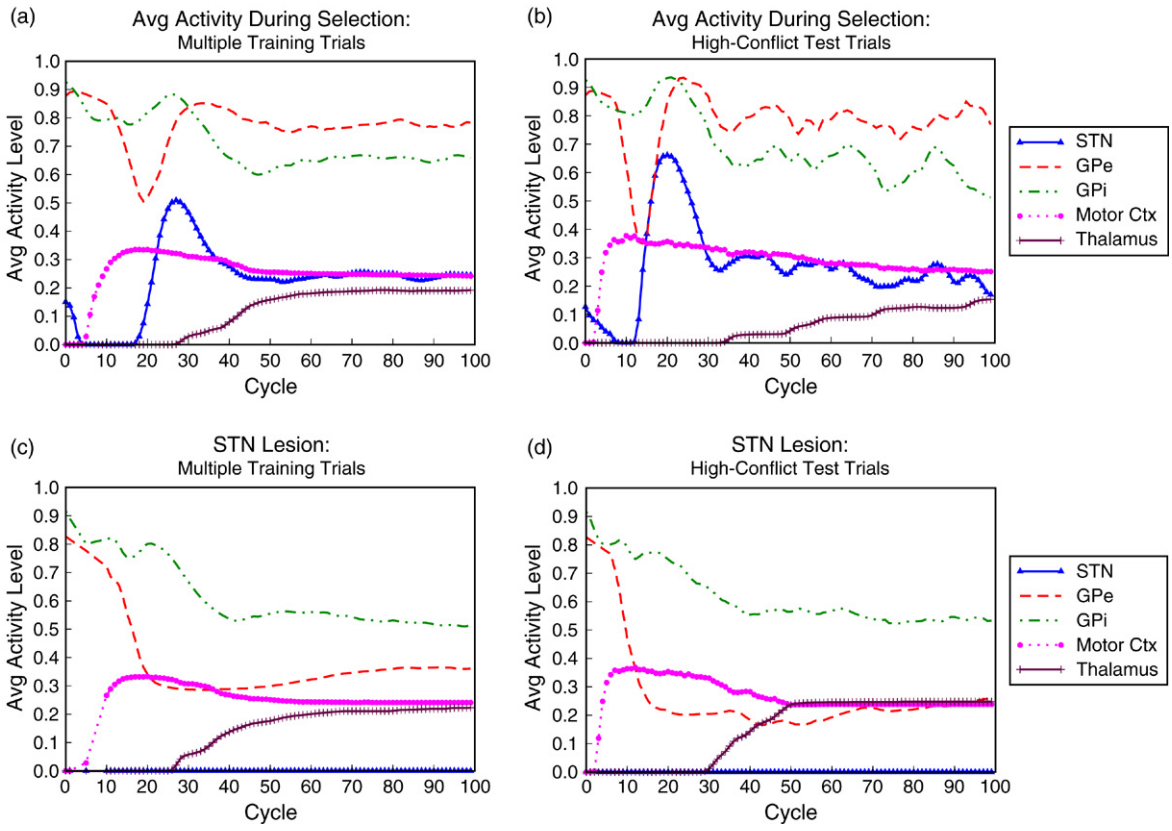

Fig. 6. Average unit activity during response selection as a function of network settling cycles. Data are averaged across units within each area, and across 100 trials of the response selection task depicted in Fig. 5. (a) Intact network. Dynamics are similar to those in a single trial (Fig. 3(b)), but transitions are less clear-cut, because they occur at somewhat different latencies across multiple trials. Nevertheless, virtually all responses during low-conflict training trials were selected by cycle 50, as can be seen by asymptotic thalamic activity. (b) High-conflict test trials (see Fig. 5(a)). Note earlier and larger STN burst due to multiple conflicting cortical responses. Gradual thalamus activity indicates that the STN Global NoGo signal prevented many responses from being selected until later in network settling. (c) In STN lesioned networks, response selection time is similar to that of intact networks for low-conflict training trials. (d) A lack of STN Global NoGo surge in STN-lesioned networks causes premature responding in high-conflict trials; all responses are selected between cycles 30 and 50.

and decision making. Consistent with other BG models (Brown et al., 2004; Gurney et al., 2001), the STN provides a "Global NoGo" signal that suppresses all responses. But the current simulations revealed that this signal is dynamic, such that it is evoked upon initial response initiation, is then inhibited, and is finally reactivated during termination, as has been observed in various physiological observations (Bevan, Magill, Terman, Bolam, & Wilson, 2002; Kolomiets et al., 2001; Magill et al., 2004; Maurice et al., 1998; Nambu et al., 2002; Wichmann et al., 1994). Further, the degree and duration of STN activity is directly driven by the amount of response conflict present in cortical motor representations where conflict emerges as a function of prior learning. These STN dynamics are therefore adaptive in preventing premature responding when multiple competing responses are activated.

In the following sections, I discuss how the model corresponds to various neural and behavioral data in the context of response selection and decision making.

4.1. Model correspondence to BG activity

In accord with physiological observations (Magill et al., 2004; Nambu et al., 2000), a burst in STN activity occurs during the initial stages of response selection. This burst is elicited via the "hyperdirect" pathway, as multiple competing responses are activated in premotor cortical areas. If multiple responses have been associated with adaptive behavior for the particular stimulus context, then these will be more active and will therefore more strongly excite the STN. The resulting Global NoGo signal delays responding until the competition is adequately resolved. When a particular response is facilitated (and the others suppressed), this Global NoGo signal is shut off. This occurs due to a combination of less overall top-down activity from cortex (since alternative responses are no longer active), and inhibition of the STN via GPe activity. Subsequent striatal NoGo activity can then lead to a second surge of STN activity (via GPe inhibition), supporting termination of responses. This overall functionality is consistent with physiological recordings of the STN during response selection and cortical stimulation. Further, recent imaging studies reveal that the human STN is particularly active during high-conflict trials in which prepotent responses are to be inhibited (Aron & Poldrack, 2006).

4.2. Relationship to models of optimal decision making

The very same BG dynamics may serve to optimize response times and decision thresholds depending on task demands, as

has been studied in the context of two-alternative forced choice tasks. Abstract neural and mathematical models of optimal decision making suggest that agents must first integrate over processing noise in order to extract the best possible decision before making a response (e.g., Brown et al. (2005), Ratcliff, Van Zandt, and McKoon (1999) and Usher and McClelland (2001)). These models make contact with electrophysiological findings showing that the firing rates of "decision" neurons in monkey motor areas gradually increase during forced choice tasks, and when these rates cross a decision threshold a choice is executed (Gold & Shadlen, 2002; Schall, 2003). Interestingly, both animals and humans can dynamically and optimally adjust the behavioral threshold for when to execute a response so as to maximize their reward rate, in terms of correct responses per unit time (Bogacz, Brown, Moehlis, Holmes, & Cohen, in press; Simen, Cohen, & Holmes, in this issue).

The adaptive and dynamic STN functionality simulated here is consistent with the formal requirements described by the above models. The STN Global NoGo signal effectively achieves this function by allowing both cortical and striatal signals to accumulate and compete before determining which response to facilitate. Because STN activity is dynamically modulated by the level of response conflict (which itself is determined by prior reward associations), the STN may contribute to optimally maximizing the response time in a given choice situation such that faster responses are achieved for low-conflict decisions, whereas more integration of information can occur for high-conflict decisions. Thus, the STN is predicted to play an important role in classical speed–accuracy trade off effects. This account is also consistent with theoretical perspectives positing that response conflict is represented in rostral anterior cingulate cortex (e.g. Botvinick, Braver, Barch, Carter, and Cohen (2001) and Yeung, Botvinick, and Cohen (2004)); this area is thought to correspond to the monkey rostral cingulate motor zone (Picard & Strick, 1996, 2001), and may represent the activation of multiple competing responses. Notably, this area has direct projections to the STN (Orieux, Francois, Feger, & Hirsch, 2002).

4.3. Effects of STN manipulation on rat choice behavior

The model is also consistent with behavioral findings in rats showing that STN lesions worsen, while STN stimulation improves, premature responding in high-conflict choice selection paradigms (Baunez & Robbins, 1997; Baunez et al., 2001; Desbonnet et al., 2004). This is exactly what is predicted by the dynamic Global NoGo signal hypothesized to depend on the STN during response selection. Nevertheless, a recent finding seem to challenge this account. Winstanley, Baunez, Theobald, and Robbins (2005) found that lesions to the STN *decreased* impulsivity in a delay discounting paradigm, such that rats were more likely to make responses leading to greater long-term rewards instead of those leading to short-term reward. However, the present model would suggest that rather than the STN lesion causing enhanced valuation of the delayed reward *per se*, it may have simply prevented the rat from reliably choosing the option that it actually would have wanted to choose (i.e. the immediate reward), if it only could have had more time to "consider" both options. Given that both choices had some reward value, the lesioned rats may have simply been more likely to choose whichever option they first considered in any given trial, resulting in relatively more choices for the delayed option. Future empirical studies are therefore needed to test this account.

4.4. Neural activity during selection and oscillations following DA depletion

The model also provides evidence for biological plausibility at the neural systems level. First, consistent with neurophysiological observations and as discussed in Frank (2005a), premotor activity is observed within a trial *prior* to that of the striatum (e.g., Alexander and Crutcher (1990b) and Crutcher and Alexander (1990); see also Mink (1996)). This pattern is especially true for well-learned responses, in which premotor cortex can activate the correct response without requiring striatal facilitation. However, *learning*-related activity is observed in striatum prior to premotor cortex (Delgado et al., 2005; Pasupathy & Miller, 2005) — since initial changes in Go/NoGo representations are required before premotor cortex can "stamp-in" the habitual response.

Moreover, in the current simulations, dopamine depletion produced emergent oscillatory activity in the STN and BG output nuclei; these oscillations are reliably observed in DA-depleted animals and humans and are thought to be the source of Parkinson's tremor (Bergman et al., 1994, 1998; Levy et al., 2000; Raz et al., 2000). Similar oscillations were previously described and more extensively explored in a biophysically detailed conductance model of GPe–STN interactions (Terman, Rubin, Yew, & Wilson, 2002). Although that model did not include a striatum, oscillations were induced by applying a constant external inhibitory current to GPe neurons, so as to simulate enhanced NoGo activity in the DA depleted state. While the current model is not as detailed at the GPe/STN unit level (including just three ionic currents; see Appendix A), this simpler implementation enabled tractable investigation of systems-level dynamics among multiple BG and cortical structures and their roles in learning and decision making. That we still observe DA-dependent oscillations provides support for the plausibility of the approach. Further, oscillatory activity in the current model was substantially reduced or eliminated with simulated lesions to STN or motor cortical areas, as has also been shown empirically (Bergman et al., 1990; Magill et al., 2001; Ni et al., 2000). Nevertheless, it is acknowledged that there are multiple possible configurations that could lead to oscillatory behavior in complex interactive circuits such as the BG. Thus the current simulations do not prove that they replicate the exact conditions under which oscillatory activity and tremor is observed in Parkinson's disease, but merely provide a plausible scenario. In particular, it is not currently known whether oscillations stem from DA depletion to the entire BG, or if they would occur with restricted DA depletion within any given area (such as the STN or striatum). It is therefore important to be explicit about the cause of oscillations, so that multiple alternatives could be tested.

In the model, oscillations in various BG nuclei occur via emergent network dynamics following striatal DA depletion. The lack of D2 receptor stimulation leads to overactive NoGo units (which are normally inhibited by dopamine), resulting in excessive inhibition of GPe and, in turn, disinhibition of the STN. Subsequent burst firing in STN propagates back to excite GPe, which now becomes more active and thus begins to inhibit STN. The cycle repeats, leading to oscillations in STN and GPe. Oscillatory activity in GPi is also observed, reflecting afferent activity from both STN and GPe. Note we did not need to simulate DA depletion within STN, GPe and GPi themselves for oscillations to occur; instead they arise from DA depletion to the striatum. (As in Terman et al. (2002), the strength of GPe activity is critical for oscillatory behavior; when the strength of the GPe–STN projection was weakened, oscillations were no longer observed.)

Therefore, in both models, the critical source of oscillations is burst firing within STN after disinhibition by GPe. However, STN disinhibition alone is not sufficient to produce oscillations; additional mechanisms are required to generate bursting. In the current model, top-down excitatory projections are required to generate burst firing only in STN units which are concurrently disinhibited by GPe *and* excited by cortex. Alternatively, this burst firing can occur via intracellular rebound firing mechanisms within STN itself Terman et al. (2002). The current systems-level account is consistent with data showing that top-down cortical projections play a key role in the generation of STN bursts and associated oscillations (Bevan et al., 2002; Magill et al., 2001). Moreover, the model makes the unique prediction that increased response conflict in cortex should drive further burst firing in STN under DA depletion, and that Parkinson's tremor should be exacerbated under these conditions. More generally, because cortical regions representing conflict also become increasingly engaged as participants make errors in cognitive tasks (e.g., Yeung et al. (2004) and Frank et al. (2005)), we predict that detectable increases in tremor should be observed following negative reinforcement.

4.5. Model limitations and future directions

While the current model is an advancement over previous simulations, in that it learns to select among four competing responses (instead of just two) and incorporates the dynamic contribution of the STN, there nevertheless remain several important limitations.

4.5.1. Integrating actor and critic functions of the BG

In this and the previous model (Frank, 2005a), we have not addressed the important question of how reward and loss information is computed by systems upstream of midbrain dopamine neurons; instead we simply assumed this function by externally increasing and decreasing simulated DA levels during positive reinforcement and negative reinforcement, and then examined their effects on learning and decision making in the BG network. Said otherwise, the current work focuses on the *actor* functions of the BG, and simply assumes the *critic* function. In parallel work, we are investigating how interactions between the amygdala and ventral striatal BG regions can support the critic function by learning to associate stimuli with affective states and driving dopaminergic firing in the midbrain (O'Reilly, Frank, Hazy, and Watz (in press); see also Brown, Bullock, and Grossberg (1999) and Houk, Adams, and Barto (1995)). This work provides a biologically explicit mechanism for the widely acknowledged relationship between the firing patterns of dopamine neurons during conditioning paradigms, and those predicted by the abstract mathematical *temporal differences* reinforcement learning algorithm (Schultz et al., 1997; Sutton, 1988). Preliminary simulations demonstrate that the current actor BG model can learn successfully when dopamine firing is computed by the critic model, rather than applying DA values externally.

4.5.2. The roles of serotonin

As described here, the model does not investigate functions of other neuromodulators beside dopamine, including serotonin and norepinephrine, both of which are thought to play key roles in reinforcement learning and decision making (Aston-Jones & Cohen, 2005; Daw, Kakade, & Dayan, 2002; Harley, 2004). Particularly relevant to the current model, a role for serotonin (5-HT) has been implicated in impulsive behavior (Walderhaug et al., 2002; Winstanley, Theobald, Dalley, Glennon, & Robbins, 2004). It is likely that these effects are partially mediated by serotonergic processes within the STN. Serotonergic neurons of the dorsal raphe nucleus innervate the STN (Lavoie & Parent, 1990), where they are excitatory via densely expressed 5-HT2c receptors (Pompeiano, Palacios, & Mengod, 1994; Stanford, Kantaria, Chahal, Loucif, & Wilson, 2005; Xiang, Wang, & Kitai, 2005). In preliminary exploratory simulations of 5-HT function within the STN of the current model, background 5-HT levels modulate the gain of STN neural activity and enhance sensitivity to cortical input. Thus, according to the model, serotonin may effectively enhance STN Global NoGo signals so as to prevent premature responding and slow responding. Consistent with this hypothesis, blockade of 5-HT2c receptors leads to increased premature responding and decreased latency to make a correct response in rats (Winstanley et al., 2004). Further research is necessary to determine whether this serotonergic effect is mediated selectively in the STN.

Others have emphasized a potential computational role for serotonin in negative reinforcement (NoGo) learning, via opponent processes with dopamine (Daw et al., 2002). We have argued that while there is some evidence to support this assertion, low levels of DA (or transient DA dips) are still necessary for BG-mediated NoGo learning, and that 5-HT effects may be more likely mediated in prefrontal cortex rather than BG (e.g., Clarke, Dalley, Crofts, Robbins, and Roberts (2004) and Frank and Claus (2006)). According to our modeling framework, 5-HT modulation of prefrontal working memory representations of recent punishments would impact the extent to which behavioral modifications are made on a trial-to-trial basis in response to negative feedback. Recent evidence supports this conclusion, showing that blockade of

5-HT reuptake in humans increases trial-to-trial punishment sensitivity (Chamberlain et al., 2006). Furthermore, 5-HT could also indirectly affect BG-dependent NoGo learning by lowering DA release, via inhibitory 5-HT receptors onto DA neurons (Nocjar, Roth, & Pehek, 2002), which could effectively lead to longer pauses in DA firing during negative reinforcement.

4.5.3. Response inhibition

Although we have emphasized the dynamic Global NoGo function of the STN, we have not discussed its potential role in explicit response inhibition (i.e., canceling a planned motor command). In the present simulations, cortical response conflict drives a Global NoGo signal that slows responding. It is plausible that activation in right inferior frontal gyrus, which is actively involved in inhibiting a response, can do so by exciting the STN and preventing responding altogether (Aron & Poldrack, 2006). Future simulations are required to formally investigate the dynamics of BG and cortical areas that contribute to successful and failed response inhibition, and their neurochemical modulation.

4.5.4. Frontal reward representations in decision making

The current BG model cannot address other more advanced aspects of human decision making, which would require the inclusion of other frontal brain regions. For example, in other simulations we have shown that the ventromedial and orbitofrontal cortices (OFC) are critical for learning to make decisions that depend on accurate estimates of the expected values of decisions. This is because in addition to the BG's specialization in slowly integrating the relative *probabilities* of reinforcement of alternative decisions, the OFC is specialized to also incorporate graded differences in the relative *magnitudes* of potential reinforcement, together with probabilities integrated on a more recent timescale (Frank & Claus, 2006). In these cases, active OFC working memory representations encode the magnitudes of recent positive and negative reinforcement experiences and may be critical for over-riding the prepotent frequency associations that the BG system is particularly specialized to extract, especially when contingencies change unexpectedly. Future work will examine the role of medial OFC projections directly to the STN (Maurice et al., 1998). It is possible that in addition to receiving information about the level of response conflict from premotor/cingulate areas, the STN may also receive OFC information about the expected magnitude of reward outcomes from each of these response alternatives. This functionality would provide a neurobiological mechanism for response threshold adaptation by reward rate (see Simen et al. (in this issue)), and could also explain the enhanced choice of delayed over immediate rewards in STN-lesioned rats (described above) (Winstanley et al., 2005). Finally, in addition to these OFC functions, other BG-PFC circuits may be critical for more complex human decision making tasks that require explicit planning and computations of if–then scenarios (Frank et al., 2001; Hazy, Frank, & O'Reilly, 2006; Houk & Wise, 1995; O'Reilly & Frank, 2006).

5. Conclusion

How do the present simulations provide insight into the problem of when the subthalamic nucleus is beneficial for cognition, compared with situations in which too much STN activity may impair cognitive function? A preliminary answer to this question may be that the STN is useful in situations that would otherwise lead to "jumping the gun" on decision making processes, by preventing premature choices. However, when excessive hesitancy is experienced, the present model would suggest turning off your STN, especially when adequate information is not available to indicate which choice is better. Future computational work may help us better understand both the therapeutic and deleterious effects of STN stimulation on motor and cognitive processes in Parkinson's disease and related disorders.

Acknowledgements

I thank Randy O'Reilly, Adam Aron, Patrick Simen, Todd Braver and Jonathan Cohen for helpful discussion.

Appendix A

The model can be obtained by emailing the author at mfrank@u.arizona.edu. For animated video captures of model dynamics during response selection and learning, please see http://www.u.arizona.edu/~mfrank/BGmodel_movies.html.

A.1. Implementational details

The model is implemented using the Leabra framework (O'Reilly, 2001; O'Reilly & Munakata, 2000). For model parameters, see Table 1. Leabra uses point neurons with excitatory, inhibitory, and leak conductances contributing to an integrated membrane potential, which is then thresholded and transformed via an $x/(x+1)$ sigmoidal function to produce a rate code output communicated to other units (discrete spiking can also be used, but produces noisier results). Each layer uses a k-Winners-Take-All (kWTA) function that computes an inhibitory conductance that keeps roughly the k most active units above firing threshold and keeps the rest below threshold.

The membrane potential V_m is updated as a function of ionic conductances g with reversal (driving) potentials E as follows:

$$\Delta V_m(t) = \tau \sum_c g_c(t)\overline{g}_c(E_c - V_m(t)) \quad (1)$$

with three channels (c) corresponding to: e excitatory input; l leak current; and i inhibitory input. Following electrophysiological convention, the overall conductance is decomposed into a time-varying component $g_c(t)$ computed as a function of the dynamic state of the network, and a constant $\overline{g}_c$ that controls the relative influence of the different conductances. The equilibrium potential can be written in a simplified form by setting the excitatory driving potential (E_e) to 1 and the leak and inhibitory driving potentials (E_l and E_i) of 0:

$$V_m^{\infty} = \frac{g_e\overline{g}_e}{g_e\overline{g}_e + g_l\overline{g}_l + g_i\overline{g}_i} \quad (2)$$

Table 1
BG model parameters

Param	Value	Param	Value	Param	Value	Param	Value	Param	Value	Param	Value
E_l	0.15	$\overline{g}_l$	0.10	E_i	0.15	$\overline{g}_i$	1.0	E_e	1.00	$\overline{g}_e$	1.0
V_{rest}	0.15	Θ	0.25	γ	600	k_{hebb}	0.01	ϵ	0.001		
Striatum ($k = 4$)		$\overline{g}_l$	1.0*	Θ, +DA	0.32*	γ	2500*	γ, +DA	10000*	γ, −DA	300*
GPi		E_l	0.28*	$\overline{g}_l$	3.0*	V_{rest}	0.26*				
GPe		E_l	0.26*	$\overline{g}_l$	1.0*	$\overline{g}_i$	2.5*	V_{rest}	0.26*		
STN		E_l	0.2*	$\overline{g}_l$	1.0*	V_{rest}	0.25*				
Thal		$\overline{g}_i$	1.7*	$\overline{g}_e$	0.5*						
Premotor ($k = 3$)		ϵ	1e−5*	k_{hebb}	1*	V_m noise	$\mu = 0.0015$	V_m noise	$\sigma = 0.0015$		

The first two rows indicate standard default parameters used in 100's of simulations with Leabra software; these parameters are used in the model except where noted with an * for specialized functions of the BG layers. Striatal units have a higher firing threshold θ and higher gain γ during DA bursts ("+DA"), and lower γ during DA dips, to simulate contrast enhancement and reduction (Frank, 2005a). GP and STN units have higher than normal E_l, $\overline{g}_l$ and V_{rest}, leading to tonic baseline activity in the absence of synaptic input, and GPe units have high maximal inhibitory currents $\overline{g}_i$. Thal units have high $\overline{g}_i$ and low $\overline{g}_e$, enabling a default strong inhibition from BG output and only allowing top-down excitatory activity if disinhibited, thereby serving a gating function. Premotor units have Gaussian noise added to the membrane potential, learn with a slow learning rate via purely Hebbian learning. k (kWTA) parameters are shown for striatum and premotor areas, which have within-layer lateral inhibition.

which shows that the neuron is computing a balance between excitation and the opposing forces of leak and inhibition. This equilibrium form of the equation can be understood in terms of a Bayesian decision making framework (O'Reilly & Munakata, 2000).

The excitatory net input/conductance $g_e(t)$ or η_j is computed as the proportion of open excitatory channels as a function of sending activations times the weight values:

$$\eta_j = g_e(t) = \langle x_i w_{ij} \rangle = \frac{1}{n} \sum_i x_i w_{ij}. \tag{3}$$

The inhibitory conductance is computed via the kWTA function described in the next section, and leak is a constant.

Activation communicated to other cells (y_j) is a thresholded (Θ) sigmoidal function of the membrane potential with gain parameter γ:

$$y_j(t) = \frac{1}{\left(1 + \frac{1}{\gamma[V_m(t) - \Theta]_+}\right)} \tag{4}$$

where $[x]_+$ is a threshold function that returns 0 if $x < 0$ and x if $X > 0$. Note that if it returns 0, we assume $y_j(t) = 0$, to avoid dividing by 0. As it is, this function has a very sharp threshold, which interferes with graded learning mechanisms (e.g., gradient descent). To produce a less discontinuous deterministic function with a softer threshold, the function is convolved with a Gaussian noise kernel ($\mu = 0$, $\sigma = 0.005$), which reflects the intrinsic processing noise of biological neurons:

$$y_j^*(x) = \int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}\sigma} \mathrm{e}^{-z^2/(2\sigma^2)} y_j(z - x) \mathrm{d}z \tag{5}$$

where x represents the $[V_m(t) - \Theta]_+$ value, and $y_j^*(x)$ is the noise-convolved activation for that value.

A.2. Inhibition within and between layers

Inhibition *between* layers (i.e. for GABAergic projections between BG layers) is achieved via simple unit inhibition, where the inhibitory current g_i for the unit is determined from the net input of the sending unit.

For *within* layer lateral inhibition (used in Striatum and premotor cortex), Leabra uses a kWTA (k-Winners-Take-All) function to achieve inhibitory competition among units within each layer (area). The kWTA function computes a uniform level of inhibitory current for all units in the layer, such that the $k + 1$th most excited unit within a layer is generally below its firing threshold, while the kth is typically above threshold. Activation dynamics similar to those produced by the kWTA function have been shown to result from simulated inhibitory interneurons that project both feedforward and feedback inhibition (O'Reilly & Munakata, 2000). Thus, although the kWTA function is somewhat biologically implausible in its implementation (e.g., requiring global information about activation states and using sorting mechanisms), it provides a computationally effective approximation to biologically plausible inhibitory dynamics.

kWTA is computed via a uniform level of inhibitory current for all units in the layer as follows:

$$g_i = g_{k+1}^{\Theta} + q(g_k^{\Theta} - g_{k+1}^{\Theta}) \tag{6}$$

where $0 < q < 1$ (0.25 default used here) is a parameter for setting the inhibition between the upper bound of g_k^{Θ} and the lower bound of g_{k+1}^{Θ}. These boundary inhibition values are computed as a function of the level of inhibition necessary to keep a unit right at threshold:

$$g_i^{\Theta} = \frac{g_e^* \bar{g}_e (E_e - \Theta) + g_l \bar{g}_l (E_l - \Theta)}{\Theta - E_i} \tag{7}$$

where g_e^* is the excitatory net input.

Two versions of kWTA functions are typically used in Leabra. In the kWTA function used in the Striatum, g_k^{Θ} and g_{k+1}^{Θ} are set to the threshold inhibition value for the kth and $k + 1$th most excited units, respectively. Thus, the inhibition is placed to allow k units to be above threshold, and the remainder below threshold.

The premotor cortex uses the *average-based* kWTA version, g_k^{Θ} is the average g_i^{Θ} value for the top k most excited units, and

g^{Θ}_{k+1} is the average of g^{Θ}_{i} for the remaining $n-k$ units. This version allows for more flexibility in the actual number of units active depending on the nature of the activation distribution in the layer and the value of the q parameter (which is set to default value of 0.6). This flexibility is necessary for the premotor units to have differential levels of activity during settling (depending on whether or not a single response has been facilitated), and also allows greater activity in high-conflict trials.

A.3. Learning

Synaptic connection weights were trained using a reinforcement learning version of Leabra. The learning algorithm involves two phases, and is more biologically plausible than standard error backpropagation. In the *minus phase*, the network settles into activity states based on input stimuli and its synaptic weights, ultimately "choosing" a response. In the *plus phase*, the network resettles in the same manner, with the only difference being a change in simulated dopamine: an increase of SNc unit firing from 0.5 to 1.0 for correct responses, and a decrease to zero SNc firing for incorrect responses (Frank, 2005a).

For learning, Leabra uses a combination of error-driven and Hebbian learning. The error-driven component is the symmetric midpoint version of the GeneRec algorithm (O'Reilly, 1996), which is functionally equivalent to the deterministic Boltzmann machine and contrastive Hebbian learning (CHL), computing a simple difference of a pre and postsynaptic activation product across these two phases. For Hebbian learning, Leabra uses essentially the same learning rule used in competitive learning or mixtures-of-Gaussians which can be seen as a variant of the Oja normalization (Oja, 1982). The error-driven and Hebbian learning components are combined additively at each connection to produce a net weight change.

The equation for the Hebbian weight change is:

$$\Delta_{\text{hebb}} w_{ij} = x_i^+ y_j^+ - y_j^+ w_{ij} = y_j^+ (x_i^+ - w_{ij}) \tag{8}$$

and for error-driven learning using CHL:

$$\Delta_{\text{err}} w_{ij} = (x_i^+ y_j^+) - (x_i^- y_j^-) \tag{9}$$

which is subject to a soft-weight bounding to keep within the 0–1 range:

$$\Delta_{\text{sberr}} w_{ij} = [\Delta_{\text{err}}]_+ (1 - w_{ij}) + [\Delta_{\text{err}}]_- w_{ij}. \tag{10}$$

The two terms are then combined additively with a normalized mixing constant k_{hebb}:

$$\Delta w_{ij} = \epsilon [k_{\text{hebb}} (\Delta_{\text{hebb}}) + (1 - k_{\text{hebb}})(\Delta_{\text{sberr}})]. \tag{11}$$

References

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., & Crutcher, M. D. (1990a). Functional architecture of basal ganglia circuits: Neural substrates of parallel processing. *Trends in Neuroscience*, *13*, 266–271.

Alexander, G. E., & Crutcher, M. D. (1990b). Preparation for movement: Neural representations of intended direction in three motor areas of the monkey. *Journal of Neurophysiology*, *64*, 133–150.

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.

Amirnovin, R., Williams, Z. M., Cosgrove, G. R., & Eskandar, E. N. (2004). Visually guided movements suppress subthalamic oscillations in Parkinson's disease patients. *Journal of Neuroscience*, *24*(50), 11302–11306.

Aron, A. R., & Poldrack, R. A. (2006). Cortical and subcortical contributions to stop signal response inhibition: Role of the subthalamic nucleus. *Journal of Neuroscience*, *26*, 2424–2433.

Ashby, F. G., Noble, S., Ell, S. W., Filoteo, J. V., & Waldron, E. M. (2003). Category learning deficits in Parkinson's disease. *Neuropsychology*, *17*, 133–142.

Aston-Jones, G., & Cohen, J. D. (2005). An integrative theory of locus coeruleus–norepinephrine function: Adaptive gain and optimal performance. *Annual Review of Neuroscience*, *28*, 403–450.

Aubert, I., Ghorayeb, I., Normand, E., & Bloch, B. (2000). Phenotypical characterization of the neurons expressing the D1 and D2 dopamine receptors in the monkey striatum. *Journal of Comparative Neurology*, *418*, 22–32.

Bar-Gad, I., Morris, G., & Bergman, H. (2003). Information processing, dimensionality reduction and reinforcement learning in the basal ganglia. *Progress in Neurobiology*, *71*, 439–473.

Basso, M. A., & Wurtz, R. H. (2002). Neuronal activity in substantia nigra pars reticulata during target selection. *Journal of Neuroscience*, *22*(5), 1883–1894.

Baunez, C., Humby, T., Eagle, D. M., Ryan, L. J., Dunnett, S. B., & Robbins, T. W. (2001). Effects of STN lesions on simple vs choice reaction time tasks in the rat: Preserved motor readiness, but impaired response selection. *European Journal of Neuroscience*, *13*, 1609–1616.

Baunez, C., & Robbins, T. W. (1997). Bilateral lesions of the subthalamic nucleus induce multiple deficits in an attentional task in rats. *European Journal of Neuroscience*, *9*(10), 2086–2099.

Bayer, H. M., & Glimcher, P. W. (2005). Midbrain dopamine neurons encode a quantitative reward prediction error signal. *Neuron*, *47*(1), 129–141.

Beiser, D. G., & Houk, J. C. (1998). Model of cortical-basal ganglionic processing: Encoding the serial order of sensory events. *Journal of Neurophysiology*, *79*, 3168–3188.

Benazzouz, A., & Hallett, M. (2000). Mechanism of action of deep brain stimulation. *Neurology*, *55*, S13–S16.

Bergman, H., Feingold, A., Nini, A., Raz, A., Slovin, H., Abeles, M., et al. (1998). Physiological aspects of information processing in the basal ganglia of normal and Parkinsonian primates. *Trends in Neurosciences*, *21*, 32–38.

Bergman, H., Wichmann, T., & DeLong, M. R. (1990). Reversal of experimental Parkinsonism by lesions of the subthalamic nucleus. *Science*, *249*, 1436–1438.

Bergman, H., Wichmann, T., Karmon, B., & DeLong, M. R. (1994). The primate subthalamic nucleus. II. Neuronal activity in the MPTP model of Parkinsonism. *Journal of Neurophysiology*, *72*, 507–520.

Bevan, M. D., Magill, P. J., Terman, D., Bolam, J. P., & Wilson, C. J. (2002). Move to the rhythm: Oscillations in the subthalamic nucleus–external globus pallidus network. *Trends in Neurosciences*, *25*, 525–531.

Bogacz, R., Brown, E., Moehlis, J., Holmes, P., & Cohen, J. D. The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced choice tasks. *Psychological Review* (in press).

Bokura, H., Yamaguchi, S., & Kobayashi, S. (2005). Event-related potentials for response inhibition in Parkinson's disease. *Neuropsychologia*, *43*, 967–975.

Boraud, T., Bezard, E., Bioulac, B., & Gross, E. (2002). From single extracellular unit recording in experimental and human Parkinsonism to the development of a functional concept of the role played by the basal ganglia in motor control. *Progress in Neurobiology*, *66*(4), 265–283.

Botvinick, M. M., Braver, T. S., Barch, D. M., Carter, C. S., & Cohen, J. D. (2001). Conflict monitoring and cognitive control. *Psychological Review*, *108*, 624–652.

Brown, E., Gao, J., Holmes, P., Bogacz, R., Gilzenrat, M., & Cohen, J. D. (2005). Simple neural networks that optimize decisions. *International Journal of Bifurcation and Chaos*, *15*, 803–826.

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, J. W., Bullock, D., & Grossberg, S. (2004). How laminar frontal cortex and basal ganglia circuits interact to control planned and reactive saccades. *Neural Networks*, *17*, 471–510.

Brown, R. G., & Marsden, C. D. (1988). Internal versus external cues and the control of attention in Parkinson's disease. *Brain*, *111*, 323–345.

Chamberlain, S. R., Muller, U., Blackwell, A. D., Clark, L., Robbins, T. W., & Sahakian, B. J. (2006). Neurochemical modulation of response inhibition and probabilistic learning in humans. *Science*, *311*(5762), 861–863.

Charbonneau, D., Riopelle, R. J., & Beninger, R. J. (1996). Impaired incentive learning in treated Parkinson's disease. *Canadian Journal of Neurological Sciences*, *23*(4), 271–278.

Choi, W. Y., Balsam, P. D., & Horvitz, J. C. (2005). Extended habit training reduces dopamine mediation of appetitive response expression. *Journal of Neuroscience*, *25*(29), 6729–6733.

Clarke, H. F., Dalley, J. W., Crofts, H. S., Robbins, T. W., & Roberts, A. C. (2004). Cognitive inflexibility after prefrontal serotonin depletion. *Science*, *304*(5672), 878–880.

Cools, R. (2005). Dopaminergic modulation of cognitive function — implications for L-DOPA treatment in Parkinson's disease. *Neuroscience and Biobehavioral Reviews*, 1–23.

Cools, R., Barker, R. A., Sahakian, B. J., & Robbins, T. W. (2001a). Enhanced or impaired cognitive function in Parkinson's disease as a function of dopaminergic medication and task demands. *Cerebral Cortex*, *11*, 1136–1143.

Cools, R., Barker, R. A., Sahakian, B. J., & Robbins, T. W. (2001b). Mechanisms of cognitive set flexibility in Parkinson's disease. *Brain*, *124*, 2503–2512.

Cools, R., Barker, R. A., Sahakian, B. J., & Robbins, T. W. (2003). L-Dopa medication remediates cognitive inflexibility, but increases impulsivity in patients with Parkinson's disease. *Neuropsychologia*, *41*, 1431–1441.

Crutcher, M. D., & Alexander, G. E. (1990). Movement-related neuronal activity selectively coding either direction or muscle pattern in three motor areas of the monkey. *Journal of Neurophysiology*, *64*, 151–163.

Czernecki, V., Pillon, B., Houeto, J. L., Pochon, J. B., Levy, R., & Dubois, B. (2002). Motivation, reward, and Parkinson's disease: Influence of dopatherapy. *Neuropsychologia*, *40*, 2257–2267.

Daw, N. D., Kakade, S., & Dayan, P. (2002). Opponent interactions between serotonin and dopamine. *Neural Networks*, *15*, 603–616.

Delgado, M. R., Locke, H. M., Stenger, V. A., & Fiez, J. A. (2003). Dorsal striatum responses to reward and punishment: Effects of valence and magnitude manipulations. *Cognitive, Affective, and Behavioral Neuroscience*, *3*, 27–38.

Delgado, M. R., Miller, M. M., Inati, S., & Phelps, E. A. (2005). An fMRI study of reward-related probability learning. *Neuroimage*, *24*, 862–873.

DeLong, M. R. (1990). Primate models of movement disorders of basal ganglia origin. *Trends in Neurosciences*, *13*, 281–285.

Desbonnet, L., Temel, Y., Visser-Vandewalle, V., Blokland, A., Hornikx, V., & Steinbusch, H. W. (2004). Premature responding following bilateral stimulation of the rat subthalamic nucleus is amplitude and frequency dependent. *Brain Research*, *1008*, 198–204.

Frank, M. J. (2005a). Dynamic dopamine modulation in the basal ganglia: A neurocomputational account of cognitive deficits in medicated and non-medicated Parkinsonism. *Journal of Cognitive Neuroscience*, *17*, 51–72.

Frank, M. J. (2005b) When and when not to use your subthalamic nucleus: Lessons from a computational model of the basal ganglia. In *Modelling natural action selection: Proceedings of an international workshop* (pp. 53–60).

Frank, M. J., & Claus, E. D. (2006). Anatomy of a decision: Striato-orbitofrontal interactions in reinforcement learning, decision making, and reversal. *Psychological Review*, *113*(2), 300–326.

Frank, M. J., Loughry, B., & O'Reilly, R. C. (2001). Interactions between the frontal cortex and basal ganglia in working memory: A computational model. *Cognitive, Affective, and Behavioral Neuroscience*, *1*, 137–160.

Frank, M. J., & O'Reilly, R. C. (2006). A mechanistic account of striatal dopamine function in human cognition: Psychopharmacological studies with cabergoline and haloperidol. *Behavioral Neuroscience*, *120*, 497–517.

Frank, M. J., O'Reilly, R. C., & Curran, T. (2006). When memory fails, intuition reigns: Midazolam enhances implicit inference in humans. *Psychological Science*, *17*, 700–707.

Frank, M. J., Seeberger, L. C., & O'Reilly, R. C. (2004). By carrot or by stick: Cognitive reinforcement learning in Parkinsonism. *Science*, *306*, 1940–1943.

Frank, M. J., Woroch, B. S., & Curran, T. (2005). Error-related negativity predicts reinforcement learning and conflict biases. *Neuron*, *47*, 495–501.

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. (2000). Molecular effects of dopamine on striatal projection pathways. *Trends in Neurosciences*, *23*, S64–S70.

Gerfen, C. R., Keefe, K. A., & Gauda, E. B. (1995). D_1 and D_2 dopamine receptor function in the striatum: Coactivation of D_1- and D_2-dopamine receptors on separate populations of neurons results in potentiated immediate early gene response in D_1-containing neurons. *Journal of Neuroscience*, *15*, 8167–8176.

Gerfen, C. R., & Wilson, C. (1996). The basal ganglia. In L. Swanson, A. Bjorkland, & T. Hokfelt (Eds.), *Integrated systems of the CNS*: *Vol. 12*. *Handbook of chemical neuroanatomy* (pp. 371–468). Amsterdam: Elsevier.

Gold, J. I., & Shadlen, M. N. (2002). Banburismus and the brain: Decoding the relationship between sensory stimuli, decisions, and reward. *Neuron*, *36*, 299–308.

Gotham, A. M., Brown, R. G., & Marsden, C. D. (1988). 'Frontal' cognitive function in patients with Parkinson's disease 'on' and 'off' levodopa. *Brain*, *111*, 299–321.

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.

Gurney, K. N., Humphries, M., Wood, R., Prescott, T. J., & Redgrave, P. (2004). Testing computational hypotheses of brain systems function: A case study with the basal ganglia. *Network*, *15*(4), 263–290.

Harley, T. (2004). Does cognitive neuropsychology have a future? *Cognitive Neuropsychology*, *21*(1), 3–16.

Hazy, T. E., Frank, M. J., & O'Reilly, R. C. (2006). Banishing the homunculus: Making working memory work. *Neuroscience*, *139*, 105–118.

Hebb, D. O. (1949). *The organization of behavior*. New York: Wiley.

Hernandez-Lopez, S., Bargas, J., Surmeier, D. J., Reyes, A., & Galarraga, E. (1997). D1 receptor activation enhances evoked discharge in neostriatal medium spiny neurons by modulating an L-type Ca^{2+} conductance. *Journal of Neuroscience*, *17*, 3334–3342.

Hernandez-Lopez, S., Tkatch, T., Perez-Garci, E., Galarraga, E., Bargas, J., Hamm, H., et al. (2000). D2 dopamine receptors in striatal medium spiny neurons reduce L-type Ca^{2+} currents and excitability via a novel PLCβ1-IP$_3$-calcineurin-signaling cascade. *Journal of Neuroscience*, *20*, 8987–8995.

Hikosaka, O. (1994). Role of basal ganglia in control of innate movements, learned behaviour and cognition. In G. Percheron, J. McKenzie, & J. Feger (Eds.), *The basal ganglia iv: New ideas and data on structure and function* (pp. 589–596). New York: Plenum.

Holroyd, C. B., & Coles, M. G. H. (2002). The neural basis of human error processing: Reinforcement learning, dopamine, and the error-related negativity. *Psychological Review*, *109*, 679–709.

Houk, J. C. (2005). Agents of the mind. *Biology Cybernetica*, *92*(6), 427–437.

Houk, J. C., Adams, J. L., & Barto, A. G. (1995). A model of how the basal ganglia generate and use neural signals that predict reinforcement. In J. C. Houk, J. L. Davis, & D. G. Beiser (Eds.), *Models of information processing in the basal ganglia* (pp. 233–248). Cambridge, MA: MIT Press.

Houk, J. C., & Wise, S. P. (1995). Distributed modular architectures linking basal ganglia, cerebellum, and cerebral cortex: Their role in planning and controlling action. *Cerebral Cortex*, *5*, 95–110.

Jackson, G. M., Jackson, S., 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.

Jiang, H., Stein, B. E., & McHaffie, J. G. (2003). Opposing basal ganglia processes shape midbrain visuomotor activity bilaterally. *Nature*, *419*, 982–985.

Joel, D., & Weiner, I (1999). Striatal contention scheduling and the split circuit scheme of basal ganglia-thalamocortical circuitry: From anatomy to behaviour. In R. Miller, & J. R. Wickens (Eds.), *Conceptual advances in brain research: Brain dynamics and the striatal complex* (pp. 209–236). Harwood Academic Publishers.

Karachi, C., Yelnik, J., Tande, D., Tremblay, L., Hirsch, E. C., & Francois, C. (2004). The pallidosubthalamic projection: An anatomical substrate for nonmotor functions of the subthalamic nucleus in primates. *Movement disorders*.

Kawaguchi, Y., Wilson, C. J., & Emson, P. C. (1990). Projection subtypes of rat neostriatal matrix cells revealed by intracellular injection of biocytin. *Journal of Neuroscience*, *10*(10), 3421–3438.

Kish, S. J., Shannak, K., & Hornykiewicz, O. (1988). Uneven pattern of dopamine loss in the striatum of patients with idiopathic Parkinson's disease. *New England Journal of Medecine*, *318*, 876–880.

Knowlton, B. J., Mangels, J. A., & Squire, L. R. (1996). A neostriatal habit learning system in humans. *Science*, *273*, 1399.

Kolomiets, B. P., Deniau, J. M., Mailly, P., Menetrey, A., Glowinski, J., & Thierry, A. M. (2001). Segregation and convergence of information flow through the cortico-subthalamic pathways. *Journal of Neuroscience*, *21*(15), 5764–5772.

Lavoie, B., & Parent, A. (1990). Immunohistochemical study of the serotoninergic innervation of the basal ganglia in the squirrel monkey. *Journal of Comparative Neurology*, *299*(1), 1–16.

Lei, W., Jiao, Y., Del Mar, N., & Reiner, A. (2004). Evidence for differential cortical input to direct pathway versus indirect pathway striatal projection neurons in rats. *Journal of Neuroscience*, *24*(38), 8289–8299.

Levesque, M., & Parent, A. (2005). The striatofugal fiber system in primates: A reevaluation of its organization based on single-axon tracing studies. *Proceedings of The National Academy of Sciences of the United States of America*.

Levy, R., Hutchison, W. D., Lozano, A. M., & Dostrovsky, J. O. (2000). High-frequency synchronization of neuronal activity in the subthalamic nucleus of Parkinsonian patients with limb tremor. *Journal of Neuroscience*, *20*, 7766.

Maddox, W. T., & Filoteo, J. V. (2001). Striatal contributions to category learning: Quantitative modeling of simple linear and complex non-linear rule learning in patients with Parkinson's disease. *Journal of the International Neuropsychological Society*, *7*, 710–727.

Magill, P. J., Bolam, J. P., & Bevan, M. D. (2001). Dopamine regulates the impact of the cerebral cortex on the subthalamic nucleus–globus pallidus network. *Neuroscience*, *106*, 313–330.

Magill, P. J., Sharott, A., Bevan, M. D., Brown, P., & Bolam, J. P. (2004). Synchronous unit activity and local field potentials evoked in the subthalamic nucleus by cortical stimulation. *Journal of Neurophysiology*, *92*(2), 700–714.

Mahon, S., Casassus, G., Mulle, C., & Charpier, S. (2003). Spike-dependent intrinsic plasticity increases firing probability in rat striatal neurons in vivo. *Journal of Physiology*, *550*(Pt 3), 947–959.

Maurice, N., Deniau, J. -M., Glowinski, J., & Thierry, A. -M. (1998). Relationships between the prefrontal cortex and the basal ganglia in the rat: Physiology of the corticosubthalamic circuits. *Journal of Neuroscience*, *18*, 9539.

McAuley, J. H. (2003). The physiological basis of clinical deficits in Parkinson's disease. *Progress in Neurobiology*, *69*, 27–48.

Mehta, M. A., Swainson, R., Ogilvie, A. D., Sahakian, B. J., & Robbins, T. W. (2001). Improved short-term spatial memory but impaired reversal learning following the dopamine D2 agonist bromocriptine in human volunteers. *Psychopharmacology*, *159*, 10–20.

Meissner, W., Leblois, A., Hansel, D., Bioulac, B., Gross, C. E., Benazzouz, A., et al. (2005). Subthalamic high frequency stimulation resets subthalamic firing and reduces abnormal oscillations. *Brain*, *128*(Pt 10), 2372–2382.

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.

Middleton, F. A., & Strick, P. L. (2002). Basal-ganglia 'projections' to the prefrontal cortex of the primate. *Cerebral Cortex*, *12*, 926–935.

Miller, W., & Delong, M. R. (1987). Altered tonic activity of neurons in the globus pallidus and subthalamic nucleus in the primate MPTP model of Parkinsonism. In M. B. Carpenter, & A. Jayaraman (Eds.), *The basal ganglia*: *Vol. II* (pp. 415–427). New York: Plenum Press.

Mink, J. W. (1996). The basal ganglia: Focused selection and inhibition of competing motor programs. *Progress in Neurobiology*, *50*, 381–425.

Nambu, A., Kaneda, K., Tokuno, H., & Takada, M. (2002). Organization of corticostriatal motor inputs in monkey putamen. *Journal of Neurophysiology*, *88*, 1830–1842.

Nambu, A., Tokuno, H., Hamada, I., Kita, H., Imanishi, M., Akazawa, T., et al. (2000). Excitatory cortical inputs to pallidal neurons via the subthalamic nucleus in the monkey. *Journal of Neurophysiology*, *84*, 289–300.

Ni, Z., Bouali-Benazzouz, R., Gao, D., Benabid, A., & Benazzouz, A. (2000). Changes in the firing pattern of globus pallidus neurons after the degeneration of nigrostriatal pathway are mediated by the subthalamic nucleus in rat. *European Journal of Neuroscience*, *12*, 4338–4344.

Nocjar, C., Roth, B. L., & Pehek, E. A. (2002). Localization of 5-HT(2A) receptors on dopamine cells in subnuclei of the midbrain A10 cell group. *Neuroscience*, *111*, 163–176.

Oja, E. (1982). A simplified neuron model as a principal component analyzer. *Journal of Mathematical Biology*, *15*, 267–273.

O'Reilly, R. C. (1996). Biologically plausible error-driven learning using local activation differences: The generalized recirculation algorithm. *Neural Computation*, *8*(5), 895–938.

O'Reilly, R. C. (2001). Generalization in interactive networks: The benefits of inhibitory competition and Hebbian learning. *Neural Computation*, *13*, 1199–1242.

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., Frank, M. J., Hazy, T. E., & Watz, B. PVLV: The primary value and learned value Pavlovian learning algorithm. *Behavioral Neuroscience* (in press).

O'Reilly, R. C., & Munakata, Y. (2000). *Computational explorations in cognitive neuroscience: Understanding the mind by simulating the brain*. Cambridge, MA: MIT Press.

Orieux, G., Francois, C., Feger, J., & Hirsch, E. C. (2002). Consequences of dopaminergic denervation on the metabolic activity of the cortical neurons projecting to the subthalamic nucleus in the rat. *Journal of Neuroscience*, *22*, 8762–8770.

Pan, W. -X., Schmidt, R., Wickens, J. R., & Hyland, B. I. (2005). Dopamine cells respond to predicted events during classical conditioning: Evidence for eligibility traces in the reward-learning network. *Journal of Neuroscience*, *25*(26), 6235–6242.

Parent, A., & Hazrati, L. (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.

Parkinson, J. A., Dalley, J. W., Cardinal, R. N., Bamford, A., Fehnert, B., Lachenal, G., et al. (2002). Nucleus accumbens dopamine depletion impairs both acquisition and performance of appetitive Pavlovian approach behaviour: Implications for mesoaccumbens dopamine function. *Behavioral Brain Research*, *137*, 149–163.

Pasupathy, A., & Miller, E. K. (2005). Different time courses for learning-related activity in the prefrontal cortex and striatum. *Nature*, *433*, 873–876.

Picard, N., & Strick, P. L. (1996). Motor areas of the medial wall: A review of their location and functional activation. *Cerebral Cortex*, *6*(3), 342–353.

Picard, N., & Strick, P. L. (2001). Imaging the premotor areas. *Current Opinion Neurobiology*, *11*(6), 663–672.

Pompeiano, M., Palacios, J. M., & Mengod, G. (1994). Distribution of the serotonin 5-HT2 receptor family mRNAs: Comparison between 5-HT2A and 5-HT2C receptors. *Brain Research. Molecular Brain Research*, *23*(1–2), 163–178.

Pothos, E. N., Davila, V., & Sulzer, D. (1998). Presynaptic recording of quanta from midbrain dopamine neurons and modulation of the quantal size. *Journal of Neuroscience*, *18*(11), 4106–4118.

Ratcliff, R., Van Zandt, T., & McKoon, G. (1999). Connectionist and diffusion models of reaction time. *Psychological Review*, *106*, 261.

Raz, A., Vaadia, E., & Bergman, H. (2000). Firing patterns and correlations of spontaneous discharge of pallidal neurons in the normal and the tremulous 1-methyl-4-phenyl-1,2,3,6-tetrahydropyridine vervet model of Parkinsonism. *Journal of Neuroscience*, *20*, 8559–8571.

Redgrave, P., Prescott, T. J., & Gurney, K. (1999). The basal ganglia: A vertebrate solution to the selection problem? *Neuroscience*, *89*(1009).

Ridley, R. M., Haystead, T. A., & Baker, H. F. (1981). An analysis of visual object reversal learning in the marmoset after amphetamine and haloperidol. *Pharmacology Biochemistry and Behaviour*, *14*(3), 345–351.

Robertson, G. S., Vincent, S. R., & Fibiger, H. C. (1992). D_1 and D_2 dopamine receptors differentially regulate c-fos expression in striatonigral and stiratopallidal neurons. *Neuroscience*, *49*, 285–296.

Rogers, R. D., Sahakian, B. J., Hodges, J. R., Polkey, C. E., Kennard, C., & Robbins, T. W. (1998). Dissociating executive mechanisms of task control following frontal damage and Parkinson's disease. *Brain*, *121*, 815–842.

Rubchinsky, L. L., Kopell, N., & Sigvardt, K. A. (2003). Modeling facilitation and inhibition of competing motor programs in basal ganglia subthalamic nucleus–pallidal circuits. *Proceedings of the National Academy of Sciences*, *100*, 14427–14432.

Salin, P., Hajji, M. D., & Kerkerian-Le Goff, L. (1996). Bilateral 6-hydroxydopamine-induced lesion of the nigrostriatal dopamine pathway reproduces the effects of unilateral lesion on substance P but not on enkephalin expression in rat basal ganglia. *European Journal of Neuroscience*, *8*, 1746–1757.

Samejima, K., Ueda, Y., Doya, K., & Kimura, M. (2005). Representation of action-specific reward values in the striatum. *Science*, *310*, 1337–1340.

Sato, F., Parent, M., Levesque, M., & Parent, A. (2000). Axonal branching pattern of neurons of the subthalamic nucleus in primates. *Journal of Comparative Neurology*, *424*, 142–152.

Satoh, T., Nakai, S., Sato, T., & Kimura, M. (2003). Correlated coding of motivation and outcome of decision by dopamine neurons. *Journal of Neuroscience*, *23*, 9913–9923.

Schall, J. D. (2003). Neural correlates of decision processes: Neural and mental chronometry. *Current Opinion in Neurobiology*, *13*, 182–186.

Schultz, W. (2002). Getting formal with dopamine and reward. *Neuron*, *36*, 241–263.

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

Shohamy, D., Myers, C. E., Geghman, K. D., Sage, J., & Gluck, M. A. (2006). L-Dopa impairs learning, but spares generalization, in Parkinson's disease. *Neuropsychologia*, *44*(5), 774–784.

Shohamy, D., Myers, C. E., Grossman, S., Sage, J., & Gluck, M. A. (2005). The role of dopamine in cognitive sequence learning: Evidence from Parkinson's disease. *Behavioural Brain Research*, *156*(2), 191–199.

Simen, P., Cohen, J., & Holmes, P. Rapid decision threshold modulation by reward rate in a neural network. *Neural Networks* (in this issue). doi: 10.1016/j.neunet.2006.05.038.

Smith, A. G., Neill, J. C., & Costall, B. (1999). The dopamine D2/D3 receptor agonist 7-OH-DPAT induces cognitive impairment in the marmoset. *Pharmacology, Biochemistry and Behavior*, *63*(2), 201–211.

Smith-Roe, S. L., & Kelley, A. E. (2000). Coincident activation of NMDA and dopamine D1 receptors within the nucleus accumbens core is required for appetitive instrumental learning. *Journal of Neuroscience*, *22*, 7737–7742.

Stanford, I. M., Kantaria, M. A., Chahal, H. S., Loucif, K. C., & Wilson, C. L. (2005). 5-Hydroxytryptamine induced excitation and inhibition in the subthalamic nucleus: Action at 5-HT(2C), 5-HT(4) and 5-HT(1A) receptors. *Neuropharmacology*.

Sutton, R. S. (1988). Learning to predict by the method of temporal differences. *Machine Learning*, *3*, 9–44.

Swainson, R., Rogers, R. D., Sahakian, B. J., Summers, B. A., Polkey, C. E., & Robbins, T. W. (2000). Probabilistic learning and reversal deficits in patients with Parkinson's disease or frontal or temporal lobe lesions: Possible adverse effects of dopaminergic medication. *Neuropsychologia*, *38*(596).

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.

Usher, M., & McClelland, J. L. (2001). The time course of perceptual choice: The leaky, competing accumulator model. *Psychological Review*, *108*, 550–592.

Walderhaug, E., Lunde, H., Nordvik, J. E., Landro, N. I., Refsum, H., & Magnusson, A. (2002). Lowering of serotonin by rapid tryptophan depletion increases impulsiveness in normal individuals. *Psychopharmacology*, *164*, 385–391.

Wichmann, T., Bergman, H., & DeLong, M. R. (1994). The primate subthalamic nucleus. I. Functional properties in intact animals. *Journal of Neurophysiology*, *72*, 494–506.

Winstanley, C. A., Baunez, C., Theobald, D. E. H., & Robbins, T. W. (2005). Lesions to the subthalamic nucleus decrease impulsive choice but impair autoshaping in rats: The importance of the basal ganglia in Pavlovian conditioning and impulse control. *European Journal of Neuroscience*, *21*(11), 3107–3116.

Winstanley, C. A., Theobald, D. E. H., Dalley, J. W., Glennon, J. C., & Robbins, T. W. (2004). 5-HT2A and 5-HT2C receptor antagonists have opposing effects on a measure of impulsivity: Interactions with global 5-HT depletion. *Psychopharmacology*, *176*(3–4), 376–385.

Witt, K., Pulkowski, U., Herzog, J., Lorenz, D., Hamel, W., Deuschl, G., et al. (2004). Deep brain stimulation of the subthalamic nucleus improves cognitive flexibility but impairs response inhibition in Parkinson's disease. *Archives of Neurology*, *61*, 697–700.

Wu, Y., Richard, S., & Parent, A. (2000). The organization of the striatal output system: A single-cell juxtacellular labeling study in the rat. *Neuroscience Research*, *38*(1), 49–62.

Xiang, Z., Wang, L., & Kitai, S. T. (2005). Modulation of spontaneous firing in rat subthalamic neurons by 5-ht receptor subtypes. *Journal of Neurophysiology*, *93*(3), 1145–1157.

Yeung, N., Botvinick, M. M., & Cohen, J. D. (2004). The neural basis of error detection: Conflict monitoring and the error-related negativity. *Psychological Review*, *111*(4), 931–959.