

A Computational Model of How the Basal Ganglia Produce Sequences

Gregory S. Berns
University of Pittsburgh

Terrence J. Sejnowski
Howard Hughes Medical Institute, The Salk Institute, and
University of California, San Diego

Abstract

■ We propose a systems-level computational model of the basal ganglia based closely on known anatomy and physiology. First, we assume that the thalamic targets, which relay ascending information to cortical action and planning areas, are tonically inhibited by the basal ganglia. Second, we assume that the output stage of the basal ganglia, the internal segment of the globus pallidus (GPI), selects a single action from several competing actions via lateral interactions. Third, we propose that a form of local working memory exists in the form of reciprocal connections between the external globus pallidus (GPe) and the subthalamic nucleus (STN). As a test of the model, the system was trained to learn a sequence of states that required the context of previous actions. The striatum, which was assumed to represent a conjunction of cortical states, directly selected the action in the GP during training. The STN-to-GP connection strengths were modified by an associative learning

rule and came to encode the sequence after 20 to 40 iterations through the sequence. Subsequently, the system automatically reproduced the sequence when cued to the first action. The behavior of the model was found to be sensitive to the ratio of the striatal-nigral learning rate to the STN-GP learning rate. Additionally, the degree of striatal inhibition of the globus pallidus had a significant influence on both learning and the ability to select an action. Low learning rates, which would be hypothesized to reflect low levels of dopamine, as in Parkinson's disease, led to slow acquisition of contextual information. However, this could be partially offset by modeling a lesion of the globus pallidus that resulted in an increase in the gain of the STN units. The parameter sensitivity of the model is discussed within the framework of existing behavioral and lesion data. ■

INTRODUCTION

The basal ganglia are a collection of subcortical structures that are relatively large in primates, particularly in humans. Although much is now known about both the types of neurons that comprise these structures and their connectivity, relatively little is known about the overall function of the basal ganglia. Lesion studies both in lower primates and in humans consistently point to a role in motor function, yet it is known that several parts of the basal ganglia receive massive projections from the prefrontal cortex, suggesting a role in planning and cognition. In this paper, we study an anatomically motivated computational model that integrates experimental data from the molecular to the behavioral levels and suggests a function for each part of the basal ganglia.

Approximately 80% of both the striato-pallidal and pallido-thalamic neurons are GABAergic, a prevalence of inhibition that is uncommon in the central nervous system, yet no satisfactory explanation exists for this finding. In particular, the direct striato-pallidal-thalamic

pathway is comprised of two GABA neurons in series with each other. A fundamental question is what advantage this arrangement confers over a single excitatory synapse. Findings suggest the maintenance of cortical topography throughout the basal ganglia, which has raised the possibility that parallel streams of information project through the structure, but with relatively little integration being performed (Alexander, DeLong, & Strick, 1986). However, it is also known that the input stage of the basal ganglia, the caudate and putamen (collectively referred to as the striatum), receives inputs from almost the entire cortex. There is also a convergence in neuron number from the striatum to the output stage, the globus pallidus. These two levels of massive convergence suggest that the basal ganglia is involved in integrating many types of information within each stream to either plan or select an action from the many competing possibilities represented in the cortex.

The primary source of dopamine in the brain is found in the substantia nigra pars compacta (SNc) and ventral tegmental area (VTA), both of which have close ties with

the other basal ganglia structures and themselves are often considered part of the basal ganglia. Dopamine has been implicated in reward-driven learning (Schultz, Apicella, & Ljungberg, 1993; Schultz, Apicella, Scarnati, & Ljungberg, 1992), and the VTA is known to be a self-stimulation site (Cooper, Bloom, & Roth, 1996). The role of dopamine seems to be closely related to motor behavior and the need to perform an action in so-called operant tasks, where rewards are contingent upon acting. Yet dopamine has also been implicated in cognitive deficits, especially in regards to schizophrenia, where one predictor of pharmacological efficacy of an antipsychotic drug is its dopamine-receptor affinity. In this paper, we propose a model for the role of dopamine that integrates its role in reinforcement with that of the aforementioned motor planning.

While the connectivity of the basal ganglia and its ventral extensions, the nucleus accumbens and ventral tegmental area, has been well described, there is no consensus regarding the types of computations these structures perform. Although the overall function is believed to be related to planning and executing actions, especially sequences of actions, it is not clear how the circuitry could accomplish this. Previous attempts to assign function to the basal ganglia circuit have primarily relied upon heuristic arguments without any analytical or computational analysis to test these ideas. Swerdlow and Koob (1987) attempted to understand certain aspects of psychiatric disease by proposing a model based on nested loops of activity through the ventral parts of the basal ganglia. However, their model did not attempt to characterize the behavior of the proposed circuit. Several models have been developed that attempt to integrate the role of cellular reinforcement with action selection in the basal ganglia (Barto, 1995; Berns & Sejnowski, 1996; Graybiel, Aosaki, Flaherty, & Kimura, 1994; Houk, Adams, & Barto, 1995; Montague, Dayan, & Sejnowski, 1996; Wickens & Kotter, 1995); although compelling at the level of classical conditioning, these models have not yet shown how a complex sequence of actions could be implemented by the basal ganglia. We propose in this paper that the basal ganglia project to cortical areas that implement actions and that they filter multimodal information by selecting previously learned optimal actions based on the instantaneous cortical state. We further show how the connectivity is ideally suited to the production of optimal action sequences. An earlier model, which was restricted to action selection, has appeared elsewhere (Berns & Sejnowski, 1996).

MODEL

We modeled the basal ganglia as groups of simplified neurons that corresponded to the various divisions of the structure. The primary assumption was that the basal ganglia facilitate the production of action sequences based on cortical states (for review of this hypothesis,

see Cromwell & Berridge, 1996; Marsden & Obeso, 1994). In order to do this, the basal ganglia must have both a mechanism that selects competing actions and a mechanism by which to learn sequences. Thus, our model demonstrates how the architecture of the basal ganglia is well suited to the production of action sequences and how diffuse reinforcement mechanisms allow for learning.

Selection Model

The output stage of the basal ganglia, the internal segment of the globus pallidus (GPi), is known to be almost wholly GABAergic and tonically active. Thus the GPi tonically inhibits the target thalamic nuclei (VLm, VLpc/mc, VLo, CM). Because these nuclei may also gate ascending information to their cortical targets (motor, supplementary motor, prefrontal cortices), it is reasonable that they should be tonically inhibited until the ascending information is required for action. Furthermore, release from tonic inhibition in the thalamus leads to a very rapid postinhibitory rebound. Inhibition of the corresponding GPi neurons may thus lead to a more rapid response in the thalamus than if they were directly excited. Such a rapid response would be necessary to produce rapid sequences of activity. In our model, we assumed that a mechanism exists in the GPi so that within a pool of neurons, only one is inhibited at a time, a form of “loser-take-all” (Berns & Sejnowski, 1996).

Each structure within the basal ganglia was constructed of a set of units, each of which represented some locally distributed function (Figure 1). The striatal units represented pools of spiny neurons that reflect the cortical state, which itself is comprised of both external representations (e.g., parietal areas) and internal, planning, representations (e.g., frontal cortex). The striatum mapped the cortical state onto a finite set of actions, as represented in the globus pallidus internus (GPi), or output layer. The striatal-GPi, or “direct” (Alexander & Crutcher, 1990), pathway thus allows for the direct mapping of the striatal state onto a particular action. However, if one assumes that the striatum reflects the cortical state, the direct pathway really maps the cortical activity onto a finite set of possible actions in the globus pallidus, without any automatic sequencing occurring in the basal ganglia.

It has been hypothesized that the basal ganglia are involved in the production of action sequences (Cromwell & Berridge, 1996). In order to produce a sequence efficiently, either a short-term memory must exist or the output state must be fed back to the input layer. There is evidence for the latter in the form of segregated cortical-subcortical loops (Alexander et al., 1986; Hedreen & DeLong, 1991) that presumably feed back information from the globus pallidus to the cortical area from which the inputs originated. However, the time course of activity propagation through this loop is probably too




Figure 1. Overall schematic of the basal ganglia model. The striatum (STR) was modeled as two types of neurons: a set of matrix neurons projecting to the globus pallidus and a set of striosomal (or patch) neurons projecting to the substantia nigra pars compacta (SNc) and ventral tegmental area (VTA). The striatal to globus pallidus pathway diverges with inhibitory projections to both the external (GPe) and internal (GPi) segments. Similarly, the subthalamic nucleus (STN) has divergent excitatory projections to both segments. The GPe also has a recurrent, inhibitory projection to the STN, which allows for the short-term storage of activity patterns. The GPi is assumed to be arranged topographically according to action and projects to the corresponding thalamic (Thal) neurons, which are tonically inhibited. The SNc/VTA signal represented a diffuse, synaptic reinforcement signal, which was given by the difference between convergent projections from the striatum and the GPi. The striatal projection represented a prediction of reinforcement, the summed GPi projection represented an estimate of the degree of match between the STR-GPi and STN-GPi signals, and the SNc/VTA signal was the difference between the two.

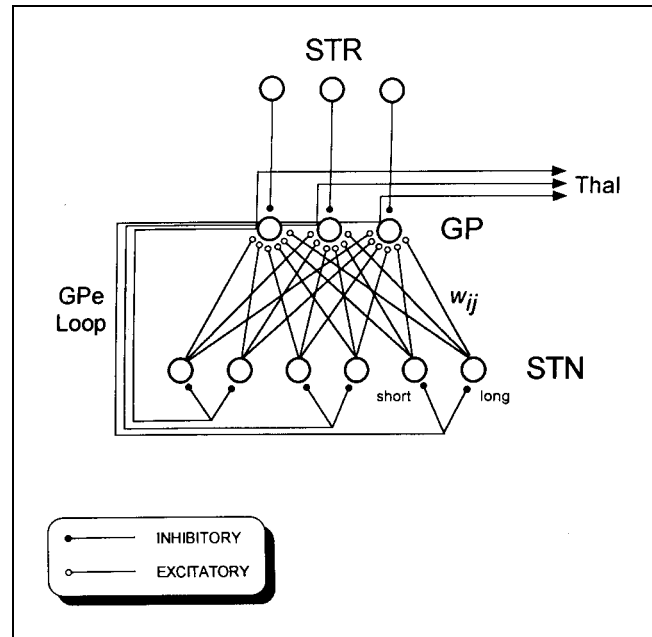


Figure 2. Detailed schematic of the network model. The globus pallidus has been simplified to one set of units. Each unit represented a pool of related neurons in a particular structure and corresponded to some discrete action. The GPe was represented by a fixed projection to corresponding STN units, and the STN-GPi projection was a matrix of weights (w_{ij}), which was modified by the learning rule given in Equation 6. The striatum was the input layer and represented a conjunction of cortical states and specified which action to select during learning by inhibiting one of the GP units. The GP layer specified the action to perform, and the GP unit with the minimum activity was assumed to be the one that selected the action. Additionally the GP layer projected topographically to the subthalamic nucleus via the GPe-STN pathway. The STN layer had two units for each action, one with a short time constant (7 msec) and one with a long time constant (90 msec). The time constants specified the time interval over which the STN activities changed significantly.

long to allow for action sequences on the millisecond time scale. Thus, if the sequencing hypothesis is true, there must be either local short-term memory or local feedback within the basal ganglia. In fact, the presence of local feedback itself could result in short-term memory.

The STN-GP loop contains recurrent projections through the GPe (Smith, Wichmann, & DeLong, 1994a), and we propose that this is a putative site of local short-term memory. This was modeled by diffuse excitatory projections from STN to GP and specific inhibitory connections from GP to STN (Figure 2). The activity of each STN unit represented the average firing frequency of a pool of neurons and was modeled as a leaky integrator:

$$\tau \frac{dB_i(t)}{dt} = -B_i(t) - \alpha G_i(t - \delta) \quad (1)$$

where $B_i(t)$ is the activity of the i -th STN unit, $G_i(t)$ is the activity of the corresponding GP unit (the sign is negative because it is an inhibitory synapse via the GPe), α is a constant specifying the magnitude of the inhibitory synapse, τ is the time constant, and δ is the synaptic delay between the GP and STN units. With $\alpha > 1$, the inhibition overrides the STN excitation. The time constant, τ , in Equation 1 represents both the actual membrane time constant and the time constant associated with activity in the recurrent STN-GPe loop. For the simulations, a nonlinear discretized version of Equation 1 was used, and the membrane potentials were converted to normalized representations of firing frequency by a sigmoidal function, σ . For each STN unit, this was determined by

$$B_i(t) = \sigma[\lambda B_i(t-1) - (1-\lambda) \alpha G_i(t-n)] \quad (2)$$

where

$$\lambda = \frac{\tau}{\tau + \Delta t} \quad (3)$$

and the sigmoid function, with a gain of γ and a bias of β , is given by

$$\sigma[x] = \frac{1}{1 + e^{-\gamma(x-\beta)}} \quad (4)$$

and t now represents the time-step number, n represents the number of time steps corresponding to the synaptic delay, and Δt is the length of the time step. λ ranges from 0 to 1. As shown in Figure 2, the activity of the GP units was determined by both striatal and STN inputs. The striatal input was represented by a fixed unit-to-unit inhibitory connection, and the STN input was weighted by the strength of the modifiable connections, w . The sum of these weighted inputs was then transformed by a sigmoid function to give the representation of firing frequency:

$$G_i(t) = \sigma \left(\sum_j w_{ij}(t) B_j(t) - \alpha S_i(t) + \eta_i \right) \quad (5)$$

where $w_{ij}(t)$ is the matrix of connection strengths between STN unit j and GP unit i at time t , $S_i(t)$ is the activity of striatal unit i at time step t , and η_i is the level of noise drawn from a uniform distribution. The action ultimately selected was assumed to be represented by the GP unit with the minimum activity, a form of loser-take-all (Berns & Sejnowski, 1996). Because the GP inhibited the thalamus, the pool of GP neurons with the lowest activity corresponded to the thalamic pool with the maximum activity.

Synaptic Plasticity

Several models of cellular reinforcement have been proposed, all of which rely on the use of a scalar signal to modify synaptic strengths (Houk et al., 1995; Montague, Dayan, Person, & Sejnowski, 1995; Montague et al., 1996; Sutton & Barto, 1990; Wickens & Kotter, 1995). These models have typically proposed that extrinsic rewards are mapped onto a structure such as the ventral tegmental area (VTA), and that dopamine is released in proportion to the extrinsic reward (Houk et al., 1995; Wickens & Kotter, 1995). Dopamine is then proposed to modify synaptic strengths. If dopamine modifies synaptic strength, it must be released in close temporal proximity to the activity that precedes it. Problems with previous models arise because extrinsic reward occurs at times

remote from the activity that produced the action. This problem can be partially overcome by using a temporal-difference model, in which dopamine is released in proportion to errors in prediction of reward (Houk et al., 1995; Montague et al., 1995; Sutton, 1988; Sutton & Barto, 1990), but this still requires either the storage of activity traces for subsequent modification of synaptic strength or the use of a diffusible messenger that regulates local dopamine release (Montague & Sejnowski, 1994).

We propose an alternative learning model. Because of the dichotomy of information flow through both the internal and external segments of the globus pallidus, it is possible for the striatum to “train” a sequence. The striatal activity represents a mapping of the cortical state onto a finite set of possible actions, which are implemented in the GPi, and in the untrained state, the striatum can directly select the action via the direct pathway. However, the striatum also projects to the GPe, as does the STN. Because the GPe-STN loop can also store short-term activity traces, the connection strengths can be modified through associative learning mechanisms such that the GPe-STN can produce a sequence of states that are learned from the striatum.

We used a modified Hebbian learning rule (Hebb, 1949) given by

$$\Delta w_{ij}(t) = \rho (e(t) G_i(t) - S_i(t)) B_j(t) \quad (6)$$

where $\Delta w_{ij}(t)$ is the weight change in the connection from STN unit j to GP unit i at time t , ρ is the learning rate, $e(t)$ is a scalar error signal from the SNc/VTA at time t , $G_i(t)$ is the activity of GP unit i , and $B_j(t)$ is the activity of STN unit j . The first term, $\rho e(t) G_i(t) B_j(t)$, represents a classical Hebbian synapse with presynaptic activity from STN and postsynaptic activity from GP, modulated by the error term, $e(t)$. The second term, $-\rho S_i(t) B_j(t)$, represents a different form of plasticity that decreases the synaptic efficacy when the postsynaptic cell is actively inhibited by a GABAergic synapse. This is an important difference because it distinguishes the situations when a postsynaptic cell is inactive because of lack of presynaptic activity from that in which the postsynaptic cell is inactive because it has been actively inhibited, in this case by the striatum. In the latter case, any concurrently active excitatory synapse decreases in strength, whereas in the former case no change occurs.

The error signal, $e(t)$, was generated by the mismatch between the direct and indirect pathways and was assumed to originate from the dopaminergic neurons in the substantia nigra (SNc) and ventral tegmental area (VTA). As shown in Figure 3, the striatum sends a convergent projection to the error unit, and the GP also sends a convergent projection to the error unit, via a hypothesized inhibitory neuron. During the training phase, the striatal inhibition overrides the STN excitation of GP, which results, for extreme levels of activity, in the logic table shown in Table 1. Thus only when the STN




Figure 3. Schematic of the error/reinforcement mechanism. The striatum (STR) sends a convergent, inhibitory projection to the dopamine containing neurons of the substantia nigra compacta (SNc) and ventral tegmental area (VTA). It is assumed that this projection arises from the striosomal compartment, and these connections (v_i) were modifiable in strength. The globus pallidus (GP) units also projected convergently to the SNc/VTA but via an inhibitory interneuron, which effectively reversed the sign of the signal. The GP projection represented the sum total of the matches between the GP activity and the striatal activity. According to Table 1, the highest match occurred when maximal subthalamic nucleus (STN) activation coincided with minimal striatal inhibition. The output of the error unit (e), computed as the difference between the weighted striatal signal and the GP signal, was projected back to both the STR-SN (v_i) and STN-GP (w_{ij}) weights where it modulated weight changes according to Equations 8 and 6, respectively.

input to a GP cell is high and there is no striatal input is there significant GP activity, which functions as a type of inverse coincidence detector. By summing all the GP activities, a scalar estimate of the total degree of coincidence is produced. A striatal estimate is produced by summing the vector of striatal activities, weighted by the connections, $v_i(t)$ (Figure 3). Ideally this would match the estimate obtained from the GP, and the difference between the two represents the error signal at time step t :

$$e(t) = \sum_i (G_i(t) - v_i(t)S_i(t)) \quad (7)$$

where $v_i(t)$ is the connection strength from striatal unit i to the error unit. Furthermore, the connection strengths, v_i , are also modified by a rule similar to that of Equation 6, so eventually the error signal tends toward zero:

$$\Delta v_i(t) = \rho e(t)S_i(t) \quad (8)$$

Table 1. Logic table for the GP units given hypothetically extreme values of input from both the STN and striatum. Striatal inhibition overrides STN excitation, so maximal GP activation occurs with both maximal STN activity and minimal striatal activity.

STN	Striatal	GP
0	0	0
1	0	1
0	1	0
1	1	0

RESULTS

During the training phase, in the simulation with five units, the activities of the GP units gradually increased as the STN inputs to them increased (Figure 4). The winning unit was the one with lowest activation, and during training, the striatum directly selected the winning unit. Hence at each time step, one, and only one, unit was off. The STN weights associated with the GP units increased throughout training until the error had declined to zero. As shown in Figure 5, the weight of the connection from the STN unit that corresponded to the action preceding the one in the GP grew at a slower rate than the others. This reflected the fact that, on average, both the pre- and postsynaptic activities for this synapse were lower than the others.

During training, the internal reward, by which we mean the scalar measure of the match between the striatal and GP activities and computed by the sum of GP activities, steadily increased (Figure 6A). Most of this effect resulted from the fact that the connection strengths from the STN were increasing, driving the GP units to higher activities. Part of the effect was also due to the increasingly better match between the striatal and STN inputs to the GP, as illustrated in Table 1. The error signal, which represented the difference between the reward predicted by the striatal patch and that of the actual internal reward (from the sum of the GP activities), increased during the initial learning phase in which there was a rapid change in connection strengths. This steadily decreased as the striatal weights, v_i , became modified to predict the amount of reward associated with a particular striatal state (Figure 6B). As the error tended toward zero, further changes in connection strength ceased.

In order to test the model's ability to reproduce a sequence trained in the aforementioned manner, the striatal activity vector corresponding to the initial state in the training sequence was loaded for one time step into the striatal units. As shown in Figure 7, the subsequent GP activities reproduced this sequence well; however, the noise led to an imperfect reproduction as

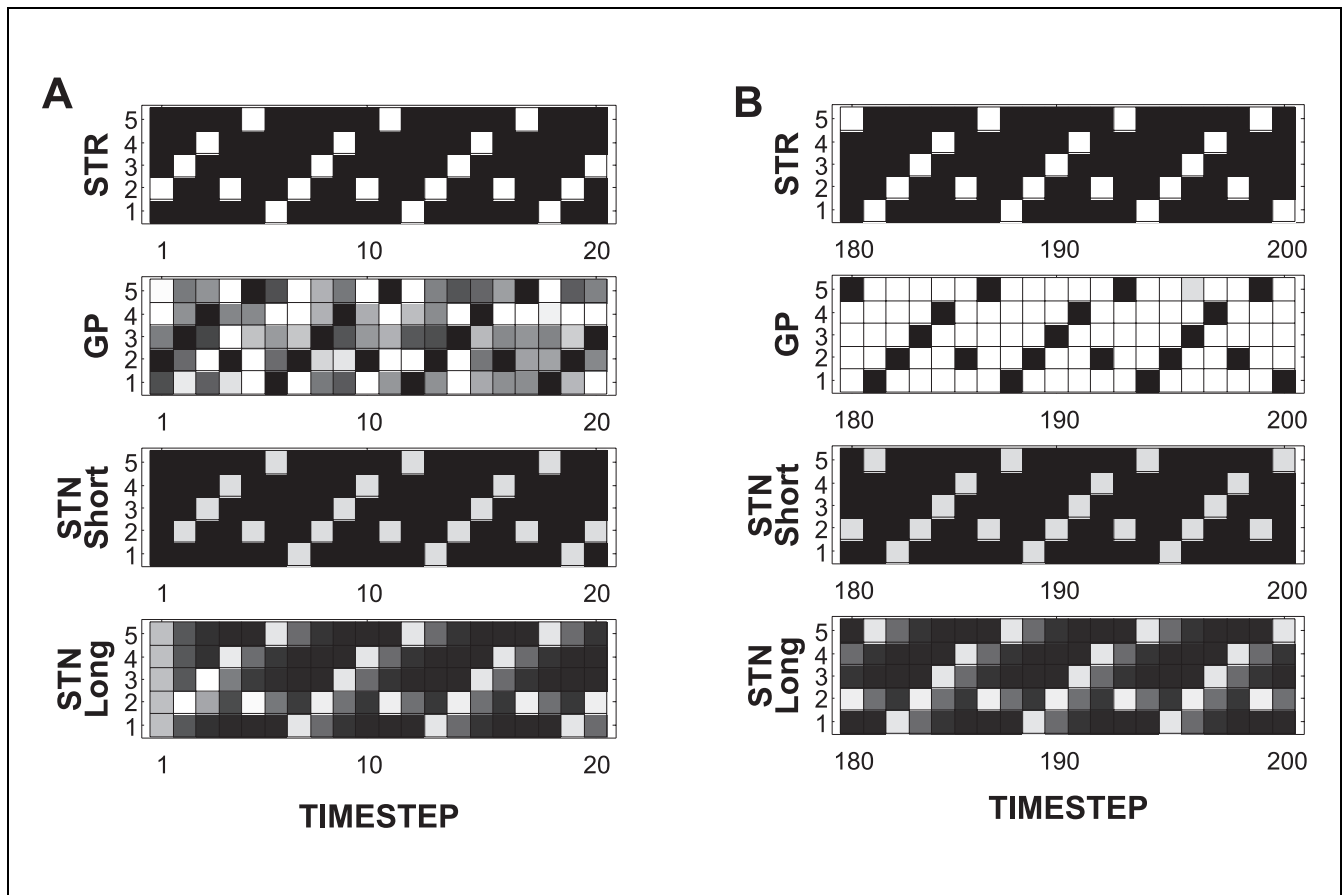


Figure 4. Unit activities during learning the sequence 1, 2, 3, 4, 2, 5. With layers of five units each, activities are shown from 0 (black) to 1 (white) for striatum (STR), globus pallidus (GP), and the two subthalamic nucleus layers with short time constant (STN Short) and long time constant (STN Long). Panel A shows the activity patterns during the initial 20 time steps of training, and Panel B shows the activity patterns after 200 time steps. Using the parameters given in Table 2, the striatum trained the globus pallidus to produce a sequence of actions. Initially, the GP activities were low and disorganized because of minimal excitation from the STN. Subsequently, the weights, and hence the GP activities, increased except for those corresponding to the action that was actively inhibited by the striatum.

evidenced by the error at time step 12. The ambiguity associated with unit 2 is resolved by the differential activations between the short and long STN units. More complex sequences, typically of length greater than 10, exceeded the model's ability to reliably reproduce them. More complex ambiguities also resulted in incorrect reproductions.

The parameter sensitivity revealed that the model is robust, but beyond certain limits various degradations occurred. The ratio of STR/GP learning rates significantly affected performance. Decreasing the ratio to 1 resulted in persistently high error signals because the STR prediction was slow to learn, which in turn resulted in STN weights that continued to increase longer. With persistently increasing weights, the model lost the ability to disambiguate the context of certain activity patterns, yielding the sequence 1, 2, 5, 1, 2, 5, With a learning ratio of 4, the same sequence was produced, but this was due to the rapid cessation of learning as the striatal weights rapidly adjusted and the error went to zero

before many of the STN to GP weights had achieved their correct values. Diminishing the degree of inhibitory override, by decreasing α to 1, resulted in maximal activation of all the GP units during training because the striatum had insufficient inhibition to directly select an action. The end result, after training, was a uniform weight matrix with all weights close to 1. With this weight matrix, the sequence could not be produced at all. Changing the gain (γ) and bias (β) parameters, with gains ranging from 2 to 8 and biases ranging from 0.1, to 0.2, did not significantly affect the production of sequences; however, certain combinations of gain and bias yielded GP activities that were subtly different.

The aforementioned sequence demonstrated how the model learned a sequence requiring the disambiguation of context. We also tested the model's ability to shift between a random sequence and a repeating 10-item sequence. This was done, in part, to test the model on a well-studied behavioral task of procedural learning (Willingham, Nissen, & Ballester, 1989). The model was pre-




Figure 5. Changes in connection strengths, w_{ij} , from learning the sequence 1, 2, 3, 4, 2, 5. The five weights from the five STN units with short time constants to GP unit 2 are shown. The three weights that increased to saturation levels were from STN units 2, 3, and 5 (i.e., those STN units that were not active prior to GP unit 2 being active). Conversely, the weights from STN units 1 and 4 did not increase significantly because when these units were active, GP unit 2 was inhibited by the striatum.

sented with 100 trials of randomly ordered stimuli (1, 2, 3, or 4), then 40 repetitions of the sequence 4-2-3-1-3-2-4-3-2-1, followed by another 100 trials of random stimuli. In order to compare the GP output to previously reported reaction times, the GP output was linearly transformed by

$$R(t) = 1 - \frac{1}{N} \sum_{i=1}^N G_i(t) \quad (9)$$

where N was the number of GP units (4 in this case). $R(t)$ represented a normalized reaction time at time step t and ranged from 0 to 1. This linearly scaled the match between the direct and indirect pathways, with the better the match, the lower the reaction time. As shown in Figure 8A, the reaction time initially declined even with a random sequence and then rapidly reached a stable level with the introduction of the repeating sequence. It stabilized at the value 0.25 because the inherent structure of the sequence allowed for maximal activation of all GP units except the one being selected. When the random sequence was reintroduced, the normalized reaction time became slightly longer.

We also used this paradigm to model the effects of Parkinson's disease and the subsequent improvement of symptoms from pallidotomy (Figure 8, parts B and C). Parkinson's disease was modeled by decreasing the learning rate (ρ in Equation 6) from 0.025 to 0.005, reflecting the overall decline in dopamine that is found




Figure 6. Levels of the “reward” (A) from the GP and the error signal from the SNc/VTA (B) during learning the sequence 1, 2, 3, 4, 2, 5. The reward was computed as the sum of the GP activities and was proportional to how well the GP activity vector matched the inverse of the striatal activity vector. As the system learned to produce the sequence, the match, and hence the reward, increased. The error signal, which was computed by Equation 7, represented the difference between a weighted sum of the striatal activity and the sum of the GP activity. The weights on the striatal activities were modified by the error signal, and thus the difference ultimately converged to zero. Note that the variance also decreased.

in Parkinson's disease. This resulted in substantially slower learning, as evidenced by the lower slope in Figure 8B, but because the GP activations were generally lower, the effect of noise was also more prominent. The effects of the decreased learning rate could largely be ameliorated by increasing the gain of both the GP and STN units from 4 to 8. As the gain was increased, units that were previously marginally active became maximally active, and thus the first term in Equation 6 increased, partially offsetting the decreased learning rate. This suggests that a potential mechanism for the efficacy of pallidotomy is in the alteration of the gain of pools of neurons in both the STN and GP. One prediction is that even though the rate of learning is partially restored, the effect of noise still remains.




Figure 7. Globus pallidus activities after learning the sequence 1, 2, 3, 4, 2, 5 and being cued to the beginning of the sequence by inhibiting GP unit 1 for one time step. Activities are shown from 0 (black) to 1 (white) for striatum (STR), globus pallidus (GP), and the two subthalamic nucleus layers with short time constant (STN Short) and long time constant (STN Long). Because of the STN-GP weights that were learned, the system subsequently reproduced the sequence without further input from the striatum. The GP layer, although noisy, correctly produces the sequence with the correct unit having minimal activity at each time step.

DISCUSSION

We have presented a model of the function of the basal ganglia that was based closely on known neuroanatomy, and we have demonstrated how such a neuron-level model can both learn and reproduce action sequences. This differs from other connectionist models based on back-propagation (Cleeremans & McClelland, 1991). Although these have been quite successful at modeling behavioral data, they are not anatomically or physiologically motivated. Our model attempted to incorporate several proposed functions in the basal ganglia. Action selection was assumed to occur through lateral influences in the globus pallidus, whereas sequences were produced via recurrent projections in the GPe-STN loop. Furthermore, both these proposals were based on the

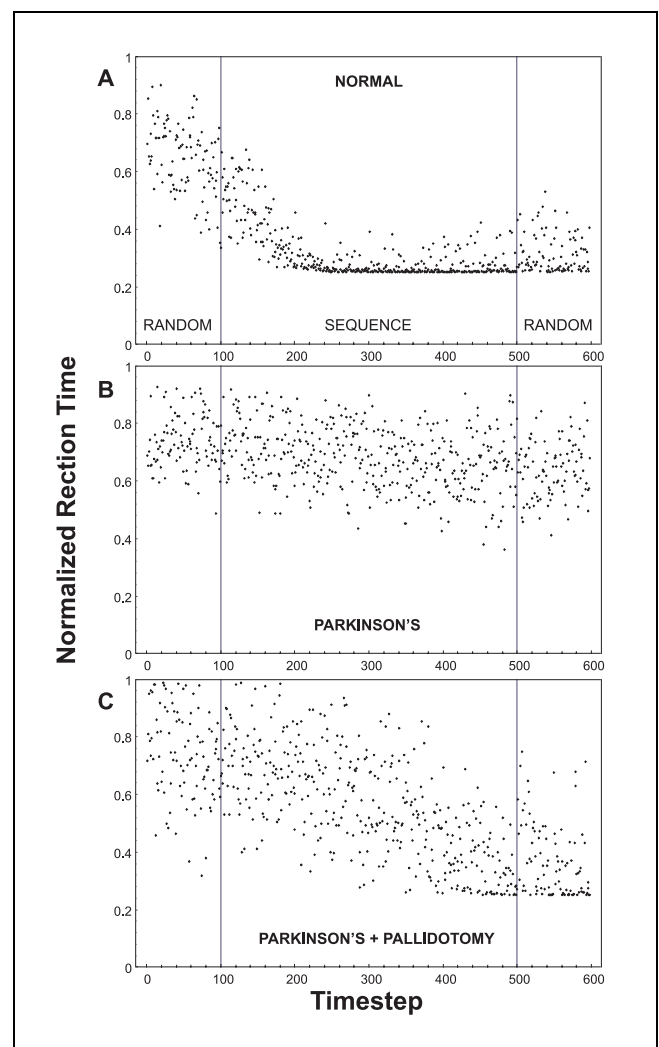


Figure 8. Modeling the effects of Parkinson's disease (PD) and subsequent pallidotomy. The model was presented with 100 trials of random responses followed by 400 trials of stimuli from a fixed, repeating sequence (4-2-3-1-3-2-4-3-2-1) (Willingham et al, 1989) and then by another 100 random trials. The output of the globus pallidus (GP) units was transformed to represent a normalized reaction time by a linear mapping and ranged from 0 to 1 (see text). In the Normal model ($\rho = 0.025$, gain = 4, bias = 0.2), the GP output became more organized, and consequently reaction time improved even with the random sequence. When the repeating sequence was introduced, the GP activation became even more organized, and reaction time improved further until the nonselected GP units were maximally activated. When the random sequence was reintroduced ($t = 500$), the activation pattern was no longer maximal, and the reaction time worsened slightly. Parkinson's disease was modeled by decreasing the overall learning rate ($\rho = 0.005$, gain = 4, bias = 0.2). With these parameters, the rate of improvement in reaction time was slower and showed a greater variance of response for the same level of noise in the system. Pallidotomy was modeled by increasing the gain of the GP and STN units ($\rho = 0.005$, gain = 8, bias = 0.2). Although the variance still remained high, the rate of reaction time improvement was partially improved.

assumption that segregated streams of information flow through the basal ganglia (Alexander et al., 1986). We formalized this segregation by representing each potential action by a separate unit, a so-called grandmother cell; however, this was chosen for computational efficiency. Each of the individual units in our model more realistically represents a pool of neurons devoted to a particular action, but the segregation requirement for pools of neurons remains.

Although there is good evidence for cortical topography being maintained throughout the basal ganglia (Alexander & Crutcher, 1990; Alexander et al., 1986; Goldman-Rakic & Selemon, 1986; Parent, 1990), it is also known that the segregation is not complete. The massive convergence from striatum to globus pallidus (Wilson, 1990) alone requires that inputs cannot remain completely segregated (Flaherty & Graybiel, 1994; Hedreen & DeLong, 1991). However, our model is consistent with this convergence. We modeled the striatum as an input stage in which diverse areas of the cortex map onto subsets of neuron pools. In this manner, diverse sensorimotor modalities are combined with higher representations of both context and timing information from the prefrontal areas. It is the functional mapping from striatum to globus pallidus and thalamus that remains segregated. In other words, the segregation may reflect the final cortical targets, not the afferents. Work with retrograde transneuronal transport of herpes simplex virus injected into the cortical motor areas suggests that the output stages of the basal ganglia are indeed organized into discrete channels that correspond to their targets (Hoover & Strick, 1993).

The basal ganglia have long been assumed to facilitate action sequencing (Cromwell & Berridge, 1996), and despite the fact that there is not a great deal of experimental evidence for the sequencing hypothesis, our model does support this notion. In fact, a review of human behavioral syndromes associated with various basal ganglia lesions did not find any reports of sequencing disturbance (apraxia) (Bhatia & Marsden, 1994). An alternative, but related, hypothesis is that the basal ganglia monitor automatic behavior and alter it in novel contexts (Marsden & Obeso, 1994). Our model focused on the role of the pallidal-subthalamic loop in facilitating sequence production, but the output of this network may be used in other ways. For example, the sequence information could be used to monitor the reliability of sequential predictions maintained elsewhere such as in the cortex or the striatum. In our model, this actually occurs in the degree of match between the direct and indirect pathways.

The primary requirement for automatic sequence production was the presence of recurrent connections that both relay information regarding previous actions and result in short-term memory. Until recently, the presence of recurrent projections has not been a prominent finding in basal ganglia anatomy; however, such path-

ways exist locally in the primate pallidal-subthalamic loop (Smith et al., 1994a; Smith, Wichmann, & DeLong, 1994b). It has previously been argued that the most obvious recurrent pathway is the thalamo-cortical projection that presumably closes the cortical-subcortical loop (Alexander et al., 1986; Berns & Sejnowski, 1996; Graybiel et al., 1994; Houk et al., 1995; Middleton & Strick, 1994), but the time taken for impulses to traverse this loop appears to be on the order of 100 msec (Wilson, 1990), which would be too long for rapid automatic sequencing. Within the basal ganglia, the striatum itself may have memory properties, and there is evidence that striatal lesions impair innate grooming sequences (Cromwell & Berridge, 1996). In slice preparations, striatal cells show bistable activity, existing in either a low state or a high, subthreshold, state from which action potentials irregularly appear (Kawaguchi, Wilson, & Emson, 1989; Wilson, 1986) and may represent context information for the conjunction of either cortical states or, in the case of grooming, other brainstem inputs. Our model does not directly address the role of the striatum in action sequencing other than representing a conjunction of inputs, but rather the model demonstrates how the GP-STN loop could automate sequence learning, which may be different from the mechanisms for innate sequences.

Similarly, we have not attempted to incorporate the role of the prefrontal cortex (PFC) in the model, even though it is closely connected to the basal ganglia. A large body of literature has implicated the prefrontal cortex, especially the lateral areas, in tasks requiring the maintenance of information for short periods of time (Fuster, 1973; Goldman-Rakic, 1987). Both the architecture and size of the prefrontal areas suggest that the PFC performs far more complex operations than sequencing, although this may be part of its function. As noted above, the contribution of the basal ganglia may be in a monitoring role, following along with highly automated sequences and detecting deviations from the expected predictions. Another important difference is the time scale on which these two structures may operate. PFC neurons can hold information on the order of seconds (Fuster, 1993), whereas the proposed local memory in the basal ganglia may be operating on the scale of 10 to 100 msec.

We have proposed that the GPe-STN loop stores short-term traces of activity and functions as working memory for the production of action sequences. In a study of pallidal activity during sequential arm movements, Mushiaki and Strick (1995) found subsets of pallidal neurons that only displayed activity changes, generally decreases, during specific phases of the movement sequence. These neurons did not display the same changes when the same task was explicitly guided, suggesting a role in both sequencing and memory. The function of the subthalamic nucleus has remained elusive and has traditionally been conceived in terms of tonic excitation of

the globus pallidus to prevent unwanted activity from being initiated. Although no direct evidence exists for the STN's role in working memory, both *in vivo* recordings and lesion data can be interpreted within this framework. Studies by Wichmann, Bergman, and DeLong (1994), demonstrated that STN neurons fire around the onset of movement, which would be considered a relatively late increase in activity. This would, however, be consistent with a storage function. Similarly, pallidal activity changes have been found to occur too late for movement initiation (Mink & Thach, 1991). STN lesions have not consistently been shown to decrease the firing rate of GP neurons, which one would expect if its function were solely tonic excitation, but instead they regularize the GP firing pattern (Ryan & Sanders, 1993). Conversely, lesioning the globus pallidus results in an increase in STN bursting cells (Ryan, Sanders, & Clark, 1992). The latter lesion may represent a large-scale inhibition of the GP, somewhat analogous to the effect of the striatum on the GP, but on a local scale. The appearance of bursting would be consistent with a memory storage function as activity oscillates through the remaining cells in the GPe-STN loop. STN lesions have been found to cause a condition in which animals are subject to involuntary, large movements of whole limbs (hemiballismus). Conceptualizing this within the framework of a local memory function, one would hypothesize that the loss of local memory that occurs with a subthalamic lesion results in the loss of context for movements. In this case, striatal representations cannot be disambiguated without the context of preceding actions and thus may result in the activation of large pools of neurons, with the effect being a large involuntary movement when the intention was a small specific one.

A major assumption of our model is that the dopamine-containing neurons of both the substantia nigra pars compacta (SNc) and the ventral tegmental area (VTA) modulate synaptic efficacy in response to intrinsic reward. There is strong evidence that dopamine plays an important role in at least extrinsic, reward-driven, learning. In studies by Schultz et al. (Ljungberg, Apicella, & Schultz, 1992; Schultz et al., 1993; Schultz et al., 1992), it has been demonstrated that dopamine neurons of the VTA fire transiently during the learning of an operant task when a reward is given, but that after the task is learned, the dopamine neurons do not fire in response to reward. Furthermore, when reward is withheld after the task is learned, the dopamine neurons show a transient depression in activity (Schultz et al., 1993). This result is consistent with the existence of a projection to these neurons that contains predictive information regarding future reward. In a previous model of mesencephalic dopamine systems, it was proposed that dopamine is released in proportion to a temporal-difference of reward prediction, but this model required a complete temporal representation of preceding stimuli in the form of a tapped-delay line (Montague et al., 1996).

A similar type of representation appears in the present model in the form of multiple time constants within the GP-STN loop. Furthermore, we propose that in addition to firing in response to extrinsic rewards, the dopamine neurons also fire in response to intrinsic rewards, by which we mean internal representations of how well the basal ganglia are performing and which we have modeled as originating from the GPi. We hypothesize that the striatonigral (or in the ventral striatum, the accumbal-VTA projection) carries a prediction of reward (Barto, 1995; Sutton & Barto, 1990), and because of its known projections to the dopamine-containing structures, this arises from the striosomal compartment (Graybiel, 1990; Houk et al., 1995).

As implemented in our model, the dopamine activity can be either positive or negative. We have modeled the role of dopamine as modulatory rather than as either an excitatory or inhibitory influence. From Equation 7, the quantity $e(t)$ represents this signal and specifies both the magnitude and direction by which the synapse efficacy changes. Although it is not formalized in the equation, the tacit assumption in allowing $e(t)$ to be either positive or negative is that a tonic level of dopamine activity is required to maintain the synaptic efficacy. Thus a depression from maintained activity would lead to a decrease in synaptic efficacy, whereas an increase in dopamine activity would lead to an increase in synaptic efficacy, assuming the correlation of both pre- and postsynaptic activities. There is evidence that the release of dopamine in the striatum is biphasic and depends on the learned availability of actions. In an experiment in which mice were allowed to escape foot shock, there was an increase in the accumbal concentration of the dopamine metabolite 3-methoxytyramine (3-MT) but a decrease in 3-MT if the mice were not allowed to escape (Cabib & Puglisi-Allegra, 1994), indicating that dopamine release both increases and decreases according to learned behaviors.

The further requirement imposed by Equation 6 is that the dopamine-modulated synaptic change only occurs at those synapses where there was a preceding correlation between pre- and postsynaptic activity. Thus in our model, dopamine modulates a classical Hebbian synapse. Presently both long-term potentiation (LTP) and long-term depression (LTD) are the best candidates for activity-dependent changes in synaptic efficacy, although there may be other mechanisms. The focus for LTP has centered on the hippocampus and its role in memory, but both LTP and LTD have also been reported in the ventral striatum (Calabresi, Maj, Pisani, Mercuri, & Bernardi, 1992; Kombian & Malenka, 1994). Kombian and Malenka reported that LTP was produced in the core of the nucleus accumbens by tetanic stimulation of the cortical input and that the LTP was mediated mainly by non-NMDA receptors. They further showed that LTD occurred in the same neurons via the NMDA receptor, suggesting that depending on the intracellular calcium

concentration, either LTP or LTD could occur in a single striatal cell. Dopamine is known to act on intracellular calcium (Cooper et al., 1996), and the areas of the brain with the highest calmodulin-dependent phosphodiesterase concentration correspond to those areas with heavy dopaminergic innervation (Polli & Kincaid, 1994). It may be that dopamine modulates the LTP/LTD behavior of both striatal and pallidal cells by altering the concentration of intracellular calcium, but it is likely that other second messenger systems are also involved.

The second component of synaptic plasticity in the model distinguished between active inhibition of the postsynaptic cell and inactivity from lack of presynaptic excitation. The classical Hebbian synapse requires only coincident pre- and postsynaptic activity without regard to the different contributors to the postsynaptic potential. In order to differentiate active inhibition, one must postulate the existence of a second messenger that weakens any synapse that is concurrently active. There is experimental evidence for long-term depression (LTD) of synaptic activity when presynaptic excitatory impulses are coupled with GABAergic activity (Thiels, Barriónuevo, & Berger, 1994; Wagner & Alger, 1995; Yang, Connor, & Faber, 1994). The underlying mechanism could involve decreases in intracellular calcium. In our model, this effect became significant for reorganization of weights when changing from one sequence to another. The model predicts that any process that weakens the striatal inhibition of the globus pallidus should lead to an impairment in sequence shifting. Processes that might do this include striatal degeneration (e.g., Huntington's disease and striatal-nigral degeneration). Conversely, a process that facilitates GABAergic transmission (e.g., benzodiazepines) would lead to an increase in sequence shifting, possibly manifested as an inability to concentrate on a task.

The learning and reproduction of a sequence of states in our model suggests mechanisms by which the primate basal ganglia may automate action sequences. Beginning from a state in which the weights were all zero, the model rapidly converged to the proper set of connections necessary to reproduce the sequence. However, altering the ratio of striatal to pallidal learning rates, which can be roughly interpreted as a differential effect of dopamine at these sites, led either to a saturation of the weights (low ratio) or a rapid cessation of learning (high ratio). In both cases, this impaired the sequence learning. Diminishing the degree of inhibitory override from the striatum, which would be analogous to the loss of spiny neurons in Huntington's disease, led to the saturation of all the STN-GP weights. This resulted not only in the inability to produce any sequence, but it generated random actions that were driven by the level of presynaptic noise in the system. This could be analogous to the choreiform movements typical of Huntington's disease.

In the model of Parkinson's disease, a dopamine deficit

was modeled by decreasing the overall learning rate. As shown in Figure 8, this resulted in slower weight changes and consequently an impairment in procedural learning. Parkinson's disease patients have been thought to suffer from deficits in both motor initiation and shifting as well as cognitive shifting (Cools, Van Den Bercken, Horstink, Van Spaendonck, & Berger, 1984; Marsden & Obeso, 1994), although it is not clear whether shifting deficits are separate from the overall slowing observed on motor tasks. The specific sequence shifting task that we modeled was selected to test the model against existing experimental data in Parkinson's disease (Pascual-Leone, et al., 1993). Using a linear mapping for R (Equation 9), the model results closely match the deficits observed in Parkinson's disease. Not only does it capture the slower rate of improvement, but it also displays the same increase in variance in R . In the model, the "signal" was relatively attenuated because of the lower connection weights and their slow rate of change, and thus the same level of noise in the system resulted in a larger variance of R . This model suggests that so-called switching deficits may simply be an aspect of slow learning, which is consistent with studies finding no basic switching deficit (Brown & Marsden, 1988; Downes, Sharp, Costall, Sagar, & Howe, 1993). By slow learning, we really mean a decrease in plasticity at the neuron level, which would be particularly evident with tasks requiring responses to new information. This could manifest itself as a switching deficit if the task involved new material or was designed to require rapid learning, or as a prediction deficit if external feedback was absent (Flowers, 1978). Our model also suggests that there is a differential effect between unlearning a previous response set and learning a new one. Because striatal inhibition overrides pallidal excitation, connection weights will generally decrease faster than they increase, and this effect will be more pronounced when the overall learning rate is decreased, as in the model of Parkinson's disease. This is consistent with previous findings of Parkinsonian impairments of both shifting away from a previously learned dimension and shifting to a new dimension (Owen, et al., 1993).

It is curious then that a focal lesion of the globus pallidus can alleviate some of the symptoms of Parkinson's disease, but there is substantial evidence for the efficacy of the pallidotomy procedure in Parkinson's disease. It has been proposed that the procedure works by lesioning a hyperactive globus pallidus, thus restoring balance (Iacono, et al., 1995; Marsden & Obeso, 1994). Our model suggests that the pallidotomy lesion, rather than simply restoring normal tonic inhibition, alters the input-output relationship in pools of both pallidal and STN neurons. This partially offsets the decreased learning rate because of an increased signal-to-noise ratio (i.e., better separation of the units that should or should not be active). Thus those units that are activated with the increased gain have maximal activity, which maximizes

the first term in Equation 6 and offsets the decreased learning rate. By lesioning neurons within a pallidal pool, the effective gain of the corresponding STN neurons may be increased by effectively uncoupling them from each other. If one assumes that each of the neurons contributing to the pool is fundamentally a binary (off/on) unit, uncoupling them from each other allows each individual cell to have a greater signal-to-noise ratio (i.e., each STN cell is better able to detect significant inputs because they are not being averaged with other cells). In an animal model of Parkinson's disease, when the STN was lesioned, the percentage of pallidal cells responding to an external stimulus decreased (Wichmann, Bergman, & DeLong, 1994b)

We have proposed a systems-level model of the basal ganglia that attempts to bridge the gap between the anatomy and the function of automatic sequence production. Local working memory was postulated to exist in the form of activity patterns in the GPe-subthalamic feedback loop, to which damage is predicted to cause incorrect movements because of a loss of ability to disambiguate the action context on a short time scale. Furthermore, diseases that alter the striatal inhibition of the globus pallidus are predicted to either decrease the ability to shift sequences or cause false shifting. In order to learn sequences, we have proposed that the striatum trains the globus pallidus and subthalamic nucleus to produce sequences of states, and dopamine modulates the synaptic efficacy to achieve this. In order to efficiently shift sequences, a mechanism by which synaptic efficacy is weakened depends on the GABAergic modulation of long-term depression. Although many details of the model are based on assumptions, the results of the model fit a large body of behavioral and physi-

ological data, and it suggests a framework for conceptualizing basal ganglia function.

METHODS

All simulations were programmed in Mathematica, version 2.2 (Wolfram, 1991), running on a 486-33, Pentium-90, or a Sparc 10. For 100 iterations, the simulation time ranged from approximately 10 sec on the Sparc 10 to 1 min on the 486-33. Several simulations were performed with varying numbers of units and sequence complexity. Values for the equation parameters are shown in Table 2. The learning rate for the STR→VTA weights was 2 times that of the STN→GP weights. Two sets of STN units were used, one with a time constant, τ , of 7 msec (Nambu & Llinas, 1994), which according to Equation 3 and a time-step interval of 10 msec, corresponded to a λ of 0.4; the other set of STN units had a τ of 90 msec, corresponding to a λ of 0.9. The synaptic delay was taken to be one time step ($n = 1$). Each unit in the simulation represented a pool of neurons, with each pool corresponding to a particular action. As a test of the model, the sequence (1, 2, 3, 4, 2, 5) was learned using 5 units in the STR and GP layers and 10 units in the STN layer (two integration times for each unit). This sequence requires short-term memory in order to disambiguate the action following Unit 2. All weights were constrained to the range 0 to 1. Larger simulations, with up to 10 units were also performed.

Acknowledgments

We thank Francis Crick for discussion during the development of this model and comments on early drafts of this paper and P. Read Montague and Peter Dayan for comments on the manuscript. This work was supported by funds from the American Psychiatric Association and Ely Lilly to G. S. B., and the Wasie Foundation to T. J. S.

Reprint requests should be sent to Gregory S. Berns, Western Psychiatric Institute & Clinic, University of Pittsburgh, 3811 O'Hara St., Pittsburgh, PA 15213, or via e-mail: berns+@pitt.edu.

REFERENCES

- Alexander, G. E., & Crutcher, M. D. (1990). Functional architecture of basal ganglia circuits: Neural substrates of parallel processing. *Trends in Neuroscience*, *13*, 266-271.
- 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.
- Barto, A. G. (1995). Adaptive critics and the basal ganglia. In J. C. Houk, J. L. Davis, & D. G. Beiser (Eds.), *Models of information processing in the basal ganglia* (pp. 215-232). Cambridge, MA: MIT Press.
- Berns, G. S., & Sejnowski, T. J. (1996). How the basal ganglia make decisions. In A. Damasio, H. Damasio, & Y. Christen (Eds.), *The neurobiology of decision making* (pp. 101-113). Berlin: Springer-Verlag.
- Bhatia, K. P., & Marsden, C. D. (1994). The behavioural and

Table 2. Parameter values for the model.

Parameter ^a	Value
ρ	0.05 (<i>w</i>) or 0.1 (<i>v</i>)
γ	4
β	0.1
Δt	10 msec
n	1
λ	0.4 or 0.9
α	10
η	0.5

^a ρ = overall learning rate, γ = gain, β = bias, Δt = length of simulated time step, n = number of time steps corresponding to synaptic delay, λ is defined by Equation 3, α = relative effect of an inhibitory synapse to an excitatory one, η = magnitude of random synaptic noise.

- motor consequences of focal lesions of the basal ganglia in man. *Brain*, *117*, 859–876.
- Brown, R. G., & Marsden, C. D. (1988). Internal versus external cues and the control of attention in Parkinson's disease. *Brain*, *111*, 323–345.
- Cabib, S., & Puglisi-Allegra, S. (1994). Opposite responses of mesolimbic dopamine system to controllable and uncontrollable aversive experiences. *Journal of Neuroscience*, *14*, 3333–3340.
- Calabresi, P., Maj, R., Pisani, A., Mercuri, N. B., & Bernardi, G. (1992). Long-term synaptic depression in the striatum: Physiological and pharmacological characterization. *Journal of Neuroscience*, *12*, 4224–4233.
- Cleeremans, A., & McClelland, J. L. (1991). Learning the structure of event sequences. *Journal of Experimental Psychology: General*, *120*, 235–253.
- Cools, A. R., Van Den Bercken, J. H. L., Horstink, M. W. I., Van Spaendonck, K. P. M., & Berger, H. J. C. (1984). Cognitive and motor shifting aptitude disorder in Parkinson's disease. *Journal of Neurology, Neurosurgery, and Psychiatry*, *47*, 443–453.
- Cooper, J. R., Bloom, F. E., & Roth, R. H. (1996). *The biochemical basis of neuropharmacology*, 7th ed. New York: Oxford University Press.
- Cromwell, H. C., & Berridge, K. C. (1996). Implementation of action sequences by a neostriatal site: A lesion mapping study of grooming syntax. *Journal of Neuroscience*, *16*, 3444–3458.
- Downes, J. J., Sharp, H. M., Costall, B. M., Sagar, H. J., & Howe, J. (1993). Alternating fluency in Parkinson's disease. An evaluation of the attentional control theory of cognitive impairment. *Brain*, *116*, 887–902.
- Flaherty, A. W., & Graybiel, A. M. (1994). Input-output organization of the sensorimotor striatum in the squirrel monkey. *Journal of Neuroscience*, *14*, 599–610.
- Flowers, K. (1978). Lack of prediction in the motor behavior of Parkinsonism. *Brain*, *101*, 35–52.
- Fuster, J. M. (1973). Unit activity in prefrontal cortex during delayed-response performance: neuronal correlates of transient memory. *Journal of Neurophysiology*, *36*, 61–78.
- Fuster, J. M. (1993). Frontal lobes. *Current Opinion in Neurobiology*, *3*, 160–165.
- Goldman-Rakic, P. S. (1987). Circuitry of primate prefrontal cortex and regulation of behavior by representational memory. In V. B. Mountcastle (Ed.), *Handbook of physiology. Section 1: The nervous system* (Vol. V, pp. 373–417). Bethesda, MD: American Physiological Society.
- Goldman-Rakic, P. S., & Selemon, L. D. (1986). Topography of corticostriatal projections in nonhuman primates and implications for functional parcellation of the neostriatum. In E. G. Jones & A. Peters (Eds.), *Cerebral cortex. Sensory-motor areas and aspects of cortical connectivity* (Vol. 5, pp. 447–466). New York: Plenum.
- Graybiel, A. M. (1990). Neurotransmitters and neuromodulators in the basal ganglia. *Neuroscience*, *13*, 244–254.
- Graybiel, A. M., Aosaki, T., Flaherty, A. W., & Kimura, M. (1994). The basal ganglia and adaptive motor control. *Science*, *265*, 1826–1831.
- Hebb, D. O. (1949). *The organization of behavior*. New York: Wiley.
- Hedreen, J. C., & DeLong, M. R. (1991). Organization of striatopallidal, striatonigral, and nigrostriatal projections in the macaque. *Journal of Comparative Neurology*, *304*, 569–595.
- Hoover, J. E., & Strick, P. L. (1993). Multiple output channels in the basal ganglia. *Science*, *259*, 819–821.
- 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. 249–270). Cambridge, MA: MIT Press.
- Iacono, R. P., Shima, F., Lonser, R. R., Kuniyoshi, S., Maeda, G., & Yamada, S. (1995). The results, indications, and physiology of posteroventral pallidotomy for patients with Parkinson's disease. *Neurosurgery*, *36*, 1118–1127.
- Kawaguchi, Y., Wilson, C. J., & Emson, P. C. (1989). Intracellular recording of identified neostriatal patch and matrix spiny cells in a slice preparation preserving cortical inputs. *Journal of Neurophysiology*, *62*, 1052–1068.
- Kombian, S. B., & Malenka, R. C. (1994). Simultaneous LTP of non-NMDA- and LTD of NMDA-receptor-mediated responses in the nucleus accumbens. *Nature*, *368*, 242–246.
- Ljungberg, T., Apicella, P., & Schultz, W. (1992). Responses of monkey dopamine neurons during learning of behavioral reactions. *Journal of Neurophysiology*, *67*, 145–163.
- Marsden, C. D., & Obeso, J. A. (1994). The functions of the basal ganglia and the paradox of stereotaxic surgery of Parkinson's disease. *Brain*, *117*, 877–897.
- Middleton, F. A., & Strick, P. L. (1994). Anatomical evidence for cerebellar and basal ganglia involvement in higher cognitive function. *Science*, *266*, 458–461.
- Mink, J. W., & Thach, W. T. (1991). Basal ganglia motor control. II. Late pallidal timing relative to movement onset and inconsistent pallidal coding of movement parameters. *Journal of Neurophysiology*, *65*, 301–329.
- Montague, P. R., Dayan, P., Person, C., & Sejnowski, T. J. (1995). Bee foraging in an uncertain environment using predictive Hebbian learning. *Nature*, *376*, 725–728.
- Montague, P. R., Dayan, P., & Sejnowski, T. J. (1996). A framework for mesencephalic dopamine systems based on predictive Hebbian learning. *Journal of Neuroscience*, *16*, 1936–1947.
- Montague, P. R., & Sejnowski, T. J. (1994). The predictive brain: Temporal coincidence and temporal order in synaptic learning mechanisms. *Learning & Memory*, *1*, 1–33.
- Mushiake, H., & Strick, P. L. (1995). Pallidal activity during sequential arm movements. *Journal of Neurophysiology*, *74*, 2754–2758.
- Nambu, A., & Llinas, R. (1994). Electrophysiology of globus pallidus neurons in vitro. *Journal of Neurophysiology*, *72*, 1127–1139.
- Owen, A. M., Roberts, A. C., Hodges, J. R., Summers, B. A., Polkey, C. E., & Robbins, T. W. (1993). Contrasting mechanisms of impaired attentional set-shifting in patients with frontal lobe damage or Parkinson's disease. *Brain*, *116*, 1159–1175.
- Parent, A. (1990). Extrinsic connections of the basal ganglia. *Trends in Neuroscience*, *13*, 254–258.
- Pascual-Leone, A., Grafman, J., Clark, K., Stewart, M., Massaquoi, S., Lou, J.-S., & Hallett, M. (1993). Procedural learning in Parkinson's disease and cerebellar degeneration. *Annals of Neurology*, *34*, 594–602.
- Polli, J. W., & Kincaid, R. L. (1994). Expression of a calmodulin-dependent phosphodiesterase isoform (PDE1B1) correlates with brain regions having extensive dopaminergic innervation. *Journal of Neuroscience*, *14*, 1251–1261.
- Ryan, L. J., & Sanders, D. J. (1993). Subthalamic nucleus lesion regularizes firing patterns in globus pallidus and substantia nigra pars reticulata neurons in rats. *Brain Research*, *626*, 327–331.
- Ryan, L. J., Sanders, D. J., & Clark, K. B. (1992). Auto- and cross-correlation analysis of subthalamic nucleus neuronal activity in neostriatal- and globus pallidus-lesioned rats. *Brain Research*, *583*, 253–261.
- Schultz, W., Apicella, P., & Ljungberg, T. (1993). Responses of monkey dopamine neurons to reward and conditioned

- stimuli during successive steps of learning a delayed response task. *Journal of Neuroscience*, *13*, 900-913.
- Schultz, W., Apicella, P., Scarnati, E., & Ljungberg, T. (1992). Neuronal activity in monkey ventral striatum related to the expectation of reward. *Journal of Neuroscience*, *12*, 4595-4610.
- Smith, Y., Wichmann, T., & DeLong, M. R. (1994a). The external pallidum and the subthalamic nucleus send convergent synaptic inputs onto single neurones in the internal pallidal segment in monkey: anatomical organization and functional significance. In G. Percheron, J. S. McKenzie, & J. Feger (Eds.), *The basal ganglia IV* (Vol. 41, pp. 51-62). New York: Plenum.
- Smith, Y., Wichmann, T., & DeLong, M. R. (1994b). Synaptic innervation of neurones in the internal pallidal segment by the subthalamic nucleus and the external pallidum in monkeys. *Journal of Comparative Neurology*, *343*, 297-318.
- Sutton, R. S. (1988). Learning to predict by the methods of temporal differences. *Machine Learning*, *3*, 9-44.
- Sutton, R. S., & Barto, A. G. (1990). Time-derivative models of Pavlovian reinforcement. In M. Gabriel & J. Moore (Eds.), *Learning and computational neuroscience: Foundations of adaptive networks* (pp. 497-538). Cambridge, MA: MIT Press.
- Swerdlow, N. R., & Koob, G. F. (1987). Dopamine, schizophrenia, mania, and depression: Toward a unified hypothesis of cortico-striato-pallido-thalamic function. *Behavioral and Brain Sciences*, *10*, 197-245.
- Thiels, E., Barrionuevo, G., & Berger, T. W. (1994). Excitatory stimulation during postsynaptic inhibition induces long-term depression in hippocampus in vivo. *Journal of Neurophysiology*, *72*, 3009-3016.
- Wagner, J. J., & Alger, B. E. (1995). GABAergic and developmental influences on homosynaptic LTD and depotentiation in rat hippocampus. *Journal of Neuroscience*, *15*, 1577-1586.
- Wichmann, T., Bergman, H., & DeLong, M. R. (1994a). The primate subthalamic nucleus. I. Functional properties in intact animals. *Journal of Neurophysiology*, *72*, 494-506.
- Wichmann, T., Bergman, H., & DeLong, M. R. (1994b). The primate subthalamic nucleus. III. Changes in motor behavior and neuronal activity in the internal pallidum induced by subthalamic inactivation in the MPTP model of parkinsonism. *Journal of Neurophysiology*, *72*, 521-530.
- Wickens, J., & Kotter, R. (1995). Cellular models of reinforcement. In J. C. Houk, J. L. Davis, & D. G. Beiser (Eds.), *Models of information processing in the basal ganglia* (pp. 187-214). Cambridge, MA: MIT Press.
- Willingham, D. B., Nissen, M. J., & Bullemer, P. (1989). On the development of procedural knowledge. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *15*, 1047-1060.
- Wilson, C. J. (1986). Postsynaptic potentials evoked in spiny neostriatal projection neurons by stimulation of ipsilateral and contralateral neocortex. *Brain Research*, *367*, 201-213.
- Wilson, C. J. (1990). Basal ganglia. In G. M. Shepherd (Ed.), *The synaptic organization of the brain* (pp. 279-316). New York: Oxford University Press.
- Wolfram, S. (1991). *Mathematica. A system for doing mathematics by computer*. (2nd ed.). Redwood City, CA: Addison-Wesley.
- Yang, X.-D., Connor, J. A., & Faber, D. S. (1994). Weak excitation and simultaneous inhibition induce long-term depression in hippocampal CA1 neurons. *Journal of Neurophysiology*, *71*, 1586-1590.