

How the Basal Ganglia Make Decisions

G. S. Berns^{1,2} and T. J. Sejnowski¹

Abstract

The primate basal ganglia are a collection of subcortical structures that have long been considered part of the extrapyramidal motor system, the part of the motor system concerned with automatic aspects of movement. Despite a large amount of data regarding their anatomy and physiology, the role of the basal ganglia in both action planning and decision making remains enigmatic. Anatomical labeling studies have suggested that the striatum receives projections from the cerebral cortex that coarsely preserves topography, and that the basal ganglia maintain a segregation of information streams (Goldman-Rakic and Selemon 1986; Hoover and Strick 1993; Parent 1990). We suggest that the connectivity of the basal ganglia is ideally suited to selecting optimal actions for given cognitive and sensory states. We demonstrate how a computational network of pools of neurons connected in the arrangement found in the basal ganglia can perform what we term a "winner-lose-all" function. The winner-lose-all mechanism refers to the fact that the neurons of the output stage of the basal ganglia, the internal segment of the globus pallidus (GPi), are tonically active and are inhibited when corresponding striatal afferents fire. Thus the GPi neuron that is selected is actually inhibited because it loses rather than wins the competition. Diffuse excitatory projections from the subthalamic nucleus prevent all but the winning pallidal neuron pool from being inhibited. Because the thalamic targets of the GPi projection in turn feed back to the approximate cortical area of the originating afferent, this cortical-subcortical loop is ideally suited not only for the aforementioned action-selection, but also for the generation of sequences appropriate for given cortical states. We demonstrate how the circuitry of the basal ganglia can learn to select the best action for different cortical states and how the feedback representation of the action-selection leads to the generation of sequences of actions.

¹Howard Hughes Medical Institute, Computational Neurobiology Laboratory, Salk Institute for Biological Studies, P.O. Box 85800, San Diego, CA 92186-5800, USA

²Present Address: Western Psychiatric Institute and Clinic, University of Pittsburgh Medical Center, 3811 O'Hara Street, Pittsburgh, PA 15213, USA

Introduction

Decision-making occurs at several levels within the central nervous system. At the highest level, the individual consciously uses both past experience and future predictions to choose an action. At the lowest level, decisions are made that do not reach consciousness until after actions are performed, and the individual is unaware of the process leading up to the particular choice. We propose that the function of a large collection of subcortical structures, the basal ganglia, is to perform automatic decision-making. The basal ganglia have classically been considered primarily part of the extrapyramidal motor system, that is, the part of the motor system concerned with automatic movement, but a wealth of new data now supports an expanded role for the basal ganglia that includes an analogous function for cognitive processes.

The basal ganglia are a collection of subcortical structures whose anatomy and physiology have been characterized in great detail. The input stage, the striatum, receives a diverse input from virtually the entire neocortex. The striatum, which itself is subdivided into the caudate and putamen, preserves the topography of the glutamatergic, excitatory afferents from the cortex. Within the striatum, a high degree of modularity exists so that discrete compartments, termed *matrisomes* and *striosomes*, have interconnectivity within compartments, but not between them. This modularity will turn out to be crucial to the model we propose. The striatal projection neurons, which are GABAergic and inhibitory, project to the globus pallidus, which itself is comprised of an internal and external segment. *Striosomes* also project reciprocally to the substantia nigra, the brain's primary source of dopamine. The external segment of the globus pallidus (GPe) projects via GABAergic inhibitory neurons almost exclusively to the subthalamic nucleus (STN). The STN also receives an excitatory cortical input. The STN projects via diffuse excitatory neurons to the internal segment of the globus pallidus (GPi). Thus the GPi receives an inhibitory projection from the striatum and an excitatory projection from the STN. These opposing influences in the GPi will be central to our proposed mechanism of decision-making in the basal ganglia. The GPi, which is the output stage of the basal ganglia, projects via GABAergic neurons to the ventrolateral thalamus, which in turn projects back to the cortex, though not necessarily to the same neurons that formed the original loop.

The anatomy of the basal ganglia is unique in the brain in that 80% of the neurons are inhibitory, with γ -aminobutyric acid (GABA) as the main neurotransmitter. In particular, the striato-pallido-thalamic pathway is comprised of two GABAergic neurons in series, provoking the question of what the functional advantage of such an arrangement is over a monosynaptic excitatory pathway. Some degree of cortical topography is preserved throughout all the structures of the basal ganglia, with the exception of the substantia nigra (Parent 1990). This finding originally gave rise to the idea of multiple parallel streams of information circulating through the basal ganglia and back to the cortex without any significant integration (Alexander and Crutcher 1990; Alexander et al. 1986). However, the caudate and putamen receive afferents from diverse cortical areas, and then

project in approximately a 1000:1 ratio to the globus pallidus (Wilson 1990). These two levels of convergence indicate either a substantial amount of integration or the multiplexed use of a structure by separate cortical areas. If integration is occurring, it most likely is computed in functionally discrete compartments. The existence of striosomes and matrisomes, in which virtually no neurons cross compartment boundaries, supports the idea of information integration, but only within defined regions corresponding to a particular action (Flaherty and Graybiel 1994; Graybiel 1990). The original hypothesis for parallel loops has now been modified, based on transneuronal retrograde tracing, to one of multiple segregated output channels (Hoover and Strick 1993). We took as our starting point the evidence for segregated channels corresponding to different possible actions, and show how the circuitry of the basal ganglia can effectively integrate several sensory and cognitive modalities to select one "best" action at any given moment.

Model

We modeled the basal ganglia as groups of simplified units interconnected according to the various divisions of the structure. Each unit corresponded to a locally distributed representation of some function, perhaps corresponding to a single matrisome in the striatum. These units were assumed to be functionally segregated and topographically related to the cortical targets that implemented the actions. The output stage of the basal ganglia, the GPi, is known to be almost wholly GABAergic and tonically active. Thus in our model the GPi tonically inhibits the target thalamic nuclei (VLm, VLpc/mc, VLo, CM). Because these nuclei 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. Given that the GPi neurons must be tonically active in order for one information stream to be selected, one GPi unit must be turned off, thus allowing for disinhibition of the corresponding thalamic target. In connectionist terminology, this is thought of as a "winner-take-all" mechanism, but here it is more accurately termed a "winner-lose-all" because the selected GPi target turns off. Because the GPi unit must be inhibited, the afferent projections to the GPi, from the neostriatum, must themselves be inhibitory.

To implement the winner-lose-all mechanism, only one unit should be allowed to become inhibited. This was achieved by lateral excitation from the STN. As shown in Figure 1, the striatal, GABAergic neurons project in a parallel fashion to the GPi, but the striatal neurons also project in a convergent fashion to the GPe. The GPe, which is also tonically active, inhibits the STN. The first striatal neuron to reach firing threshold disinhibits its corresponding thalamic target (via the GPi), but it also disinhibits the STN (via the GPe). The firing of the STN excites a larger group of GPi neurons, thus preventing any other streams from being disinhibited. Because there was an extra delay associated with traversing the GPe/STN route, the so-called indirect pathway (Alexander and Crutcher 1990), any stream to be selected must inhibit the corresponding GPi target before the STN fires.

project in approximately a 1000:1 ratio to the globus pallidus (Wilson 1990). These two levels of convergence indicate either a substantial amount of integration or the multiplexed use of a structure by separate cortical areas. If integration is occurring, it most likely is computed in functionally discrete compartments. The existence of striosomes and matrisomes, in which virtually no neurons cross compartment boundaries, supports the idea of information integration, but only within defined regions corresponding to a particular action (Flaherty and Graybiel 1994; Graybiel 1990). The original hypothesis for parallel loops has now been modified, based on transneuronal retrograde tracing, to one of multiple segregated output channels (Hoover and Strick 1993). We took as our starting point the evidence for segregated channels corresponding to different possible actions, and show how the circuitry of the basal ganglia can effectively integrate several sensory and cognitive modalities to select one "best" action at any given moment.

Model

We modeled the basal ganglia as groups of simplified units interconnected according to the various divisions of the structure. Each unit corresponded to a locally distributed representation of some function, perhaps corresponding to a single matrisome in the striatum. These units were assumed to be functionally segregated and topographically related to the cortical targets that implemented the actions. The output stage of the basal ganglia, the GPi, is known to be almost wholly GABAergic and tonically active. Thus in our model the GPi tonically inhibits the target thalamic nuclei (VLm, VLpc/mc, VLo, CM). Because these nuclei 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. Given that the GPi neurons must be tonically active in order for one information stream to be selected, one GPi unit must be turned off, thus allowing for disinhibition of the corresponding thalamic target. In connectionist terminology, this is thought of as a "winner-take-all" mechanism, but here it is more accurately termed a "winner-lose-all" because the selected GPi target turns off. Because the GPi unit must be inhibited, the afferent projections to the GPi, from the neostriatum, must themselves be inhibitory.

To implement the winner-lose-all mechanism, only one unit should be allowed to become inhibited. This was achieved by lateral excitation from the STN. As shown in Figure 1, the striatal, GABAergic neurons project in a parallel fashion to the GPi, but the striatal neurons also project in a convergent fashion to the GPe. The GPe, which is also tonically active, inhibits the STN. The first striatal neuron to reach firing threshold disinhibits its corresponding thalamic target (via the GPi), but it also disinhibits the STN (via the GPe). The firing of the STN excites a larger group of GPi neurons, thus preventing any other streams from being disinhibited. Because there was an extra delay associated with traversing the GPe/STN route, the so-called indirect pathway (Alexander and Crutcher 1990), any stream to be selected must inhibit the corresponding GPi target before the STN fires.

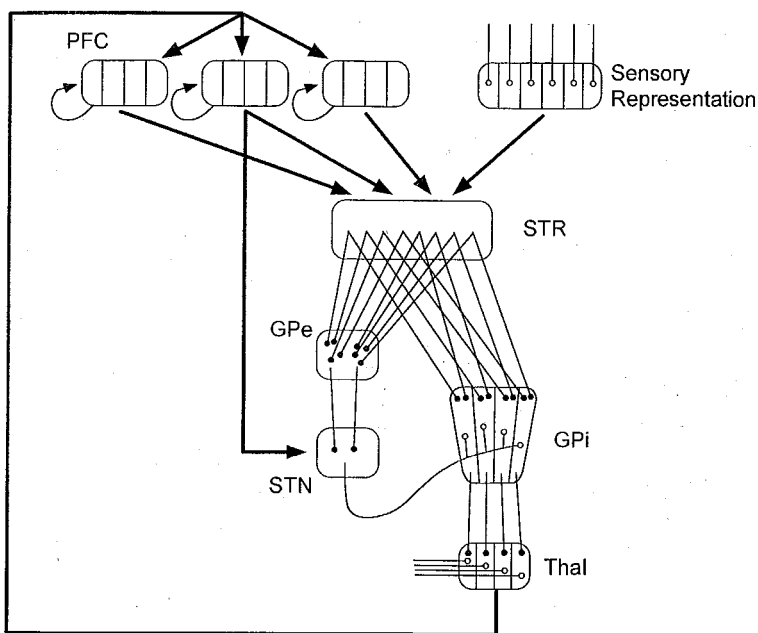
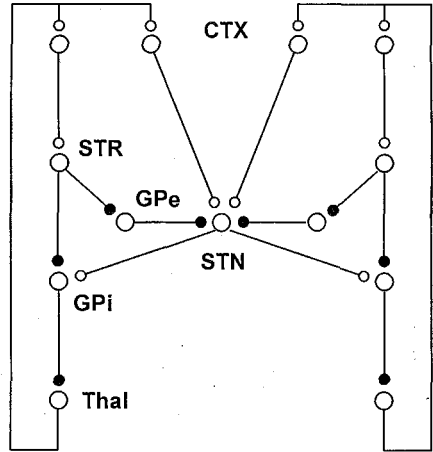


Fig. 1. Schematic diagram of the proposed winner-take-all model of basal ganglia function. A multimodal sensory and cognitive map from the cortex projects in a convergent manner to the striatum. The prefrontal cortex (PFC) is represented as a site of working memory and contains copies of the thalamic output integrated over different time scales. The sensory representation roughly represents the parietal cortex and sends a processed representation of the environment to the striatum. The striatum (STR) sends two projections: a highly convergent one to the globus pallidus externus (GPe) and a moderately convergent one to the globus pallidus internus (GPi). The first pool of striatal neurons to reach firing threshold inhibits its target in the GPi. Inhibition of the GPi neuron leads to disinhibition of the corresponding thalamic neuron (Thal) and consequent gating of ascending information to the corresponding cortical motor neurons. Simultaneously, inhibition of the GPe neuron leads to disinhibition of the subthalamic (STN) neuron, which then diffusely excites the other GPi neurons. The diffuse excitation prevents all but the first action from being selected. The cortico-STN pathway acts as negative feedback to prevent ongoing activity within the loop, ensuring that selected actions are turned off

The STN receives a prominent excitatory projection from the cortex. We propose that the function of this pathway is to inhibit actions that have recently been selected. As shown in Figure 2, the STN receives both striatal and cortical projections from two segregated loops. In the context of each individual loop, the recurrent thalamo-cortico-STN path acts to inhibit preexisting activity (note that the loop has only one inhibitory neuron, and thus it acts as negative feedback). In the context of multiple loops, the STN also serves to inhibit competing actions as described above.

The prefrontal cortex (PFC) assumes a prominent role in our model. One of the PFC's functions is thought to be related to working memory (Fuster 1993). As illustrated in Figure 1, we modeled the PFC as representing multiple copies of the

Fig. 2. Diagram of two segregated cortical-subcortical loops. The STN diffusely excites the GPi units, which is necessary for the winner-lose-all, or action-selection function. It also acts as the focal point for negative feedback within each loop, preventing chaotic oscillatory activity which would lead to either tremor in a motor loop or persistent thoughts within a cognitive loop



thalamus, in effect, multiple copies of the thalamic representation of the winning action. Each copy, however, represents an integration of activity on a different time scale. Thus the PFC contains several representations of previous winning actions, some representing recent activity and others representing the average activity over a longer time period. By virtue of the PFC-striatal projection, the decision-making function incorporates information about previously selected actions and thus can lead to the production of sequences of winning actions.

When a winner is selected, the action is assumed to be implemented through neurons in the cortical motor areas. The action results in some external event so that the model receives feedback regarding the appropriateness of the action performed. This feedback results in affective significance such that areas of the brain like the limbic cortex or the hypothalamus respond in an appropriate manner. Signals from these affective areas are presumed to reach several of the monoamine systems, which then release modulatory neurotransmitters. Montague et al. (1994) proposed a model of reinforcement learning in which the diffuse monoamine systems, especially the dopamine system, could modify synaptic strengths. They suggested that dopamine is released in response to deviations from learned predictions of future reward. Postulating that a diffuse neurotransmitter such as dopamine facilitates the change of synaptic strengths, they then demonstrated a plausible mechanism by which extrinsic rewards and penalties could be translated into the learning of specific behaviors. The primary source of dopamine, a neurotransmitter whose release is known to be closely linked to reward-driven behavior, is located in the ventral tegmental area (VTA), a structure also known to have intimate connections with the basal ganglia. Because of the close relationship with the action-selection model described above, we modeled the effect of dopamine as that of a neuromodulator that changes cortico-striatal synapse strength in close accordance with the Montague-Dayan-Sejnowski model (1994).




Fig. 3. Schematic diagram of a striosome. The striosome computes a temporal difference of its inputs, which represent the predicted future reward associated with a particular action. The temporal difference estimates the correction to an earlier prediction, i.e., it approximates the error in the prior estimate. The striosome projects to the ventral tegmental area (VTA) and substantia nigra (SNc), where the temporal difference is combined with a representation of the actual reward received. The difference in these two inputs is diffusely projected back to the striatum via dopamine neurons to modify cortico-striatal synapses

The dopamine signal from the VTA represents a prediction error for all future rewards based on the current state of the brain. This pool of neurons receives a projection from the limbic system, which conveys the instantaneous reward or affective significance of an event. Another afferent comes from the patch component of the striatum. This striosome computes a modified temporal difference (Barto et al. 1990; Sutton 1988; Sutton and Barto 1990) of the expected reward associated with the action selected by the winner-lose-all mechanism (Fig. 3). There are two possibilities for conveying the information to the striatal patch about the selected action. One possibility is through the cortex; the other possibility, which we have used, relies upon the thalamo-striatal pathway, and in particular, the intralaminar nuclei.

Simulations

Each unit was modeled in terms of its firing rate. At each time step, a function was applied to each unit that took the sum of all its respective inputs, each of which represented the firing rate of an afferent unit multiplied by a factor proportional to the synaptic weight. For example,

$$y_i = g \left[\sum_j w_{ij} x_j \right] \quad (1)$$

where y_i is the output firing rate of cell i , w_{ij} is the synaptic strength from cell j to cell i , x_j is the firing rate of input cell j , and g is a sigmoidal function that has both upper and lower saturations so that a cell can neither have a negative firing rate nor exceed a maximum rate (see Fig. 4). Explicitly, the sigmoid function is given by:

$$g(x) = \frac{1}{1 + e^{-\text{gain}(x+\text{bias})}} \quad (2)$$

where *gain* controls the slope of the inflection point and *bias* controls the shift. Table 1 provides the values for *gain* and *bias* for each of the units in the model.

The activity of each unit in the model was computed every millisecond, which corresponded to one time step. At each time step, activity was computed sequentially from the cortex to the thalamus. The input to the cortical units was




Fig. 4. General form of the input-output function of model units. The output represents firing rate. The input represents a sum of the afferent firing rates, but with each afferent weighted by the strength of its synapse to the unit. The sigmoidal function both prevents negative firing rates and allows for a maximum firing rate. The slope of the inflection point is determined by *gain* in Equation 2, and the location of the inflection point by *bias*. Here, *gain* = 4 and *bias* = -0.5

Table 1. Values for *gain* and *bias* in Equation 2 for the different synapses in the model

Synapse ^a	Gain	Bias
CTX → STN	+4	-0.5
STR → STN	+4	-0.5
STR → GPi	-4	-0.25
STN → GPi	+8	-0.25
GPi → Thal	-8	-0.5

^a CTX, cortex; STR, striatum; STN, subthalamic neuron; GPi, internal segment of the globus pallidus; Thal, thalamus.

either an external input at the beginning of a run or the activity of the corresponding thalamic unit, delayed by 100 ms. When the cortico-striatal pathway is stimulated *in vitro*, an EPSP with irregular spiking is observed in the striatum, which is followed by another EPSP and spike 100 ms later (Wilson 1990), suggesting a 100 ms transit time for the whole loop. In order to model working memory in the PFC, the cortical units additionally had 20 ms time constants associated with their activity, allowing for short-term storage. The striatal activity was then computed based on the cortical activity. The GPi activity was computed by the sum of the striatal and STN inputs, with the appropriate signs, but the STN input to the GPi was delayed by 10 ms to account for the extra synapses in the indirect pathway. The output of the GPi was then used as the input to the thalamus.

Two types of model were examined. One model represented a single recurrent loop and the other represented two competing loops. In the single-loop model, we studied the role of the STN in stopping ongoing activity. In the two-loop model, we studied both the role of the STN in inhibiting competing actions and the sensitivity to firing thresholds in the striatum in relationship to other parameters of the model.

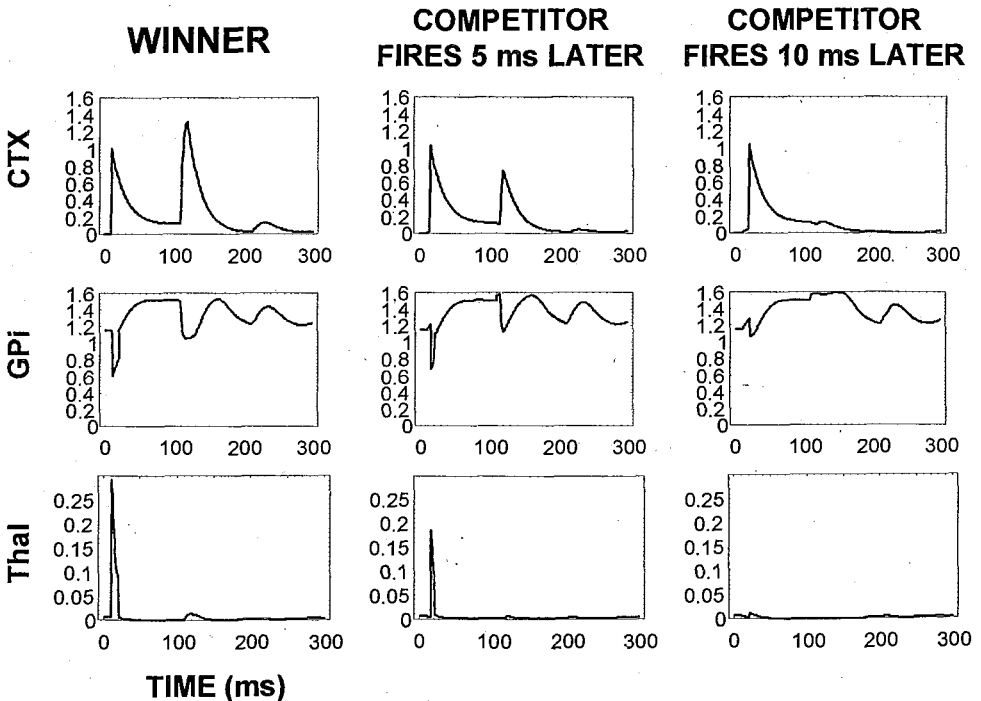


Fig. 5. Activity of the cortical (CTX), globus pallidus (GPi), and thalamic units (Thal) of two competing actions. The winning action (first column) is the first striatal unit to fire (at $t = 0$). There is a large thalamic discharge. A competing action is shown firing at both 5 and 10 ms after the first one. For longer delays, the transient decrease in GPi activity becomes progressively less, as does the corresponding thalamic disinhibition. This occurs because the winning action has already begun to diffusely excite the GPi units and thus prevents subsequent disinhibition of corresponding thalamic units

As an example of the winner-lose-all function, the efficacy of the STN in inhibiting competing actions is shown in Figure 5. The winning action corresponds to the striatal unit that fires first and results in the inhibition of the GPi unit, which in turn disinhibits the corresponding thalamic unit. A competing action in the two-loop model (Fig. 2) reaches firing threshold several milliseconds after the first one (Δt). When the second striatal unit fires, it has varying effects on the corresponding thalamic target because the first action has already started to diffusely excite the GPi units. When the second striatal unit fires at both 5 ms and 10 ms after the first, a markedly attenuated response is observed in its thalamic target.

In the context of a single recurrent pathway from cortex to striatum to globus pallidus to thalamus and back to cortex, the STN played a crucial role in preventing chaotic oscillations. As shown in Figure 6, the activity within a single loop without an STN unit was characterized by oscillations of activity at a frequency corresponding to the delay associated with the thalamo-cortex-striatum pathway (approximately 100 ms or 10 Hz). Without the STN unit, the presence of an initial stimulus led to oscillatory activity of chaotic amplitude, and in the absence of an initial stimulus, the loop activity was chaotic, but confined to a single basin of attraction (not shown). The addition of the STN unit effectively added negative feedback to the loop and prevented the ongoing oscillations after an initial stimulus. In effect, the STN acts as a brake following action selection.

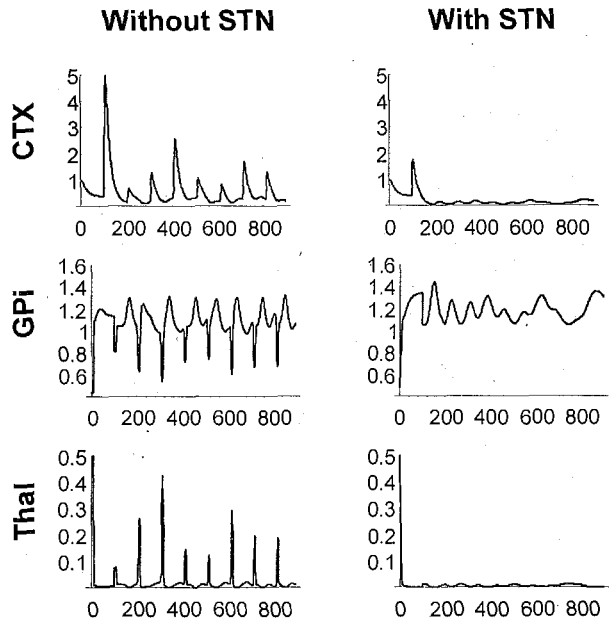


Fig. 6. Activity of cortical (CTX), globus pallidus (GPi), and thalamic units (Thal) with and without an STN unit in a single circuit. Without the STN, an initial stimulus to the cortical unit results in oscillations of chaotic amplitude throughout the circuit. With the STN, an initial cortical stimulus is subsequently inhibited by virtue of the GPi excitation. This prevents further periodic disinhibition of the thalamus (note that there is only one initial thalamic discharge)

Discussion

The proposal that the basal ganglia aid in decision-making was based largely on the observation that, given the existing anatomy of the basal ganglia, the connectivity of excitatory and inhibitory neurons is the simplest possible way to construct a winner-lose-all circuit. This proposed function is a fundamental prediction of our model that needs to be tested. The primary assumption for the winner-lose-all mechanism is the existence of streams of information that remain segregated from striatum to thalamus. We have 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 and Crutcher 1990; Alexander et al. 1986; Goldman-Rakic and 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 and Graybiel 1994; Hedreen and Delong 1991). However, our model is consistent with this convergence. We modeled the striatum as an input stage in which diverse areas of 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 and Strick 1993).

To perform a winner-lose-all function, it was necessary to have a diffuse input to several neuron pools. We hypothesize that the "indirect" pathway (Alexander and Crutcher 1990) through the GPe and STN performs this function. Anatomically consistent with this hypothesis is the finding that the subthalamic inputs to the GPi are more diffuse than the striatal inputs (Hazrati and Parent 1992; Parent and Hazrati 1993). At the cellular level, the subthalamopallidal fibers appear as plexuses with varicosities, thus suggesting a diffuse excitatory function. A further requirement of our model is that subthalamic excitation overrides striatal inhibition. Hazrati and Parent (1992; 1993) reported that there is indeed a tendency for the subthalamic inputs to terminate proximally on the GPi neurons, whereas the striatal inputs terminate more distally in the dendritic tree. Such an arrangement would be consistent with excitation overriding inhibition.

The results of the model suggest that the efficacy of the winner-lose-all mechanism can be modulated by changing both the striatal resting membrane potential and the firing threshold. The resting membrane potential of both matrix and patch cells has been reported to be approximately -70 mV (Kawaguchi et al. 1989);

however, long-lasting subthreshold depolarizations are commonly observed during which spikes are irregularly initiated (Wilson 1986). Thus the effective difference between resting and threshold potentials may be quite small. Conversely, the large subthreshold depolarizations may be a result of widespread synchrony in the striatal afferents. This latter hypothesis would be more consistent with our model, which requires that several cortical neurons fire simultaneously to bring a striatal cell to threshold.

Lesions of the basal ganglia are known to produce various movement disorders. Lesions of the subthalamic nucleus cause a condition termed hemiballismus, which is characterized by violent hemi-body movements. In our model, loss of the STN resulted in total failure of the winner-lose-all mechanism. Without any lateral excitation, all the pallidal targets that correspond to potential actions can be disinhibited. This would manifest itself perhaps as moving a whole arm when the intention was to move a finger. The presence of chaotic oscillations in our model without the STN suggests that STN lesions might lead to an inability to stop ongoing movements, which would recruit more control activity than intended.

If our model is to perform tasks more complex than simple winner selection, it must have the capability of learning, i.e., altering synaptic connections. Although they do not appear explicitly in the winner-lose-all model, the dopamine-containing neurons of the ventral tegmental area and substantia nigra are known to be involved in the motivational aspects of action (Ljungberg et al. 1992; Schultz et al. 1992). We propose that dopamine acts in the striatum to transiently change the cortico-striatal mapping, possibly by altering levels of intracellular calcium or cAMP (Garcia-Munoz et al. 1991; Kombian and Malenka 1994). 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 et al. 1992; Kombian and Malenka 1994). Dopamine is known to act on intracellular calcium (Cooper et al. 1991), and the areas of the brain with the highest calmodulin-dependent phosphodiesterase concentration correspond to those areas with heavy dopaminergic innervation (Polli and Kincaid 1994). It may be that dopamine modulates the LTP/LTD behavior of striatal cells by altering the concentration of intracellular calcium. In this manner, dopamine would act as a postsynaptic neuromodulator, but there is also evidence that dopamine acts presynaptically to modulate corticostriatal synaptic strength (Cameron and Williams 1993; Garcia-Munoz et al. 1991). Dopamine deficits, such as that seen in Parkinson's disease, are predicted to cause an impairment at the synaptic level so that cortico-striatal maps cannot be changed.

Combining the proposed decision-making function of the basal ganglia with that of dopamine in reinforcement-driven learning, we have used this model to simulate behavior in several neuropsychological tests. In a model of risk-avoidance, termed the Multiarmed Bandit, the task is to sample several decks of cards, each deck having both a different mean reward and different variance of reward.

The goal is to discover the best deck, but depending upon the consequences of negative rewards, it may be more advantageous to only sample a deck with low variance. Using dopamine to modify cortico-striatal synaptic strength, the model learns to choose the deck with the highest mean reward, but this can be shifted to low variance decks if negative rewards carry high affective significance. We have also modeled the Wisconsin Card Sorting Test, a test in which cards must be sorted according to shape, color, or number. The subject must not only discover the rules by which to sort, but also be able to change these rules. Our model has accurately simulated human behavior on this task. When noise is added to the system, the model makes errors of the failure-to-maintain-set type. Deficits of dopamine are seen to inhibit the ability to rapidly change cortico-striatal synaptic strength, and the model tends to perseverate on categories. Finally, we have simulated the learning of sequences of actions. Through the use of the recurrent thalamo-cortical connections to prefrontal cortex, the model learns to weight short-term memory traces of previous actions. By first training the model with an explicit external sequence to perform, associations are made with short-term memory so that eventually the model generalizes from performing action-selection for each point in the sequence to performing a sequence-selection when the starting point of the sequence is given.

We have proposed a systems-level model of the basal ganglia that attempts to bridge the gap between anatomy and the function of automatic decision-making. The winner-lose-all model makes specific predictions at both the cellular and behavioral levels. It predicts that lesions that prolong the conduction time through the GPe/STN pathway can be counteracted by either raising the striatal firing threshold or by hyperpolarizing striatal cells. Behaviorally, these lesions should manifest both as an inability to choose one action and as an inability to stop ongoing actions. Although these ideas are more concretely conceptualized in terms of motor function, the massive connections to the frontal cortex suggest an analogous function for cognitive actions. Thus lesions of the basal ganglia are also predicted to lead to both the inability to concentrate on one thought and the inability to stop a given train of thought.

Acknowledgments. We thank P. Read Montague and Peter Dayan for helpful discussion on reinforcement learning, Antonio Damasio for comments regarding the amygdala and its role in the motivational aspects of learning, and Francis Crick for discussion during the development of this model and comments on this paper.

References

- Alexander GE, Crutcher MD (1990) Functional architecture of basal ganglia circuits: neural substrates of parallel processing. *TINS* 13:266-271
- Alexander GE, DeLong MR, Strick PL (1986) Parallel organization of functionally segregated circuits linking basal ganglia and cortex. *Ann Rev Neurosci* 9:357-381
- Barto AG, Sutton RS, Watkins CJH (1990) Learning and sequential decision making. In: Gabriel M, Moore J (eds) *Learning and computational neuroscience: foundations of adaptive networks*. Cambridge, MIT Press, 539-602

- Calabresi P, Maj R, Pisani A, Mercuri NB, Bernard G (1992) Long-term synaptic depression in the striatum: physiological and pharmacological characterization. *J Neurosci* 12:4224–4233
- Cameron DL, Williams JT (1993) Dopamine D1 receptors facilitate transmitter release. *Nature* 366:344–347
- Cooper JR, Bloom FE, Roth RH (1991) *The biochemical basis of neuropharmacology*. New York, Oxford University Press
- Flaherty AW, Graybiel AM (1994) Input-output organization of the sensorimotor striatum in the squirrel monkey. *J Neurosci* 14:599–610
- Fuster JM (1993) Frontal Lobes. *Curr Opin Neurobiol* 3:160–165
- Garcia-Munoz M, Young SJ, Groves PM (1991) Terminal excitability of the corticostriatal pathway. I. Regulation by dopamine receptor stimulation. *Brain Res* 551:195–206
- Goldman-Rakic PS, Selemon LD (1986) Topography of corticostriatal projections in nonhuman primates and implications for functional parcellation of the neostriatum. In: Jones EG, Peters A (eds) *Cerebral cortex. Sensory-motor areas and aspects of cortical connectivity*. Plenum Press, New York, 447–466
- Graybiel AM (1990) Neurotransmitters and neuromodulators in the basal ganglia. *TINS* 13:244–254
- Hazrati LN, Parent A (1992) Convergence of subthalamic and striatal efferents at pallidal level in primates: an anterograde double-labeling study with biocytin and PHA-L. *Brain Res* 569:336–340
- Hedreen JC, DeLong MR (1991) Organization of striatopallidal, striatonigral, and nigrostriatal projections in the macaque. *J Comp Neurol* 304:569–595
- Hoover JE, Strick PL (1993) Multiple output channels in the basal ganglia. *Science* 259:819–821
- Kawaguchi Y, Wilson CJ, Emson PC (1989) Intracellular recording of identified neostriatal patch and matrix spiny cells in a slice preparation preserving cortical inputs. *J Neurophysiol* 62:1052–1068
- Kombian SB, Malenka RC (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. *J Neurophysiol* 67:145–163
- Montague PR, Dayan P, Sejnowski TJ (1994) Foraging in an uncertain environment using predictive Hebbian learning. In: Cowan JD, Tesauro G, Alspecter J (eds) *Neural information processing system 6*. San Francisco, Morgan Kaufmann, 598–605
- Parent A (1990) Extrinsic connections of the basal ganglia. *TINS* 13:254–258
- Parent A, Hazrati LN (1993) Anatomical aspects of information processing in primate basal ganglia. *TINS* 16:111–116
- Polli JW, Kincaid RL (1994) Expression of a calmodulin-dependent phosphodiesterase isoform (PDE1B1) correlates with brain regions having extensive dopaminergic innervation. *J Neurosci* 14:1251–1261
- Schultz W, Apicella P, Scarnati E, Ljungberg T (1992) Neuronal activity in monkey ventral striatum related to the expectation of reward. *J Neurosci* 12:4595–4610
- Sutton RS (1988) Learning to predict by the methods of temporal differences. *Machine Learning* 3:9–44
- Sutton RS, Barto AG (1990) Time-derivative models of pavlovian reinforcement. In: Gabriel M, Moore J (eds) *Learning and computational neuroscience: foundations of adaptive networks*. Cambridge MIT Press, 497–538
- Wilson CJ (1986) Postsynaptic potentials evoked in spiny neostriatal projection neurons by stimulation of ipsilateral and contralateral neocortex. *Brain Res* 367:201–213
- Wilson CJ (1990) Basal ganglia. In: Shepherd GM (ed) *The synaptic organization of the brain*. New York, Oxford University Press, 279–316