November 2008

Conditional Routing of Information to the Neocortex: A Network Model of Basal Ganglia Function

Andrea Stocco
Carnegie Mellon University

Christian Lebiere
Carnegie Mellon University

John R. Anderson
Carnegie Mellon University

Published In
.

Running head:

CONDITIONAL ROUTING IN THE BASAL GANGLIA

Conditional routing of information to the neocortex: A network model of basal ganglia function

(Main Text: 81962 chars, 14844 words)

Andrea Stocco

Christian Lebiere

John R. Anderson

Carnegie Mellon University

Correspondence to: Andrea Stocco

Department of Psychology

Carnegie Mellon University, Baker Hall, BP 345-E

5000 Forbes Ave.,

Pittsburgh, PA, 15213

Email: stocco@cmu.edu

Phone: +1 412 268 2787

Abstract

The basal ganglia are widely regarded as playing a central role in cognition, being involved is such general functions as action selection and reinforcement learning. Here, we present a model exploring the hypothesis that the basal ganglia implement a conditional information routing circuit, where cortical signals are gated and subsequently relayed onto different pathways projecting to the prefrontal cortex. It also accounts for the existence of two distinct branches (pallidal and nigral) of the direct pathway, by assuming that they are used to carry separately content and routing information. A possible mechanism is described by which the indirect pathway can control the release of dopamine along the nigrostriatal pathway, generating a signal that the striatum can use to produce novel stimulus-response associations by internalizing cortical representations. In a series of simulations, it is shown how the model can perform simple stimulus-response tasks, develop automatic behaviors, and provide an account of impairments in Parkinon's and Huntington's disease. Finally, it is also discussed how the proposed model's mechanisms relates to production systems, a widely adopted formalism for human cognition, and how the reward signals carried by the mesolimbic pathway can be integrated to exert fine control over action selection.

Keywords

Basal Ganglia; Neural networks, Computational modeling, Procedural learning, Sequential behavior.

.

Conditional Routing of Information within the Neocortex: A Network Model of Basal Ganglia Function

The improved understanding of individual neural circuits has spawned a number of ambitious attempts to model the basic workings of the brain (Arbib, 2003; Granger, 2006; Hawkins & Blakeslee, 2004; Houk, 2005). These attempts are widely different from each other, although all of them eventually have to deal with a few common problems. One of these is how to control the transfer of information from one processing site to another. This issue is particularly crucial for achieving integration, because it enables specialized processing circuits to access their proper inputs. It is also ultimately decisive in defining cognitive control, because deciding where individual representations are delivered shapes how behavior will be carried out and which actions will be taken.

In the field of cognitive science, so-called *cognitive architectures* are probably the most ambitious examples of general-purpose, integrated cognitive systems, aimed at providing a set of primitive functions and upon which cognitive behaviors can be built upon (Anderson, 1983; Newell, 1973). Different architectures have been proposed over the years (Anderson, 1983; Just & Carpenter, 1992; Meyer & Kieras, 1997a, 1997b; Newell, 1990). In a few of them (Anderson, 2007; Just & Varma, 2007) specific mappings between computational components and brain regions have been developed, giving the architecture a biological substrate, and shedding light on the large-scale design of the brain. These modeling attempts solve the problem of how information is transferred in different ways.

In this paper, we put forward the hypothesis that the function of directing the transfer of information is mainly performed by a particular set of brain structures: the basal ganglia. A computational model is presented, which is consistent with different functional and anatomical

features of the basal ganglia, and whose circuit implements an elaborate conditional routing system.

The proposed model has a series of advantages. Its solution to the information-routing problem has a direct mapping onto brain physiology. It also gives a unified view on the functions of the basal ganglia, and provides an account of procedural learning effects that is based solely on the circuit structure and Hebbian learning. Finally, the model provides a biological justification for the use of production systems as models of information processing in the brain, reducing the gap between commonly adopted high-level characterizations of human cognition and brain circuitry.

The remaining of this paper is articulated as follows: first, a brief survey of the anatomical and functional aspects of the basal ganglia literature is given, together with a description of currently existing models, from which the present work took many ideas. Second, the conditional routing model of the basal ganglia is described in detail, explaining its architecture, how it relates to the physiology of the circuit, and how it achieves the hypothesized functions. Third, a series of simulations are provided to highlight the model's performance and learning capabilities. Finally, possible objections and alternative accounts or models of the basal ganglia are discussed.

The Basal Ganglia

The basal ganglia are a series of subcortical nuclei located in the medial part of the brain, in proximity to the thalamus. They constitute a distinctive circuit that is evolutionarily old and remarkably similar across different species (Parent, 1986; Redgrave, Prescott, & Gurney, 1999). Its main components are the striatum (composed of the caudate nucleus and the putamen), the

internal (GPi) and external (GPe) parts of the globus pallidus, the *pars reticulata* (SNr) and *compacta* (SNc) of the substantia nigra, and the sub-thalamic nucleus (STN).

The entry point of the circuit is represented by the striatum, which receives extensive projections from the entire cortex (Gerfen, 1992; Graybiel, 2000; Parent, 1986). Cortico-striatal projections are articulated into parallel, segregated loops (Alexander, DeLong, & Strick, 1986; Gerfen, 1992; Parent & Hazrati, 1995). These projections preserve much of the topological properties of their origin, with different cortical areas targeting different compartments of the striatum (Flaherty & Graybiel, 1994).

Outgoing projections from the basal ganglia return to the cortex through the thalamus, closing the loop. These projections target almost exclusively the frontal lobes, as confirmed by anatomical tracing of brains (Alexander, DeLong, & Strick, 1986; Gerfen, 1992; Graybiel, 2000; Parent, 1986), by diffusion tensor imaging (Lehericy et al., 2004) and, indirectly, by a meta-analysis of functional imaging experiments (Postuma & Dagher, 2006). Smaller pathways originating within the basal ganglia and targeting parietal and temporal lobes have also been recently described (Clower, Dum, & Strick, 2005; Middleton & Strick, 1996).

The various nuclei of the circuit are connected to each other in a complex arrangement, which is visually summarized in Figure 1. Two striking properties of this circuit need to be commented upon. The most prominent is the existence of two independent routes connecting the striatum to the thalamus. The first one is the so-called *direct pathway* that proceeds to the thalamus through the SNr and GPi. The second is the *indirect pathway* that proceeds through the GPe and STN before merging with the direct pathway on the SNr and GPi.

The second striking characteristics of the circuit is that most of internal synapses are GABAergic and, therefore, inhibitory (Bolam, Hanley, Booth, & Bevan, 2000; Packard &

Knowlton, 2002). The exceptions include the projections coming from and returning to the cortex, and those from the STN to the SNr/GPi, all of which are glutammatergic excitatory; and the projections from SNc back to the striatum, which are dopaminergic and whose effects will be described in a later section.

Cognitive Functions of the Basal Ganglia

Movement disorders are the most obvious symptoms of basal ganglia pathologies, such as Huntington's (HD) and Parkinson's Disease (PD). In the past decades, however, converging evidence from single-cell recordings, lesion studies in humans and animals, and brain imaging studies in humans have extended the involvement of the basal ganglia to a wide range of non-motor domains. These include memory (Packard & Knowlton, 2002), attention (Ravizza & Ivry, 2001; Teicher et al., 2000), perception (Brown, Schneider, & Lidsky, 1997), language (Teichmann, Dupoux, Kouider, & Bachoud-Levi, 2006; Ullman et al., 1997), and even higher-level cognitive functions, like planning (Dagher, Owen, Boecker, & Brooks, 2001), reasoning (Frank, O'Reilly, & Curran, 2006; Goel, Buchel, Frith, & Dolan, 2000), or algebraic problem solving (Anderson, 2005; Stocco & Anderson, 2008)

The various contributions of the basal ganglia to learning are probably the most important and researched. First, they are pivotal in the acquisition of non-declarative memories. Lesions studies in animals also show that basal ganglia damage disrupts the formation of stimulus-response associations, while sparing the acquisition of declarative knowledge, which is hypothesized to depend on medial temporal lobe structures (Gabrieli, 1998; Packard & Knowlton, 2002; Packard & McGaugh, 1996). Single cell recordings in animals have also shown that the acquisition of new habits results in major changes in the spiking patterns of striatal neurons (Jog, Kubota, Connolly, Hillegaart, & Graybiel, 1999). Correspondingly, patients

with either HD or PD are impaired in learning novel skills or forming new habits (Cohen & Squire, 1980; Knowlton, Mangels, & Squire, 1996). They are also impaired in probabilistic classification tasks (Knowlton, Squire, & Gluck, 1994; Seger & Cincotta, 2005), whose non-deterministic nature and complex arrangement of stimulus-outcome associations prevents from being solved by relying on declarative memory strategies alone. Brain imaging studies in humans have shown that increases in metabolic activity in the striatum are associated with the automatization of procedures, as evidenced by experiments with sequence learning paradigms (Aizenstein et al., 2004; Peigneux et al., 2000; Rauch et al., 1997; Wu & Hallett, 2005).

A number of imaging experiments have found the basal ganglia active in working memory tasks (Braver et al., 1997; Lewis, Dove, Robbins, Barker, & Owen, 2004; Owen, Doyon, Petrides, & Evans, 1996; Rypma, Prabhakaran, Desmond, Glover, & Gabrieli, 1999). Also, genetic variations causing different quantitative expressions of dopamine receptors in the striatum correlate with both behavioral performance and brain activation in working memory tasks (Zhang et al., 2007). As expected, patients with either PD or HD are impaired in tasks tapping different forms of working memory (Gabrieli, 1998; Gabrieli, Singh, Stebbins, & Goetz, 1996; Lawrence, Sahakian, & Robbins, 1998; Owen, Iddon, Hodges, Summers, & Robbins, 1997). Conversely, the administration of dopamine, which is used to ameliorate the conditions of PD patients, also enhances their working memory performance (Cooper et al., 1992).

Another major line of research on the basal ganglia deals with their prominent role in reward processing and reward-based learning. A number of landmark single-cell recording studies in monkeys have shown that the spiking rates of dopamine neurons projecting to the striatum encode different aspects of reward-related information (Schultz, Apicella, & Ljungberg, 1993; Tobler, Fiorillo, & Schultz, 2005). In particular, their activity can be predicted on the basis

of mathematical models of reward expectancy error (Schultz, 2002; Schultz, Tremblay, & Hollerman, 2000) and uncertainty (Fiorillo, Tobler, & Schultz, 2003). Similar patterns of activations were found in brain imaging studies that adopted these paradigms with healthy adults (O'Doherty, Deichmann, Critchley, & Dolan, 2002). In fact, imaging studies have provided mounting evidence that the striatum is one of the pivotal regions underlying human decision-making tasks (Montague, King-Casas, & Cohen, 2006; O'Doherty, 2004). In the case of PD, it has been shown that dopamine agonists reverse the PD patients' learning and error patterns PD in a go/no-go task (Frank, Seeberger, & O'Reilly, 2004). The recognized role of striatum in modulating reward expectations seems to scale up to more complex tasks involving social interaction (de Quervain et al., 2004).

Computational Models of the Basal Ganglia

Interest in the cognitive contributions of the basal ganglia has spawned a wide range of modeling efforts (see Gillies & Arbuthnott, 2000, for a review). A particularly successful group of models adopts a reinforcement learning technique known as the Temporal Difference (TD) algorithm (Sutton & Barto, 1998) to model the role of basal ganglia in reward-based learning. Their success rests on the striking similarity between the short-latency dopamine signal received in the striatum and the so-called effective reinforcement signal (Schultz, 2002), a term that encodes the difference between predicted and actual reward in the algorithm. TD-learning models successfully capture many features of the dopamine neurons' activity, including its initial onset in correspondence with primary reinforcers, its propagation back to conditioned stimuli through learning, and its negative response when expected reinforcers are eventually not delivered.

Despite their success, such models have a number of shortcomings. Such models are usually agnostic with respect to some important biological features of the ganglia, such as the preponderance of inhibitory connections within the circuit or the existence of two pathways (Gillies & Arbuthnott, 2000). Also, the interpretation of the short-latency dopamine has been questioned: despite largely following the pattern of the effective reinforcement signal, it seems to be triggered under other circumstances that lay beyond those predicted by TD-learning models (Redgrave & Gurney, 2006; but see Suri & Schultz, 1999, for an interesting extension of the TD model).

Models based on the TD-learning algorithm offer an elegant solution to the sequence learning problem: when an action is predictably associated with an outcome, the reinforcement signal error shifts back to the preceding action, providing a way to organize isolated responses into a single chain to obtain a reinforcer. A number of other modeling attempts, however, took a different stance on the problem of sequence learning and automatization (Beiser & Houk, 1998; Berns & Sejnowski, 1998; Dominey, Arbib, & Joseph, 1995). They vary in both their architecture and the biological features they decided to take into consideration or ignore, as well as in the way sequential steps are encoded. For instance, the models by Dominey, Arbib and Joseph (1995) and Beiser and Houk (1998), although differing in the details, ultimately rely on a sequence representation encoded in the prefrontal cortex. In the model by Berns and Sejnowski (1998), on the other hand, the sequence order is encoded internally in the striatum. This model is also noteworthy because it explicitly addresses the functional meaning of the longer path in the indirect pathway, which is used to maintain a delayed memory of the previous state. A similar function is embedded in the model we propose, and used to provide the system with the information necessary to compile two cognitive steps into a single one.

Wickens and Arbuthnott (1993) showed that sequential patterns are spontaneously produced in a neural network model of the striatum where units are laterally interconnected. The stability of the generated patterns depends on the amount of inhibitory connections between units: more inhibition results in more stability. Lateral inhibition among striatal neurons also plays a crucial role in a particularly elegant model by Gurney, Prescott & Redgrave (2001). In their model, inhibition is instrumental in suppressing most of the signals coming from a set of input channels, filtering out those that are mediated by the few active cells. This process provides a low-level basis for the crucial role of basal ganglia in selecting among competing actions. It also hints at a more global function of the striatum in adaptively gating the incoming information.

Such a gating function has been capitalized upon by some recent models of role of the basal ganglia in working memory. In Amos (2000), the striatum performs a dynamic filtering on the currently non-relevant dimensions of a stimulus, which enables a flexible update of behavior and working memory contents. Similarly, the basal ganglia model by Frank, Loughry, and O'Reilly (2001) and O'Reilly and Frank (2006) works by letting incoming information pass through to the prefrontal areas, where it can take the place of currently held representations, or filtering it out, protecting existing information from being overwritten. The integration of striatal and working memory models made it possible to reproduce human performance in complex tasks: the Wisconsin Card Sorting Test in the case of Amos (2000) and the A-X version of the Continuous Performance Task (Cohen et al., 1997) in Frank, Loughry, and O'Reilly (2001) and O'Reilly and Frank (2006).

The Conditional Routing Model

A common interpretation is that the basal ganglia perform selection and reinforcement of actions. Selection is a general computational problem of any complex living organism, and many physiological features of the basal ganglia seem appropriate to support this function (Redgrave, Prescott, & Gurney, 1999). There is considerable variation in the literature in terms of how an "action" is defined, as well as how actions are encoded and learned within the ganglia. Both these issues are important in determining whether a model is ultimately consistent with the anatomy of the circuit.

In different frameworks, actions can be represented by individual motor responses, cognitive operations, or even larger behaviors (e.g., to run away when facing of danger). Each specific definition is challenged by the need to encompass the vast range of basal ganglia contributions to human behavior. In the presented model, an action consists in routing a restricted set of representations from the incoming cortical pathways to a different corridor of the circuit, through which it will eventually reach a correspondingly different part of the brain.

In informal descriptions of the basal ganglia circuitry, it is customary to say that certain nuclei exert an inhibitory action over their output targets. Such a description, however, is only grossly correct, since the specific ways nuclei that are consecutive along a path are wired together permits certain commands to be blocked and others to be enhanced (Wickens, 1997). In the proposed model, this structural connectivity between the nuclei plays a crucial role in shaping the functional properties of the circuit, and making the conditional routing possible. Figure 2 provides an overview of the model's architecture. For visual clarity, this figure omits the internal wiring pattern of the striatum, which is illustrated in Figure 3, and the entire nigrostriatal dopaminergic pathway (including the connections to and from the SNr) which is

shown in Figure 6. An arbitrary number of cortical regions are represented in the top white box, while the large grey box contains individual components of the basal ganglia circuit.

Before describing each of its structural parts and their associated computations in detail, it is worth to briefly illustrate how the model works. Based on the incoming signals from the cortex, the striatum performs two related operations: it selects a limited number of representations from the cortical areas, and retrieves an associated destination map for each of them. The number of representations that are effectively selected is regulated by a free parameter B: in all this simulations, $B = 1$.

Once a representation and its destination have been found, they are transferred in parallel along the direct pathway. This pathway divides in two branches, consisting of striatal projections to the SNr and the GPi, respectively. From these two intermediate nuclei, both branches eventually converge on the thalamus. In the model, the selected representation travels through the intermediate station of the SNr, and from there it is broadcasted to all the compartments of the thalamus. The destination map, on the other hand, is momentarily represented in the GPi, where it is used to selectively inhibit all the thalamic compartments but the one corresponding to the intended destination. The wiring between the SNr, GPi, and the thalamus, therefore, permits the selected contented to be moved onto the appropriate compartments. From there, they will eventually be routed to the corresponding cortical area.

A common interpretation is that the indirect pathway counters the action of the direct one. In the proposed model, however, the function of routing information is achieved by the direct pathway alone, while the indirect pathway has two other functions. First, it provides a source of activation that is needed to maintain the tonic activity of neurons in the SNr and GPi. Second, and most importantly, it provides the circuit with a transient memory of the previous

state that, because of the longer path, can be maintained for another cycle while information is being routed. In particular, the indirect pathway holds a representation of the destination map used in the previous routing operation. As it will be illustrated in a later section, this internal memory can be used to acquire new skills by combining consecutive steps.

In the model, a number of nuclei are internally organized into compartments that reflect the original organization of cortical areas. Biologically, the distinction between compartments blurs progressively within the circuit (Gerfen, 1997). This is partially reflected in the model. The striatum reflect the original cortical topology entirely. The GPi, GPe and STN nuclei, maintain the organization of compartments, but with a smaller number of units per compartment, since they are only supposed to hold maps where the active units signal a target cortical destination. As it will be illustrated later, the SNr reflects the original organization in an even lesser degree, with only part of its units being divided into compartments.

Besides *B*, a number of other parameters regulate the model's structure and size. The parameter *d* determines the number of units used to encode a cortical destination, and therefore the number of units in the GPe, GPi, and STN compartments. For reasons that will be clear in the following sections, when $d > 1$ units in the destination layers can control which part of the selected content will be transferred. In the simplest case, each of these layers can be reduced to a single unit, whose activation simply determines whether a specific cortical representations will be transferred or not. By increasing the number of units one can independently control a correspondent number of portions of the selected content. Figures 2-4, for instance, illustrate a case where $d = 2$: this makes it possible to transfer separately the part of the pattern held in either the front or the bottom row of each cortical layer. Although a biological system might have specific reasons to exert finer control over the individual parts of a representation, in the

presented architecture a larger value of d is computationally equivalent to having comparable larger number of smaller layers. Therefore, it was kept equal to 1 in all the simulations.

Finally, the number of cortical compartments and the number of units in each of them are regulated by the parameters *C* and *U*, respectively. Although they do not influence the functionality of the model, these parameters determine the size and topology of the nuclei in the circuit, and therefore the richness of the information being processed (see Table 1, and Figure 2, 3, 4). The complete list of parameters used in the model is given in Table 1. Table 2 provides a summary of the different types of units used in each component, together with their input and activation functions and thresholds.

The Striatum and the Adaptive Gating of Information

The striatum is organized into parallel compartments that mirror the cortical topology: separate compartments receive afferents from different cortical areas. Their internal organization reflects the contribution of four different types of units, which we will refer to as the input, gating, content and destination group. They divide into separate groups, whose connections are detailed in Figure 3.

Input units are those receiving afferents from the corresponding cortical region, and project to their compartment's content units. Input and content layers share the same topology, so that each content layer receives a copy of the pattern held in the corresponding input group. Content and destination units represent the output system of the striatum. They project separate pathways to the homologous compartments of the substantia nigra and of the globus pallidus, respectively. Their names reflect the two kinds of information they transmit along the direct and indirect pathways, namely the selected content and its routing destination, as it will be described in the following section.

The activation of content and destination neurons is regulated by the gating neurons, which are arguably the most essential component. They receive widespread connections from the input units in all the compartments, which enable them to work as pattern detectors for the incoming cortical signals. In turn, they project to both content and destination neurons. However, while destination neurons receive long-range connections from gating units across the striatum, content neuron receive only local projections within the same compartment.

Furthermore, destination units are activated whenever the activation from the active gating units exceeds a predetermined threshold (see Table 2 for details). Content units, on the other hand, need concurrent input from the gating units as well as from the input group to be triggered. It is only thanks to the additional activating from the gating units that the input signals can pass through and be copied on the content units.

Crucially, gating units compete with each other by means of lateral inhibition, a feature that is often included in computational models of striatum (Berns & Sejnowski, 1998; Wickens & Arbuthnott, 1993). Instead of modeling this feature directly, however, a *k*-winner-takes-all procedure (O' Reilly & Munakata, 2000) is applied to each group of gating units, ensuring that only the *k* most active units will be activated. The value of *k* is a free parameter of the model.[1]

This wiring enables a complex processing cycle. When cortical signals reach the striatum, matching gating units are concurrently activated. Those who win in the reciprocal inhibition let the incoming representation pass to the content layer and concurrently activate an associated destination. As a result, content and destination are separately held in the two striatal output systems, and eventually separately transferred along the direct pathway.

[1] Note that the gating layers use a localist representation when $k = 1$, but can use distributed one when $k > 1$

Routing operations are encoded in connections between input and gating units, input and content units, and input and destination units. One convenient way of representing routing operations is by visualizing the weights of the input and output synapses of the gating units encoding a particular operation. The weights of synapses coming from the input units represent the trigger condition. The weights of the synapses to the destination and content units indicate how the selected contents will be routed. Figure 4 provides a graphical rendition of such encoding, using a randomly generated routing operation as an example. The figure consists of two matrixes, both of which represent normalized synaptic weights: brighter spots correspond to stronger weights. The right matrix depicts a gating unit's input synapses, organized vertically by compartment. The dot above the matrix indicates the compartment the gating unit itself belongs to. The left matrix is a graphical rendition of the same gating unit output projections' weights. Gating units projects to separately the content and destination neurons: the left matrix summarizes both pieces of information by positioning the weights of the content projections in the column corresponding to the compartment indicated by the target destination units.

Biologically, input units of the striatum are supposed to represent input matrisomes, while content and destination groups represent striatal output matrisomes. Consistent with the two-state nature of neurons in the output matrisomes (Graybiel & Kimura, 1995; Wilson, 1993), content an destination units are modeled as binary units, whose activation is either 0 or 1 (see Table 2). Gating units do not have an equally straightforward biological correspondence. Their behavior summarizes the contributions of different types of striatal cells, most notably those of cholinergic interneurons, whose dendritic trees crucially receive afferents from many matrisomes (Kawaguchi, 1992), and whose tonic activity is modulated by incoming dopamine signals.

The Direct Pathway and the Thalamus

The connections projecting from the input units to the SNr, together with those from the destination groups to the internal part of GPi, constitute the direct pathway (see Figure 1). The GPi is organized into compartments that mirror the organization of the destination units in the striatum (see Figures 2 and 5). The SNr, on the other hand, does neither reflect neither the original cortical topology nor the striatal organizations. Its number of compartments is constrained by the *B* parameter, and was set to 1 for all the reported simulations. Each compartment contains *U* units, each of which receives projections from each corresponding content unit in the striatal compartments. Therefore, the SNr works as a buffer, temporarily holding a representation of the content that has been selected in the striatum. Because of its structure, the SNr constitutes the bottleneck of the system, determining the amount of information that can be transferred in a single cycle.

Both the GPi and the SNr contain binary units that are tonically active until selectively depressed by the striatal inputs. When active, they exert an overwhelming inhibitory pressure over the thalamic units, preventing them from firing. Thalamic units can be activated only by a joint reduction of the inhibition from the SNr and the GPi.

The specific wiring architecture of thalamic afferents plays a functional role in the final step of the routing process. Thalamic compartments mirror the size and arrangement of striatal input units, and each of them receives afferents from the SNr. Afferents from the GPi, on the other hand, are organized in parallel projections, so that each unit projects to all the homologous thalamic units. When units in an SNr compartment are selectively depressed by striatal connections, all the targeted thalamic neurons are partly disinhibited and predisposed to fire. A latent copy of the input representation chosen by the striatum is therefore present in all the thalamic compartments. Of all these latent patterns, one is eventually activated where the GPi

inhibition is also removed, i.e. in the compartment corresponding to the active destination matrisome. As a result, the signals that have been selected in the striatum are transferred to a different compartment in the thalamus and, from there, routed back to the cortex. This process is visually represented in Figure 5.

The Indirect Pathway

The indirect pathway is constituted by the bundle of projections that originate in the striatum and proceed through the GPe and the STN before merging with the direct pathway. The GPe shares the same structure and organization of striatal afferents as the GPi. In a similar fashion, inhibitory pathways from the GPe selectively target their counterpart in STN, which also shares the same structure. Neurons in the STN provide excitatory input to both the GPi and the SNr. These two sets of projections are organized differently: while the inputs to the GPi are still one-to-one connections, each unit in the STN targets all the units in the SNr. (see Figure 2: solid lines for one-to-one connections and dotted lines for one-to-many connections). In addition, the STN also targets the dopamine neurons in the SNc. This set of projections will be described in detail in the following section.

The indirect pathway carries a destination signal that is homologous to that conveyed by the pallidal branch of the direct pathway. Because of the longer length, however, such a signal is delayed. In particular, it is assumed that its delivery occurs one cycle later, and therefore it carries the destination map of the previous operation.

This complex arrangement has a number of functions. First, and most importantly, it maintains a short-term memory of the previous destination. A later section of this paper will illustrate how this memory can be used to acquire new procedural knowledge. Second, it provides a pace-making function to the network. In fact, in the intervals between the deliveries of

consecutive signals, the STN provides the GPe, GPi, and SNr with the constant tonic input that is needed to maintain inhibition over the thalamus. This is also instrumental in cleaning up the representations temporarily held in the direct pathway, and preparing it for the execution of the next operation.

The Dopaminergic Pathway

The fourth major functional part of the model is the dopaminergic pathway, which is illustrated in Figure 6. It involves the two parts of the substantia nigra, the SNc and SNr. The SNr has been described in a previous section. The SNc is internally organized in two parallel layers, each of which is divided into *C* compartments of *d* units each, like the GPe, GPi, and STN. The top layer represents nigral interneurons. Biologically, the dendrites of such neurons reach into the SNr, connecting to the axonal projections of striatal neurons in the SNr. The model assumes that the interneurons capture the topological organization of striatal afferents, which is otherwise lost in the SNr. They also receive local inhibitory afferents from the SNr (Hajos & Greenfield, 1994). Because of their afferents, SNc interneurons can detect which striatal compartment (and, by extension, which cortical region) contained the representation that is being transferred.

The bottom layer of neurons in the SNc represents dopaminergic neurons. They receive inhibitory input from the SNc interneurons and excitatory input from the STN. Each of them projects to the striatal neurons in the corresponding compartment (Figure 6).

Through the release of dopamine, these projections play an important role in regulating on the striatal activity and on learning. One important effect of dopamine release is that it reduces the effect of lateral inhibition among striatal neurons (Centonze et al., 2002; Delgado, Sierra, Querejeta, Valdiosera, & Aceves, 2000). In the proposed model, lateral inhibition exists

only among gating units, and is regulated by the *k*-WTA algorithm. Therefore, the model translates the activity of SNc units into adjustments in the value of *k*. In particular, an increase in the activity of SNc units causes an increment in the number of gating units active in the correspondent striatal compartment. This increase is instrumental in learning novel, skilled routing operations with practice, as it will be explained in a later section.

Relation to Production Systems

The model described in this paper can be seen as a neural implementation of a simple production system. Production systems are computational systems that rely on the execution of condition-action rules, which consists in a series of operations (the "action") to be applied when the predefined circumstances (the "condition") are verified. This behavior is not dependent on an explicit control structure that must be provided in advance, and is simply shaped by the representations they work upon. This makes production systems flexible and versatile, and, since Newell's (1973) seminal paper, they have been frequently adopted as a model of human cognition. In our network, production rules correspond to the individual routing operations learned by the model; the condition part corresponds to the cortical patterns detected by the gating units; and the action part corresponds to the performed routing operations.

Part of the flexibility of production system originates from the use of variables in the production rules, which makes them applicable to different representations. However, variables are not easily dealt with in the neural networks. A number of procedures have been proposed over the years, like a special binding space (Touretzky & Hinton, 1988), tensor product (Smolensky, 1990), and temporal binding (Ajjanagadde & Shastri, 1991). In the presented model, an alternative solution to the variable-binding problem is provided, whereby binding a variable corresponds to the process transferring the contents of a cortical region to the SNr. The

variable is the cortical location, and its bound value is the its held representation. This solution has the advantage of being tied to the specific anatomy of a neurological circuit. In a later section, it will be shown how a variable can be transformed into a constant, a process corresponding to a particular cortical pattern being stored into the striatum as a consequence of learning.

Production systems are often adopted as general models of cognition, especially in the form of cognitive architectures (Anderson, 1983, 2007; Just, Carpenter, & Varma, 1999; Meyer & Kieras, 1997a, 1997b; Newell, 1990). Among the existing models, an identical perspective on the functions of the basal ganglia is assumed in the Adaptive Control of Thought-Rational (ACT-R) cognitive archictecture (Anderson, 2007; Anderson, Fincham, Qin, & Stocco, 2008). In ACT-R, the selection and execution of rules is managed by a specific module, which is explicitly identified with the basal ganglia, and whose temporal course of activity has been successfully employed to predict the hemodynamic response in the head of the caudate nucleus. In fact, the model presented here was explicitly developed using ACT-R's procedural module as a reference for the functional properties to ensure in the circuit. It is interesting to note that a neural network implementation was made of an earlier version of ACT-R (Lebiere & Anderson, 1993). While not explicitly addressing issues of biological plausibility, the original implementation anticipated some ideas that have been developed here, as well as in other computational models of the basal ganglia (e.g., the gating functions of the circuit, as in O'Reilly & Frank, 2006).

Model Performance

An Example

An initial example will be given that illustrates how the model performs a series of routing operations and how their execution gives rise to meaningful cognitive and motor

operations. The example is an aural discrimination task that has been used as part of a dual-task experiment by Schumacher et al. (2001) and Hazeltine, Teague, and Ivry (2002). In this task, participants responded to the presentation of a tone. Tones could have three different pitches (440, 880 and 3520Hz), to which participants had to respond "one", "two", or "three", respectively. This task was later modeled by Anderson, Taatgen, and Byrne (2005) as occurring in four stages: (a) The stimulus is initially presented and automatically encoded in an aural buffer; (b) The stimulus is used as a retrieval cue and matched with previously learned stimulus-response associations in long term-memory; (c) A vocal response is retrieved that matches the attended stimulus; (d) The retrieved vocal representation is eventually programmed as a response. The authors also argued that these steps require three different brain regions: one holding a representation of the aural input, one for retrieving information from long-term memory, and a region programming vocal commands. They can be mapped onto the auditory cortex, the lateral prefrontal cortex, and the part of the motor strip responsible for programming vocal gestures, respectively.

In the simulation, a simple basal ganglia model was generated ($C = 10$, $U = 20$, $d = 1$) and connected with ten layers representing cortical regions. Three of these regions were the ones required by the task. In the vocal and auditory regions each unit was recurrently connected to itself, so that representations would passively dissipate over time. The auditory area adopted population coding for the incoming aural stimuli. The memory region was modeled as an autoassociator, retrieving a complete pattern given a partial cue.

Figure 7 illustrates the model performs such task. The figure reads top to bottom, left to right. The panels at the four corners of the figure represent the activation of the cortical units,

divided by areas, at the four stages of task execution. Two routing operations are required to perform this task: they are represented in the two middle panels.

Panel (a), in the top-left corner, corresponds to the state of the cortex when the auditory signal is first received. The first routing operation (left middle panel) is applied at stage, and consists in transferring the representation held in the auditory buffer to the prefrontal memory region. This determines the transition to stage (b), which is represented in the bottom left corner. When the prefrontal region has received the auditory cue, it responds by producing a pattern corresponding to the retrieved association to the tone. This new state corresponds to stage (c), in the top-right corner. Notice that this transition occurs within the cortex, and the basal ganglia are not involved in it. The second routing operation (right middle panel) is triggered at this point, and routes the retrieved response to the vocal region, where it can be transformed into a vocal program. This corresponds to the final stage, illustrated in panel (d), bottom-right corner.

These routing operations are visually represented in the two middle panels in Figure 7. As in the example of Figure 4, these panels illustrate how each operation is encoded in the striatum. It can be seen that the first operation is triggered by the presence of any activity in the aural region. When firing, it activates the content units corresponding to the prefrontal regions, causing the incoming aural patterns to be copied. The second one is triggered by the presence of activation in both the prefrontal and aural regions, and predisposes the units in the vocal area. It is interesting to note that the weight matrices are initially undifferentiated, so that the gating units respond to any pattern in the monitored regions, and activate all the target content units. As such, they can be applied to any stimulus/response combination in the task, the final answer being determined by the retrieved representation. This reflects an initial stage where performance

largely relies on cortical representations: in a later section it will be shown how learning relaxes this dependence, internalizes cortical representations, and speeds up execution.

General Performance and Sequential Behavior

Having the model reproducing a particular task does not provide sufficient information on its generality as an information routing device. Two series of simulations were therefore carried out to test two desired properties of the model: The ability of selecting the correct operation when a large number of independent operations have been learned; and the ability of performing a large number of learned operations in sequence.

Both tests were performed using random binary patterns as the initial cortical representations, and varying parametrically the configuration of the model by altering the numbers of compartments (C = 3, 5, 10, 15, 20), the number of units in each compartment (U = 10, 20, 30, 40, 50), and the maximum number of gating units allowed to be active (k = 1, 2, 3). In its smallest configuration, the model counted 165 units and 1,545 synapses. In its largest, 5,100 units and 1,095,100 synapses. Each configuration was tested for 100 trials.

In the first series of simulations, the model initially learned a predetermined number of operations, and was tested with a randomly selected one. The number of operations the model can learn depends on the number of gating units and how the encoding of operations is distributed among them. In each trial, this quantity was set to U / k operations. In each trial, the conditional pattern of the randomly selected operations was presented to the model, and propagated through the circuit. The thalamic output was then compared against the desired response. An error was counted whenever units outside the destination compartment were found to be active, or when the representation in the destination compartment did not match exactly the one initially present in the source cortical area.

Figure 8 (left) shows the percentage of errors for each value of *C* and *U*, averaged over different values of *k*. It can be seen that the model exhibited a robust behavior. The error peak was 2.0%, meaning that, in the most difficult condition, the model selected the correct pattern and transferred it perfectly in 98% of the trials.

The second series of simulations was aimed at testing the model's capability for sequential behavior. In this case, routing operations were generated that could be applied in sequence, and were learned by the model in random order. As in the previous test, different configurations were generated by varying the values of *U*, *C* and *k*. The number of operations was similarly set to *U* / *k*, creating series of 3 to 50 consecutive steps.

After each step in the sequences, the activation patterns in the thalamic compartments were compared against the predicted output according the criterion previously used. The number of errors was counted and eventually normalized over the length of the entire sequence, yielding an estimate of the proportion of steps that was executed incorrectly. Figure 8 (right) shows this percentage under each tested configuration. Again, the model's behavior was robust, showing a peak of just 5.2% of incorrect operations per sequence.[2]

Model Learning

So far, routing operations have been directly encoded by means of a special algorithm. The model, however, is also able to spontaneously learn new routing patterns. Computationally, learning takes the form of simple Hebbian update of the synaptic strengths between gating units and content or destination units in the striatum. Hebbian learning algorithms constitute a reasonable approximation to the biological dynamics of synaptic potentiation (Brown, Kairiss, &

[2] Sequence performance tends to be worse than on single, randomly-selected operations for two reasons. First, the patterns of operations that could be applied in sequence tend to share much contextual information, being more similar than randomly

Keenan, 1990). In the implemented version, the weight of the synapse $w_{x,y}$ between two neurons x and y is adjusted according to the equation:

$$\Delta w_{x,y} = \begin{cases} \eta\,(1 - w_{x,y})\, a_x\, a_y & \text{if } a_x\, a_y > 0 \\ -\eta\, w_{x,y} & \text{if } a_x\, a_y \leq 0 \end{cases}$$

where η is the learning rate, and a_x and a_y are the activations of neuron x and y, respectively. The term $(1 - w_{x,y})$ was introduced to bound to the growth of the synaptic weight, which is otherwise exponential and makes the basic Hebbian rule unstable (Dayan & Abbott, 2001). The update term is negative when either of the activations is zero, which accommodates the learning rule for binary units.

Simple Hebbian update among striatal neurons refines performance in two basic ways. First, it shapes the sensitivity of gating units to the incoming cortical signals. By repeated and consistent exposure, gating units can learn to expect specific patterns in cortical areas they were initially insensitive to. In a similar way, their projections to the content units specialize in imposing a specific and well-learned representation over the content units.

Despite being an unsupervised form of learning, the Hebbian rule can be strategically directed by other signals that, by shutting down or enhancing activity in the targeted units, determine whether the strength of a synapse is decreased or increased. In the proposed model, the role of such signal is played by dopamine, a neurotransmitter known to play a crucial role in learning within the striatum (Calabresi et al., 2000; Centonze et al., 2002; Centonze et al., 1999).

The striatum receives dopamine from two major pathways, the mesolimbic pathway from the ventral tegmental area (VTA), and the nigrostrial originating from the SNc. The activity of

---

generated ones. This increases the probability of the model applying the wrong operation. Second, a mistake occurring in one of

VTA neurons has been shown to correlate with reward information and provided the foundation for the TD models of basal ganglia. The role of the SNc projections is less well understood, and our model relies on the conjecture that they carry signals instrumental for the acquisition of procedural knowledge. In fact, differently from the VTA, the nigrostriatal pathway is controlled internally by the SNr and STN projections, providing an ideal substrate for a form of learning that is driven by the internal monitoring of performance and execution.

One way learning can improve execution is by eliminating intermediate processing steps that occur in the cortex, like stages (b) and (c) in the example task (see Figure 7). These extra steps can be substituted by specialized operations that bind together the initial stimuli with their associated responses. This idea has been exploited in a number of production systems' learning algorithms, like chunking (Laird, Rosenbloom, & Newell, 1986) or compilation (Taatgen & Lee, 2003).

One hallmark of these intermediate steps is that the basal circuit harvests in one cycle a representation from the very same cortical source that was the target destination in the previous cycle. In the example, where the prefrontal is responsible for memory retrievals, a cue was initially laid down in the region and the remembered vocal response was later collected from this region. The modeled wiring of the nigral nuclei allows the detection of such circumstances.

As previously discussed, SNc dopaminergic units receive concurrent excitatory projections from the STN and inhibitory projections from the SNc interneurons—which, in turn, capture the inhibitory projections from the striatum (Figure 6). Signals from these two sources reach the dopamine neurons at the same time, so that their representations either sum up or cancel each other out. When the current content (traveling on the direct pathway) is transmitted

---

the intermediate steps frequently carries over and affect subsequent steps of the sequence, increasing the proportion of errors.

from the same compartment that was the previous targeted destination (traveling on the indirect pathway), then the target SNc units benefit from an excitatory input from the STN and a reduced inhibition from the SNr, increasing their activation. Their response coincides with dopamine release, which the model translates in a chain of effects in the corresponding compartment in the striatum.

The effect of dopamine depends on the type of receptors present on the postsynaptic unit. These receptors divide into two families, D1 and D2: while the former is associated with excitatory effects, the latter associated with inhibitory effects. They have different distribution in the striatum: The model postulates that input units have predominantly D2 receptors, while output compartments have predominantly D1 receptors. Therefore, a dopamine signal from the SNc has the effect of shutting down the active input units and increasing the activity of the active content ones. In this new configuration, only gating and content units are active at the same time. Hebbian learning, therefore, weakens the association between the gating units and the current input units, and strengthens the association between the gating units and the current output units.

Dopamine also affects the number of gating units active. In a previous section it was explained how the effects of dopamine were modeled as an increase in the parameter *k*, which allows a larger number of gating units to be active at the same time. This increase in the number of units is functional to the acquisition of procedural knowledge. In fact, while dopamine-triggered Hebbian learning also affects the weight patterns of the units that were already active, it has deeper effects on those that were activated by the dopamine burst: their association with the current cortical patterns is weaker, so that the present configuration is more easily impressed on their synaptic matrices.

As an effect these newly recruited units learn to reproduce the currently active content when presented with the current cortical pattern, minus the part corresponding to the target compartment. This is a new operation that can be applied instead of transferring and harvesting a content from the same cortical areas, omitting this step altogether. Figure 9 provides a summary of this process. In the figure, the first panel illustrates the state of the different types of units before dopamine is received. The second one details the different inhibitory ("–") or excitatory ("+") effects of dopamine on the different groups of neurons. Finally, the third panel illustrates the effects of Hebbian learning, indicating were the Hebbian rule weakens ($H\downarrow$) or strengthens ($H\uparrow$) the synaptic weights.

The simple aural-vocal task described in the previous section is useful to demonstrate the effects of learning. Initially, the mappings between tone and vocal response have to be retrieved from long-term memory (see Figure 7). This intermediate step is detected in the SNc and triggers the dopamine release and the learning process described above. In the newly recruited units, the initial aural stimulus triggers a new operation that delivers an immediate vocal response. Figure 10 shows how the task can be performed only in two stages after learning has occurred: (a) The presentation of the tone, and (b) The programming of the associated motor response. The middle panel in Figure 10 depicts the input and output weight matrixes (as in the example of Figure 4) of one unit recruited in the newly learned operation. It can be noticed how, differently from the two original operations in Figure 7, the input and output patterns are not undifferentiated but instead specific to a particular stimulus/response combination.

Parkinson's and Huntington's Disease

Computational model of the basal ganglia need to confront two signature disorders of the circuit: Parkinson's and Huntington's diseases. PD is caused by the death of dopaminergic

neurons in the SNc, which drastically reduces the dopamine supply to the striatum. Its most obvious symptoms include tremors, dystonia, postural problems, and serious difficulties in initiating voluntary movements, which eventually lead to paralysis.

Since dopaminergic neurons are explicitly modeled in the SNc, a straightforward way to reproduce PD is to simulate a damage to these units. Just like the phasic burst of dopamine that is produced by SNc units translates into an increase of the parameter *k*, so the pathological depletion of SNc neurons in PD can be straightforwardly modeled as a drastic decrease of the same parameter.

This reduction has two important consequences. The first one is that less gating units can be active at any given time, reducing the total amount of activation fed to the content units, and therefore to overcome their threshold, making them less likely to fire. As a consequence, the proper content is not propagated along the direct pathway, does not reach the thalamus and hence does not update the cortex. Figure 11 illustrates such case. In the figure, the top image represents the pattern of activation in the thalamic compartments when an operation is executed under normal condition. It can be seen that only one compartment is active, containing the one pattern that has just being transferred from the cortex and properly routed through the basal ganglia circuitry. In this case, the model was set with $k = 3$. The image on the bottom left represents the landscape of thalamic activation when *k* has been lowered (in the example, $k = 2$) to simulate PD, and the very same operation is attempted. In this example, the net input to the thalamic output is falls below the activation threshold, blocking the delivery of any content to the cortex. If, like in the final stage of the example task, the transferred pattern represents a specific motor command, then the model is basically frozen in a condition resembling parkinsonian akinesia, unable to initiate a proper movement.

The model's impairment is such that any pattern is prevented from being transferred, which is an unrealistically unfavorable condition. It must be noted, however, that the model works under simplified conditions, where most units are binary and the spreading of activation occurs in discrete steps. A more realistic approach should incorporate temporal dynamics in integrate-and-fire units (Dayan & Abbott, 2001): this would mute the transfer of some (not all) patterns, and would also produce a general slowing down of the transfer process, which can be put into correspondence with the symptom of bradykinesia.

The generality of the proposed circuit, however, implies that a dysfunction of the routing mechanism should not be limited to motor programs, but extend to other kinds of representations, causing a wide range of non-motor impairments that depend on the functions of the destination cortical regions. This is consistent with the existence of widespread cognitive deficits even at the early stages of PD (Levin, Llabre, & Weiner, 1989). Given the fact that basal ganglia projections mainly target the frontal lobe, a relationship between PD and executive function disorders should be expected, and has in fact been reported in neuropsychological studies (Muslimovic, Post, Speelman, & Schmand, 2005). The relationship between different frontal lobe functions, like working memory or set-shifting, and gating functions of the basal ganglia is also highlighted in other computational models (Amos, 2000; O'Reilly & Frank, 2006)

The disruption of dopaminergic input to the striatum and its effects on the *k* parameter have a second, important consequence on the model,: it disrupts the acquisition new procedural knowledge, which crucially relies on the modulation of dopamine release to a target striatal compartment. The inability of acquiring new skills in PD has been reported in a number of studies (Jackson, Jackson, Harrison, Henderson, & Kennard, 1995; Knowlton et al., 1996). A confirmation of these impairments comes for *in vivo* imaging studies in humans. For example,

Dagher, Owen, Boecker and Brooks (2001) compared the brain activity patterns of healthy controls and mildly affected PD patients solving a set of problems with the Tower of London puzzle (Shallice, 1982), a task often used to measure planning capabilities. Participants were chosen so that the two groups had comparable behavioral performance. Compared to controls, patients exhibited less activity in the caudate nucleus, but increased activity in the hippocampus. This suggests that PD patients compensated their lack of skill learning by relying on previously acquired declarative memories (Dagher et al., 2001).

Huntington's Disease

Huntington's Diseases is caused by the death of cells within the striatum itself. In most of the cases, the disease produces motor symptoms are complementary to those of PD, with an abundance of jerky, disordered and uncontrolled movements and tics. It might, however, also result in bradykinesia. In the model, these contradictory pattern of symptoms can be explained by the different effects of losing different functional types of striatal units.

The effect of cell loss among input, content, and destination units can be simulated directly by disabling a number of randomly selected neurons. Loss of input and content units deteriorates the quality of representations being transferred. Loss of destination units, on the other hand, causes an irreversible incapability of transferring representations to the corresponding destination.

The simulation of gating units loss, on the other hand, requires additional care. As discussed in the previous sections, gating units are assumed to inhibit each other, with the more active ones dominating and silencing the other ones. Instead of modeling it directly, however, this inhibitory pressure was simplified as a *k*-WTA procedure. Since the death of gating cells

reduces the average inhibitory pressure among them, in the simulation of HD not only were some gating units disabled, but the *k*-WTA constrained was relaxed as well.

Because of this adjustment, there are two separate consequences of gating units loss. First, when the lost neurons were part of the population encoding for a specific routing pattern, cortical signals are simply not being transferred—much like in the case of PD. In the majority of cases, however, lost cells would be part of those only partially activated by the current signals. The lack of inhibition results in these cells not being suppressed still firing phasically, causing a number of content and destination units to active at the same time in different compartments.

The bottom right image on Figure 11 illustrates such an example. In the figure, the same operation is performed that was previously used to test normal and PD conditions, only this time a number of units has been disabled and the *k*-WTA constraint has been relaxed. As a result, different, impoverished end up overlapping on many thalamic compartments. This excess of activation provides a basis for involuntary and uncontrolled movements, while the degradation of transmitted patterns accounts for the deterioration of cognitive and motor abilities.

As in the case of PD, this impasse of the circuit is general and not limited to the delivery of motor programs. Correspondingly, HD patients are affected by a number of cognitive as well as motor impairments. Many of the compromised cognitive abilities (e.g., planning, working memory, and set shifting) are related to the functions of the frontal lobes, which are in fact the main target of basal ganglia projections (Lawrence et al., 1998).

In addition to a difference in their motor symptoms, the model also predicts that PD and HD should differ in the domain of skill learning. This is because the etiology of PD directly impairs the cells in the SNc that, according to the model, are responsible for triggering the learning signal. In the case of HD, however, SNc cells are spared and the learning mechanism is

in principle intact and available. This potential dissociation, however, is counterbalanced by two factors. First, part of the skill acquisition process consists in the striatum learning specific patterns that were previously available in the cortex: in HD, the striatum itself is compromised, and the extent of its damage constrains its ability to internalize and reproduce cortical representations. Second, skill acquisition only proceeds from an initial stage where task-relevant information is properly represented in the cortical regions, (as in the simplified initial stage of the example task: Figure 7), but the delivery of proper representations to the appropriate cortical regions is also affected in case of striatal damage.

Given these considerations, one not should expect more than a difference of degrees between the learning capabilities in the two conditions. Indeed, some experimental findings (e.g., Sprengelmeyer, Canavan, Lange, & Homberg, 1995) suggests that, although HD patients do exhibit skill learning impairments, their deficits are less severe than in the case of PD patients—a fact that is even more remarkable when one considers that the cognitive effects of HD are more severe and imparing than in PD.

Discussion

This paper presented a model based on the proposition that the function and structure of the basal ganglia can be understood in term of routing information among cortical areas. The model proposed a mechanistic characterization of how dopamine interacts with striatal neurons to produce novel procedural knowledge. Such a characterization was based on simple Hebbian learning mechanisms, coupled with a series of local effects in striatal compartments that are compatible with the biological effects of dopamine.

By deciding which information is transferred to and processed by the appropriate cortical region, the basal ganglia can articulate the ongoing processing activities of the brain in an

ordered sequence. Such a function does not constitute a complete characterization of the role of the basal ganglia in human cognition. However, it is general enough to explain a number of cognitive contributions of the basal ganglia, especially in memory and higher-level cognition. Furthermore, the hypothesized function emerges naturally from the structure of the circuit.

It might be argued that the model exaggerates the contribution of basal ganglia to the structured organization of information processing within the brain. Within the cortex, cortico-cortical connections constitute the large majority of the afferents (Braitenberg & Schüz, 1991). In fact, theories of the biological properties that enable complex and recursive computations in the brain usually focus on the cortex and ignore subcortical contributions (Schneider & Chein, 2003; Treves, 2005; van der Velde & de Kamps, 2006). Thalamic connections, however, have been shown to be capable of driving the activity of cortical area, even in absence of cortical amplification of the signal (Bruno & Sakmann, 2006). Indirect evidence for the role of the basal ganglia in driving cortical activity also comes from studies of functional cortical connectivity in Parkinson's Disease. Both EEG (Moazami-Goudarzi, Sarnthein, Michels, Moukhtieva, & Jeanmonod, 2008) and MEG recordings (Stoffers et al., 2008) have found that, as the disease progresses, the synchronization among cortical regions' time courses of activity increases. This suggests that, in normal conditions, the basal ganglia do interfere with the spontaneous flow of cortical activity, a fact that is consistent with their proposed role in transferring and superimposing specific representations onto the target regions of the cortex.

A well-established fact is that the basal ganglia receive projections from all of the neocortex, but project back almost exclusively to the frontal areas (Alexander et al., 1986). Such an asymmetry is apparently incompatible with the proposed information-routing function, because it implies that most of the areas providing input to the circuit are not its possible targets.

One possible explanation is that altering the activity of the frontal lobe is sufficient by itself to bias the workings of posterior parts of the brain by means of cortico-cortical projections only, making any additional projections from the basal ganglia superfluous. Another possibility is that this uneven organization of basalganglionic projections results from the functional organization of the brain. In fact, this organization mirrors a large-scale difference in functional properties of brain regions, with the frontal parts more engaged in controlling behavior by holding task-specific representations, and the posterior parts predominantly engaged in processing or representing sensory and perceptual information. Assuming this distinction, it seems rational that the basal ganglia, while gathering information from all the brain, mainly feed gated information to the prefrontal cortex, purposely not interfering with the posterior processing of sensory information. In fact, Atallah, Frank & O'Reilly (2004) have previously argued for an architectural partition of the brain in which, based on the functional macro-organization of the cortex and its ensuing computational tradeoffs, the basal ganglia and the prefrontal cortex form a joint subsystem.

The important role played by the basal ganglia in reward-based learning was not addressed in this paper. One of the assumptions underlying the model is that the mesolimbic and the nigral dopamine pathways carry different signals. In particular, the mesolimbic pathway carries information about reward, while the nigral pathway carries information relevant to procedural learning. The fact that the SNc, differently from the VTA, receives the majority of its connections from other nuclei of the basal ganglia, puts it in an ideal position for performance monitoring, which was the focus of our modeling effort.

However, the activity of striatal neurons must be ultimately affected by both kinds of signals. Any complex biological system benefits from reward signals: they are instrumental in

shaping behavior by pruning useless or harmful actions and increasing the adoption of those that yield future biological benefits. These advantages persist even when actions consist of simple routing operations, as in the presented model. The striatal integration between the learning signal from the SNc and the reward signal from the VTA might be even more finely tuned. For example, it would desirable to use reward signals to foster procedural learning under those conditions when it is more profitable, or (alternatively), prevent learning when it can be harmful.

Since both reward and procedural learning rely on dopamine, the mechanisms that have been outlined for procedural learning should in principle underlie reward-based learning as well. A straightforward point of convergence is represented by the effect of dopamine on the model's *k* parameter. An increase in *k* results in more available gating units, which could be trained to encode the same operation that is being reinforced. If the same routing pattern is encoded in more units, it is obviously more likely to fire. Correspondingly, a decrease in dopamine would reduce the number of gating units available. As an effect, the associations between the current input and content patterns and those gating units that have been shut down would be weakened. Therefore, the encoded operation would be less likely to fire again.

In any case, a detailed characterization of how a dopamine-carried reward signal enhances the likelihood of performing specific operation, and how this mechanism eventually interacts with the presently described nigrostriatal learning mechanism, constitutes a goal for future extensions of this model.

Acknowledgments

We would like to express our gratitude to Stan Dehaene of College du France, Jon M. Fincham, Ion Juvina, and Niels Taatgen of Carnegie Mellon University, and Randall O'Reilly of the University of Colorado for their help in the developing of the model and their suggestions at different stages of the manuscript.

References

Aizenstein, H. J., Stenger, V. A., Cochran, J., Clark, K., Johnson, M., Nebes, R. D., et al. (2004). Regional brain activation during concurrent implicit and explicit sequence learning. Cerebral Cortex, 14(2), 199-208.

Ajjanagadde, V., & Shastri, L. (1991). Rules and variables in neural nets. Neural Computation, 3, 121-134.

Alexander, G. E., DeLong, M. R., & Strick, P. L. (1986). Parallel organization of functionaly segregated circuits linking basal ganglia and cortex Annual Reviews of Neuroscience, 9, 357-381.

Amos, A. (2000). A computational model of information processing in the frontal cortex and basal ganglia. Journal of Cognitive Neuroscience, 12(3), 505-519.

Anderson, J. R. (1983). The architecture of cognition. Mahwah, NJ: Lawrence Erlbaum Associates.

Anderson, J. R. (2005). Human symbol manipulation within an integrated cognitive architecture. Cognitive Science, 29(3), 313-341.

Anderson, J. R. (2007). How can the human mind occur in the physical universe? (1 ed.). New York, NY: Oxford University Press.

Anderson, J. R., Albert, M. V., & Fincham, J. M. (2005). Tracing problem solving in real time: fMRI analysis of the subject-paced Tower of Hanoi. Journal of Cognitive Neuroscience, 17(8), 1261-1274.

Anderson, J. R., Fincham, J. M., Qin, Y., & Stocco, A. (2008). A central circuit of the mind. Trends in Cognitive Sciences, 12(4), 136-143.

Anderson, J. R., Taatgen, N. A., & Byrne, M. D. (2005). Learning to achieve perfect timesharing: Architectural implications of Hazeltine, Teague, and Ivry (2002). Journal of Experimental Psychology: Human Perception and Performance, 31(4), 749-761.

Arbib, M. A. (2003). Towards a neurally-inspired computer architecture. Natural Computing, 2, 1-46.

Atallah, H. E., Frank, M. J., & O'Reilly, R. C. (2004). Hippocampus, cortex, and basal ganglia: Insights from computational models of complementary learning systems. Neurobiology of Learning and Memory, 82(3), 253-267.

Beiser, D. G., & Houk, J. C. (1998). Model of cortical-basal ganglionic processing: encoding the serial order of sensory events. Journal of Neurophysiology, 79(6), 3168-3188.

Berns, G., & Sejnowski, T. J. (1998). A computational model of how the basal ganglia produce sequences. Journal of Cognitive Neuroscience, 10, 108-121.

Bolam, J. P., Hanley, J. J., Booth, P. A., & Bevan, M. D. (2000). Synaptic organisation of the basal ganglia. Journal of Anatomy, 196 (Pt 4), 527-542.

Braitenberg, V., & Schüz, A. (1991). Anatomy of the cortex. Statistics and geometry. Heidelberg, Germany: Springer Verlag.

Braver, T. S., Cohen, J. D., Nystrom, L. E., Jonides, J., Smith, E. E., & Noll, D. C. (1997). A parametric study of prefrontal cortex involvement in human working memory. Neuroimage, 5(1), 49-62.

Brown, L. L., Schneider, J. S., & Lidsky, T. I. (1997). Sensory and cognitive functions of the basal ganglia. Current Opinion in Neurobiology, 7(2), 157-163.

Brown, T. H., Kairiss, E. W., & Keenan, C. L. (1990). Hebbian synapses: Biophysical mechanisms and algorithms. Annual Review of Neuroscience, 13, 475-511.

Bruno, R. M., & Sakmann, B. (2006). Cortex is driven by weak but synchronously active thalamocortical synapses. Science, 312(5780), 1622-1627.

Calabresi, P., Gubellini, P., Centonze, D., Picconi, B., Bernardi, G., Chergui, K., et al. (2000). Dopamine and cAMP-regulated phosphoprotein 32 kDa controls both striatal long-term depression and long-term potentiation, opposing forms of synaptic plasticity. Journal of Neuroscience, 20(22), 8443-8451.

Centonze, D., Bracci, E., Pisani, A., Gubellini, P., Bernardi, G., & Calabresi, P. (2002).

Activation of dopamine D1-like receptors excites LTS interneurons of the striatum. The European Journal of Neuroscience 15(12), 2049-2052.

Centonze, D., Gubellini, P., Picconi, B., Calabresi, P., Giacomini, P., & Bernardi, G. (1999). Unilateral dopamine denervation blocks corticostriatal LTP. Journal of Neurophysiology, 82(6), 3575-3579.

Centonze, D., Usiello, A., Gubellini, P., Pisani, A., Borrelli, E., Bernardi, G., et al. (2002). Dopamine D2 receptor-mediated inhibition of dopaminergic neurons in mice lacking D2L receptors. Neuropsychopharmacology, 27(5), 723-726.

Clower, D. M., Dum, R. P., & Strick, P. L. (2005). Basal ganglia and cerebellar inputs to 'AIP'. Cerebral Cortex, 15(7), 913-920.

Cohen, J. D., Perlstein, W. M., Braver, T. S., Nystrom, L. E., Noll, D. C., Jonides, J., et al. (1997). Temporal dynamics of brain activation during a working memory task. Nature, 386(6625), 604-608.

Cohen, N. J., & Squire, L. R. (1980). Preserved learning and retention of pattern-analyzing akill in amnesia - dissociation of Knowing How and Knowing That. Science, 210(4466), 207-210.

Cooper, J. A., Sagar, H. J., Doherty, S. M., Jordan, N., Tidswell, P., & Sullivan, E. V. (1992). Different effects of dopaminergic and anticholinergic therapies on cognitive and motor function in Parkinson's disease. A follow-up study of untreated patients. Brain, 115 (Pt 6), 1701-1725.

Dagher, A., Owen, A. M., Boecker, H., & Brooks, D. J. (2001). The role of the striatum and hippocampus in planning: A PET activation study in Parkinson's disease. Brain, 124(5), 1020-1032.

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

de Quervain, D. J., Fischbacher, U., Treyer, V., Schellhammer, M., Schnyder, U., Buck, A., et al. (2004). The neural basis of altruistic punishment. Science, 305(5688), 1254-1258.

Delgado, A., Sierra, A., Querejeta, E., Valdiosera, R. F., & Aceves, J. (2000). Inhibitory control of the GABAergic transmission in the rat neostriatum by D2 dopamine receptors. Neuroscience, 95(4), 1043-1048.

Dominey, P., Arbib, M. A., & Joseph, J. (1995). A model of corticostriatal plasticity for learning oculomotor associations and sequences. Journal of Cognitive Neuroscience, 7, 311-336.

Fiorillo, C. D., Tobler, P. N., & Schultz, W. (2003). Discrete coding of reward probability and uncertainty by dopamine neurons. Science, 299(5614), 1898-1902.

Flaherty, A. W., & Graybiel, A. M. (1994). Input-output organization of the sensorimotor striatum in the squirrel monkey. Journal of Neuroscience, 14(2), 599-610.

Frank, M. J., O'Reilly, R. C., & Curran, T. (2006). When memory fails, intuition reigns: midazolam enhances implicit inference in humans. Psychological Science, 17(8), 700-707.

Frank, M. J., Seeberger, L. C., & O'Reilly, R. C. (2004). By carrot or by stick: Cognitive reinforcement learning in parkinsonism. Science, 306(5703), 1940-1943.

Gabrieli, J. D. E. (1998). Cognitive neuroscience of human memory. Annual Review of Psychology, 49, 87-115.

Gabrieli, J. D. E., Singh, J., Stebbins, G. T., & Goetz, C. G. (1996). Reduced working memory span in Parkinson's disease: Evidence for the role of a frontostriatal system in working and strategic memory. Neuropsychology, 10(3), 322-332.

Gerfen, C. R. (1992). The Neostriatal Mosaic - Multiple Levels of Compartmental Organization in the Basal Ganglia. Annual Review of Neuroscience, 15, 285-320.

Gillies, A., & Arbuthnott, G. W. (2000). Computational models of the basal ganglia. Movement Disorders, 15(5), 762-770.

Goel, V., Buchel, C., Frith, C., & Dolan, R. J. (2000). Dissociation of mechanisms underlying syllogistic reasoning. Neuroimage, 12(5), 504-514.

Granger, R. (2006). Engines of the brain: The computational instruction set of human cognition. AI Magazine, 27(2), 15-31.

Graybiel, A. M. (2000). The basal ganglia. Current Biology, 10(14), R509-R511.

Graybiel, A. M. (2000). The basal ganglia. Current Biology, 10(14), R509-511.

Gurney, K., Prescott, T. J., & Redgrave, P. (2001). A computational model of action selection in the basal ganglia. I. A new functional anatomy. Biological Cybernetics, 84(6), 401-410.

Hajos, M., & Greenfield, S. A. (1994). Synaptic connections between pars compacta and pars reticulata neurones: electrophysiological evidence for functional modules within the substantia nigra. Brain Research, 660(2), 216-224.

Hawkins, J., & Blakeslee, S. (2004). On intelligence. New York, NY: Times Books.

Hazeltine, E., Teague, D., & Ivry, R. B. (2002). Simultaneous dual-task performance reveals parallel response selection after practice. Journal of Experimental Psychology: Human Perception and Performance, 28(3), 527-545.

Houk, J. C. (2005). Agents of the mind. Biological Cybernetics, 92(6), 427-437.

Jackson, G. M., Jackson, S. R., Harrison, J., Henderson, L., & Kennard, C. (1995). Serial reaction time learning and Parkinson's disease: evidence for a procedural learning deficit. Neuropsychologia, 33(5), 577-593.

Jog, M. S., Kubota, Y., Connolly, C. I., Hillegaart, V., & Graybiel, A. M. (1999). Building neural representations of habits. Science, 286(5445), 1745-1749.

Just, M. A., & Carpenter, P. A. (1992). A capacity theory of comprehension: Individual differences in working memory. Psychological Review, 99, 122-149.

Just, M. A., Carpenter, P. A., & Varma, S. (1999). Computational modeling of high-level cognition and brain function. Human Brain Mapping, 8(2-3), 128-136.

Just, M. A., & Varma, S. (2007). The organization of thinking: What functional brain imaging reveals about the neuroarchitecture of complex cognition. Cognitive, Affective and Behavioral Neuroscience, 7(3), 153-191.

Knowlton, B. J., Mangels, J. A., & Squire, L. R. (1996). A neostriatal habit learning system in humans. Science, 273(5280), 1399-1402.

Knowlton, B. J., Squire, L. R., & Gluck, M. A. (1994). Probabilistic classification learning in

amnesia. Learning & Memory, 1(2), 106-120.

Laird, J. E., Rosenbloom, P. S., & Newell, A. (1986). Chunking in soar: The anatomy of a general learning mechanism. Machine Learning, 1(1), 11-46.

Lawrence, A. D., Sahakian, B. J., & Robbins, T. W. (1998). Cognitive functions and corticostriatal circuits: insights from Huntington's disease. Trends in Cognitive Sciences, 2(10), 379-388.

Lebiere, C., & Anderson, J. R. (1993). A connectionist Implementation of the ACT-R production system. In L. E. Associates (Ed.), Proceedings of the Fifteenth Annual Conference of the Cognitive Science Society (pp. 635-640). Mahwah, NJ.

Lehericy, S., Ducros, M., Van de Moortele, P. F., Francois, C., Thivard, L., Poupon, C., et al. (2004). Diffusion tensor fiber tracking shows distinct corticostriatal circuits in humans. Annals of Neurology, 55(4), 522-529.

Levin, B. E., Llabre, M. M., & Weiner, W. J. (1989). Cognitive impairments associated with early Parkinson's disease. Neurology, 39(4), 557-561.

Lewis, S. J., Dove, A., Robbins, T. W., Barker, R. A., & Owen, A. M. (2004). Striatal contributions to working memory: a functional magnetic resonance imaging study in humans. Eur J Neurosci, 19(3), 755-760.

Meyer, D. E., & Kieras, D. E. (1997a). A computational theory of executive cognitive processes and multiple-task performance .1. Basic mechanisms. Psychological Review, 104(1), 3-65.

Meyer, D. E., & Kieras, D. E. (1997b). A computational theory of executive cognitive processes and multiple-task performance .2. Accounts of psychological refractory-period phenomena. Psychological Review, 104(4), 749-791.

Middleton, F. A., & Strick, P. L. (1996). The temporal lobe is a target of output from the basal ganglia. Proceedings of the National Academy of Sciences of the United States of America, 93(16), 8683-8687.

Moazami-Goudarzi, M., Sarnthein, J., Michels, L., Moukhtieva, R., & Jeanmonod, D. (2008).

Enhanced frontal low and high frequency power and synchronization in the resting EEG of parkinsonian patients. NeuroImage, 41(3), 985-997.

Montague, P. R., King-Casas, B., & Cohen, J. D. (2006). Imaging valuation models in human choice. Annual Review of Neuroscience, 29, 417-448.

Muslimovic, D., Post, B., Speelman, J. D., & Schmand, B. (2005). Cognitive profile of patients with newly diagnosed Parkinson disease. Neurology, 65(8), 1239-1245.

Newell, A. N. (1973). You can't play twenty questions with nature and win. In W. G. Chase (Ed.), Visual information processing (pp. 283-308). New York, NY: Academic Press.

Newell, A. N. (1990). Unified theories of cognition. Cambridge, MA: Harvard University Press.

O'Doherty, J. P. (2004). Reward representations and reward-related learning in the human brain: insights from neuroimaging. Current Opinion in Neurobiology, 14(6), 769-776.

O'Doherty, J. P., Deichmann, R., Critchley, H. D., & Dolan, R. J. (2002). Neural responses during anticipation of a primary taste reward. Neuron, 33(5), 815-826.

O'Reilly, R. C., & Frank, M. J. (2006). Making working memory work: a computational model of learning in the prefrontal cortex and basal ganglia. Neural Computation, 18(2), 283-328.

Owen, A. M., Doyon, J., Petrides, M., & Evans, A. C. (1996). Planning and spatial working memory: a positron emission tomography study in humans. The European Journal of Neuroscience, 8(2), 353-364.

Owen, A. M., Iddon, J. L., Hodges, J. R., Summers, B. A., & Robbins, T. W. (1997). Spatial and non-spatial working memory at different stages of Parkinson's disease. Neuropsychologia, 35(4), 519-532.

Packard, M. G., & Knowlton, B. J. (2002). Learning and memory functions of the basal ganglia. Annual Review of Neuroscience, 25, 563-593.

Packard, M. G., & Mcgaugh, J. L. (1996). Inactivation of hippocampus or caudate nucleus with lidocaine differentially affects expression of place and response learning. Neurobiology of Learning and Memory, 65(1), 65-72.

Parent, A. (1986). Comparative neurobiology of the basal ganglia. New York: J. Wiley.

Parent, A., & Hazrati, L. N. (1995). Functional anatomy of the basal ganglia. 1: The cortico-basal ganglia-thalamo-cortical loop. Brain Research Reviews, 20(1), 91-127.

Peigneux, P., Maquet, P., Meulemans, T., Destrebecqz, A., Laureys, S., Degueldre, C., et al. (2000). Striatum forever, despite sequence learning variability: a random effect analysis of PET data. Hum Brain Mapp, 10(4), 179-194.

Postuma, R. B., & Dagher, A. (2006). Basal ganglia functional connectivity based on a meta-analysis of 126 positron emission tomography and functional magnetic resonance imaging publications. Cerebral Cortex, 16(10), 1508-1521.

Rauch, S. L., Whalen, P. J., Savage, C. R., Curran, T., Kendrick, A., Brown, H. D., et al. (1997). Striatal recruitment during an implicit sequence learning task as measured by functional magnetic resonance imaging. Human Brain Mapping, 5(2), 124-132.

Ravizza, S. M., & Ivry, R. B. (2001). Comparison of the basal ganglia and cerebellum in shifting attention. Journal of Cognitive Neuroscience, 13(3), 285-297.

Redgrave, P., Prescott, T. J., & Gurney, K. (1999). The basal ganglia: A vertebrate solution to the selection problem? Neuroscience, 89(4), 1009-1023.

Rypma, B., Prabhakaran, V., Desmond, J. E., Glover, G. H., & Gabrieli, J. D. E. (1999). Load-dependent roles of frontal brain regions in the maintenance of working memory. Neuroimage, 9(2), 216-226.

Schneider, W., & Chein, J. M. (2003). Controlled & automatic processing: From mechanisms to biology. Cognitive Science, 27, 525-559.

Schultz, W. (2002). Getting formal with dopamine and reward. Neuron, 36(2), 241-263.

Schultz, W., Apicella, P., & Ljungberg, T. (1993). Responses of monkey dopamine neurons to reward and conditioned stimuli during successive steps of learning a delayed response task. Journal of Neuroscience, 13(3), 900-913.

Schultz, W., Tremblay, L., & Hollerman, J. R. (2000). Reward processing in primate orbitofrontal cortex and basal ganglia. Cerebral Cortex, 10(3), 272-284.

Schumacher, E. H., Seymour, T. L., Glass, J. M., Fencsik, D. E., Lauber, E. J., Kieras, D. E., et al. (2001). Virtually perfect time sharing in dual-task performance: uncorking the central cognitive bottleneck. Psychological Science, 12(2), 101-108.

Seger, C. A., & Cincotta, C. M. (2005). The roles of the caudate nucleus in human classification learning. Journal of Neuroscience, 25(11), 2941-2951.

Shallice, T. (1982). Specific impairments of planning. Philosophical transactions of the Royal Society of London, 298(1089), 199-209.

Smolensky, P. (1990). Tensor product variable binding and the representation of symbolic structures in connectionist systems. Artificial Intelligence, 46(1-2), 159-216.

Sprengelmeyer, R., Canavan, A. G., Lange, H. W., & Homberg, V. (1995). Associative learning in degenerative neostriatal disorders: contrasts in explicit and implicit remembering between Parkinson's and Huntington's diseases. Movement Disorders, 10(1), 51-65.

Stocco, A., & Anderson, J. R. (in press). Endogenous control and task representation: An fMRI study of algebraic problem solving. Journal of Cognitive Neuroscience.

Stoffers, D., Bosboom, J. L., Deijen, J. B., Wolters, E., Stam, C. J., & Berendse, H. W. (2008). Increased cortico-cortical functional connectivity in early-stage Parkinson's disease: An MEG study. Neuroimage, 41(2), 212-222.

Sutton, R. S., & Barto, A. G. (1998). Reinforcement learning: An introduction. Cambridge, MA: MIT Press.

Taatgen, N. A., & Lee, F. J. (2003). Production compilation: A simple mechanism to model complex skill acquisition. Human Factors, 45(1), 61-76.

Teicher, M. H., Anderson, C. M., Polcari, A., Glod, C. A., Maas, L. C., & Renshaw, P. F. (2000). Functional deficits in basal ganglia of children with attention-deficit/hyperactivity disorder shown with functional magnetic resonance imaging relaxometry. Nature Medicine, 6(4), 470-473.

Teichmann, M., Dupoux, E., Kouider, S., & Bachoud-Levi, A. C. (2006). The role of the striatum in processing language rules: Evidence from word perception in Huntington's

disease. Journal of Cognitive Neuroscience, 18(9), 1555-1569.

Tobler, P. N., Fiorillo, C. D., & Schultz, W. (2005). Adaptive coding of reward value by dopamine neurons. Science, 307(5715), 1642-1645.

Touretzky, D., & Hinton, G. H. (1988). A distributed connectionist production system. Cognitive Science, 12, 423-466.

Treves, A. (2005). Frontal latching networks: A possible neural basis for infinite recursion. Cognitive Neuropsychology, 22(3-4), 276-291.

Ullman, M. T., Corkin, S., Coppola, M., Hickok, G., Growdon, J. H., Koroshetz, W. J., et al. (1997). A neural dissociation within language: Evidence that the mental dictionary is part of declarative memory, and that grammatical rules are processed by the procedural system. Journal of Cognitive Neuroscience, 9(2), 266-276.

van der Velde, F., & de Kamps, M. (2006). Neural blackboard architectures of combinatorial structures in cognition. The Behavioral and Brain Sciences, 29(1), 37-70; discussion 70-108.

Wickens, J. R. (1997). Basal ganglia: structure and computations. Network-Computation in Neural Systems, 8(4), R77-R109.

Wu, T., & Hallett, M. (2005). A functional MRI study of automatic movements in patients with Parkinson's disease. Brain, 128(Pt 10), 2250-2259.

Zhang, Y., Bertolino, A., Fazio, L., Blasi, G., Rampino, A., Romano, R., et al. (2007). Polymorphisms in human dopamine D2 receptor gene affect gene expression, splicing, and neuronal activity during working memory. Proceedings of the National Academy of Sciences of the United States of America, 104(51), 20552-20557.

Tables

Table 1: List of the parameters that determine the model's configuration and behavior.

Parameter	Meaning
C	Number of compartments in each nucleus of the circuit.
U	Number of input, content, and gating units in each striatal compartment; number of units in each compartment of SNr and Thalamus.
d	Number of destination units in each striatal compartment; number of units in each compartment of GPe, GPi, STN and SNc.
B	Number of compartments where k gating unit are allowed to be active (e.g., number of representations that can be transferred simultaneously)
k	Number of gating units that can be simultaneously active in a compartment (proportional to the amount of dopamine)
η	Hebbian learning scaling parameter
a	Balance between the input and the gating pathways to the content layers

Table 2: Net inputs, activation functions, and thresholds for all the different types of units in the model.

Unit type		Net input (N)	Activation (a)	Threshold (T)
Striatum	Input	$N^I_{i,j} = wa^{CTX}_{i,j}$	$a^I = \begin{cases} 1 & \text{if } N > T \\ 0 & \text{otherwise} \end{cases}$	$T = 0$
	Gating	$N^G_{i,j} = \sum_i \sum_j wa^I_{i,j}$	$a^G = \tanh(N)$	$T = 0$
	Content	$N^C_{i,j} = wa^I_{i,j} + \sum_j wa^G_{i,j}$	$a^C = \begin{cases} 1 & \text{if } N > T \\ 0 & \text{otherwise} \end{cases}$	$T = 0.25$
	Destination	$N^D_{i,j} = \sum_i \sum_j wa^G_{i,j}$	$a^D = \begin{cases} 1 & \text{if } N > T \\ 0 & \text{otherwise} \end{cases}$	$T = 0.5$
GPe/GPi		$N^{GPe}_{i,j} = N^{GPi}_{i,j} = wa^D_{i,j}$	$a^{GPe} = a^{GPi} = N$	$T = -k$
SNr		$N^{SNr}_{i,j} = wa^C_{i,j}$	$a^{SNr} = N$	$T = -k$
SNc		$N^{SNc}_{i,j} = wa^{SNr}_{i,j} + \sum_i wa^{STN}_{i,j}$	$A^{STN} = \tanh(N)$	T = 0
STN		$N^{STN}_{i,j} = wa^{GPe}_{i,j}$	$a^{STN} = N$	$T = 0$
Thalamic		$N^{Th}_{i,j} = \sum_i wa^{SNr}_{i,j} + \sum_j wa^{GPe}_{i,j}$	$a^{Th} = \begin{cases} 1 & \text{if } N > T \\ 0 & \text{otherwise} \end{cases}$	$T = -C$

Figure legends

Figure 1: A simplified representation of the basal ganglia circuit and connectivity (adapted from Graybiel, 2000, and Wickens, 1997). Black arrows represent excitatory (i.e., glutammatergic or dopaminergic) projections; white arrows represent inhibitory (i.e., GABA-ergic) projections.

Figure 2: Architecture of the basal ganglia model. Black arrows represent excitatory connections. White arrows represent inhibitory connections. Solid lines represent one-to-one connections (i.e., ach unit of the source layer projects only to the corresponding unit of the target layer). Dotted arrows represent one-to-many connections (i.e, each unit of the sending layer projects to every unit of the receiving layer). The particular model depicted has the following parameters: $C = 5$, $U = 6$, $d = 2$ (see Table 1 for a detailed explanation).

Figure 3: Internal connectivity of the model's striatal component. Solid arrows represent one-to-one connections, dotted arrows represent one-to-many connections. Configuration parameters are the same as in Figure 2.

Figure 4: An illustration of how routing operations are encoded in the striatum. The figure depicts the input and output normalized weight matrixes for a gating unit encoding a randomly-generated operation. The matrix on the left depicts the incoming projections from the input units, divided by compartments. The matrix on the right shows the weights of the gating unit's projections to the content units, allocated in the column corresponding to the destination unit's compartment (i.e., in the location where the content is eventually going to be routed).

Figure 5: An example of how information is routed along the Direct Pathway. Solid arrows represent one-to-one connections, dotted arrows represent one-to-many connections. Configuration parameters are the same as in Figure 2.

Figure 6: Connectivity of the model's dopaminergic pathway. Black arrows represent excitatory

connections. White arrows represent inhibitory connections. Solid lines represent one-to-one connections. Dotted arrows represent one-to-many connections (i.e, each unit of the sending layer projects to every unit of the receiving layer). Configuration parameters are the same as in Figure 2.

Figure 7. The four stages of the aural discrimination task, as they appear in form of activation in the simulated cortical areas projecting to the striatum (*top and bottom panels*), together with a representation of the two operations as encoded in the model (*middle panels*). In the middle panels, the left matrix depicts the incoming synapses weights from the input units in the striatum to one of the gating units encoding the operation. The right matrix represents synaptic weights of projections from the gating unit to the content units. Since the output projections of a gating unit are local, this matrix was generated as the outer product of the activation vector of the homologous compartment's content units and the activation vector of the destination units.

Figure 8. *Left*: Performance of the model across multiple configurations, when multiple operations have been previously encoded; *Right:* Performance of the model across multiple configurations, when multiple operations have to be applied in sequence. The number of operations loaded in the model for each configuration was set to U / k.

Figure 9: A step-by-step rendition of how dopamine affects activity in the striatum, and how these changes of activity, coupled with Hebbian learning, are used to generate new operations. *Panel 1*: Activity of various groups of neurons in the target striatal compartment before dopamine arrival. *Panel 2*: Summary of the different effect of dopamine on the targeted input, gating and destination units. *Panel 3*: Effects of Hebbian learning on the new patterns of activity.

Figure 10: The aural discrimination task, as it appears after procedural learning has occurred. The two intermediate stages have been omitted (cfr. Figure 5), and the task only requires two

stages (*top and bottom panels*) and the newly learned routing operation (*middle panel*). Note that the operations' synaptic weights are now specific to a singular stimulus-response mapping.

Figure 11: Performance of the model under normal and impaired conditions. In the example configuration, $C = 10$, $U = 20$, and $k = 3$. (*Top*) Activation of thalamic neurons after the execution of a routing operation in a normal condition. The pattern was successfully transferred onto the first thalamic compartment. (*Bottom left*) Activation of thalamic neurons during the execution of the same operation after the k value was decreased (from $k = 3$ to $k = 2$). This condition simulates the effect of Parkinson's disease. (*Bottom right*) Activation of thalamic neurons during the execution of the same operation after the lateral inhibition between gating units was relaxed ($k = U$). This condition simulates the effects of striatal cell loss in Huntington's disease.

Figure 1: Architecture of the basal ganglia

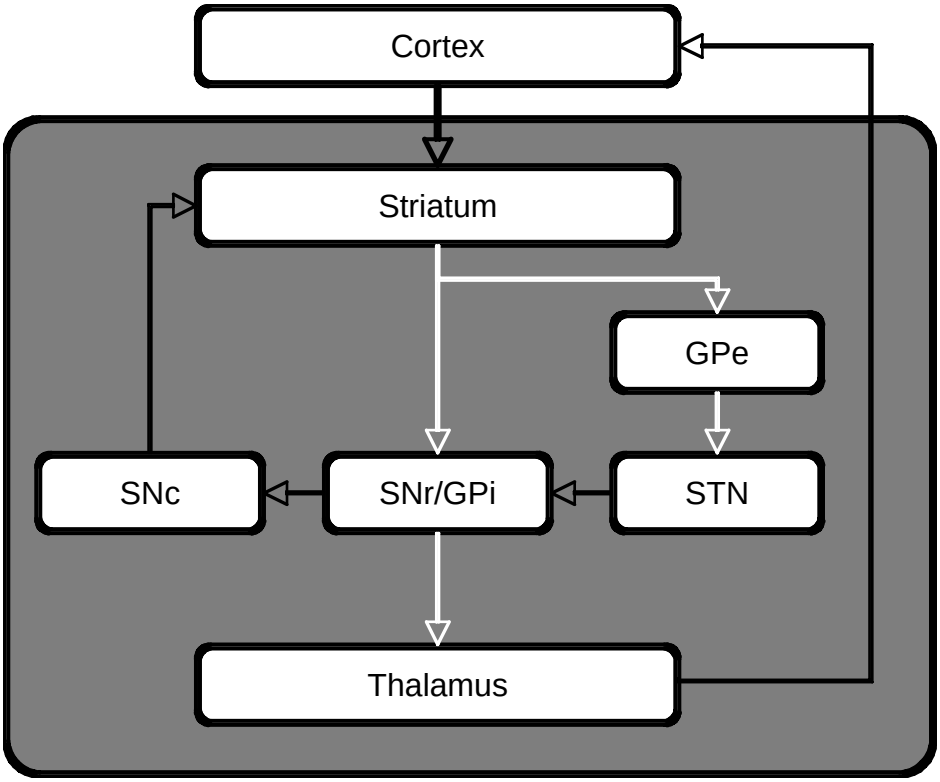

Figure 2: Architecture of the model

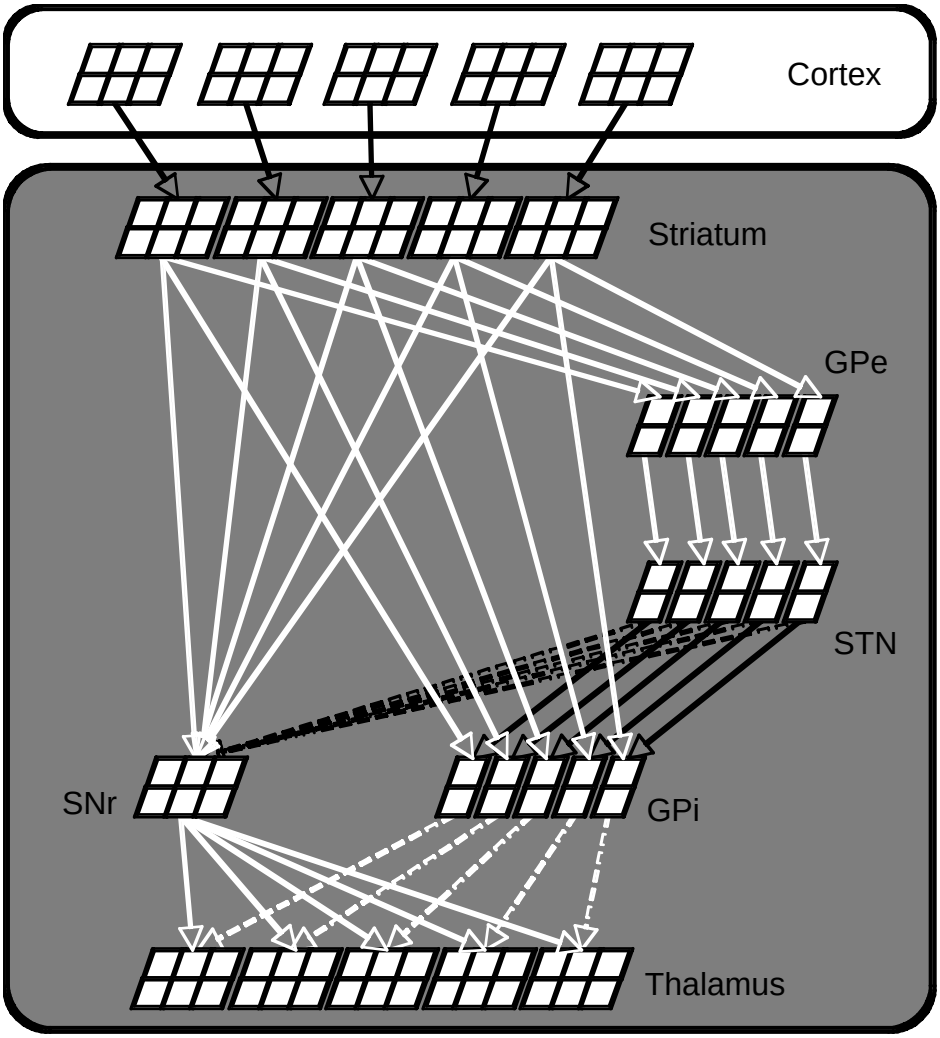

Figure 3: Internal connectivity of the Striatum

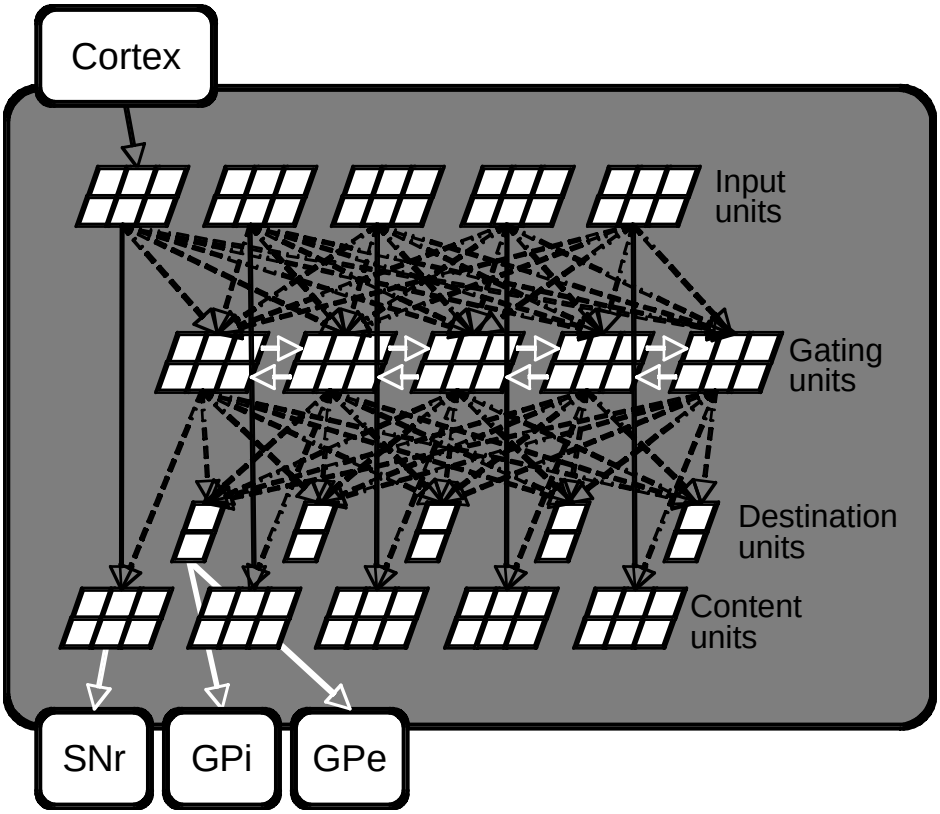

Figure 4: Representation of a routing operation in the striatum

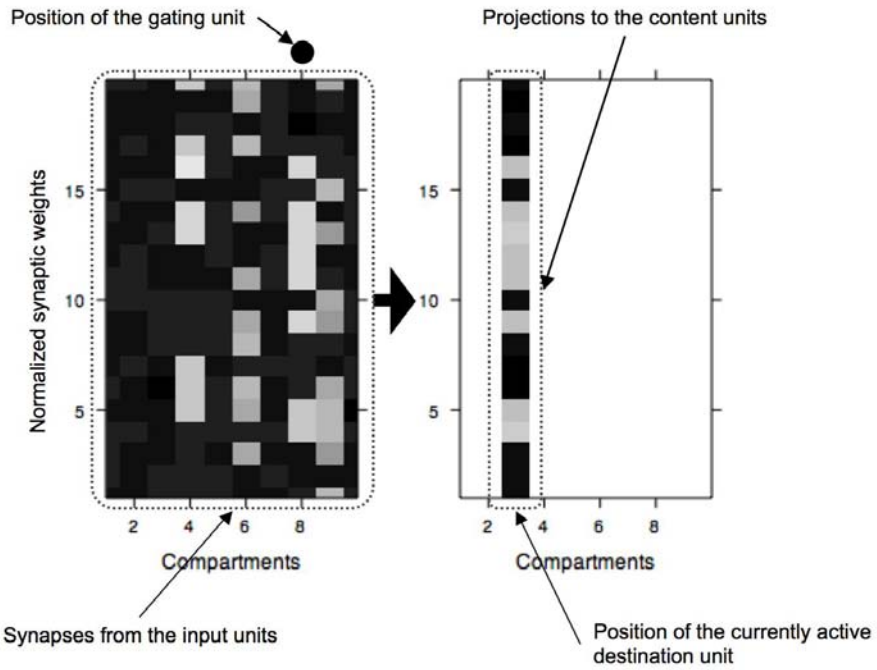

Figure 5: Transmission of information along the direct pathway

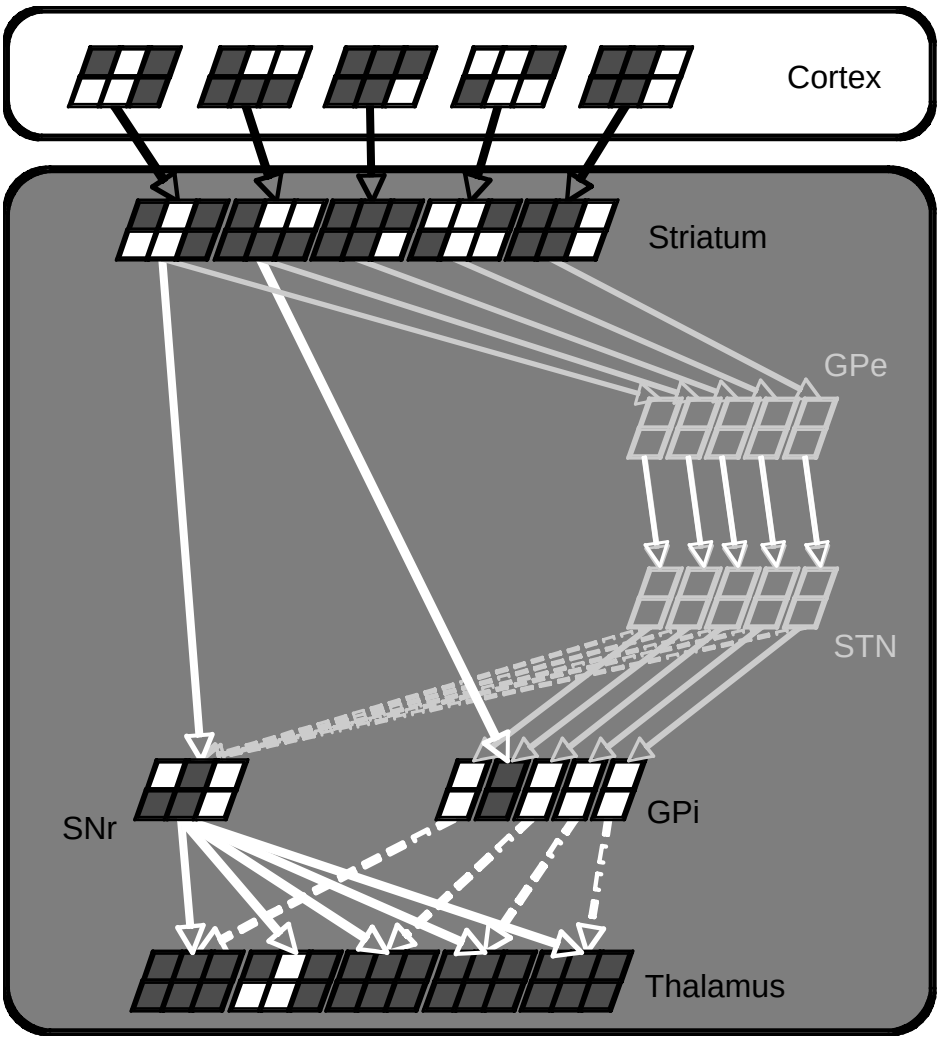

Figure 6. The dopaminergic pathway.

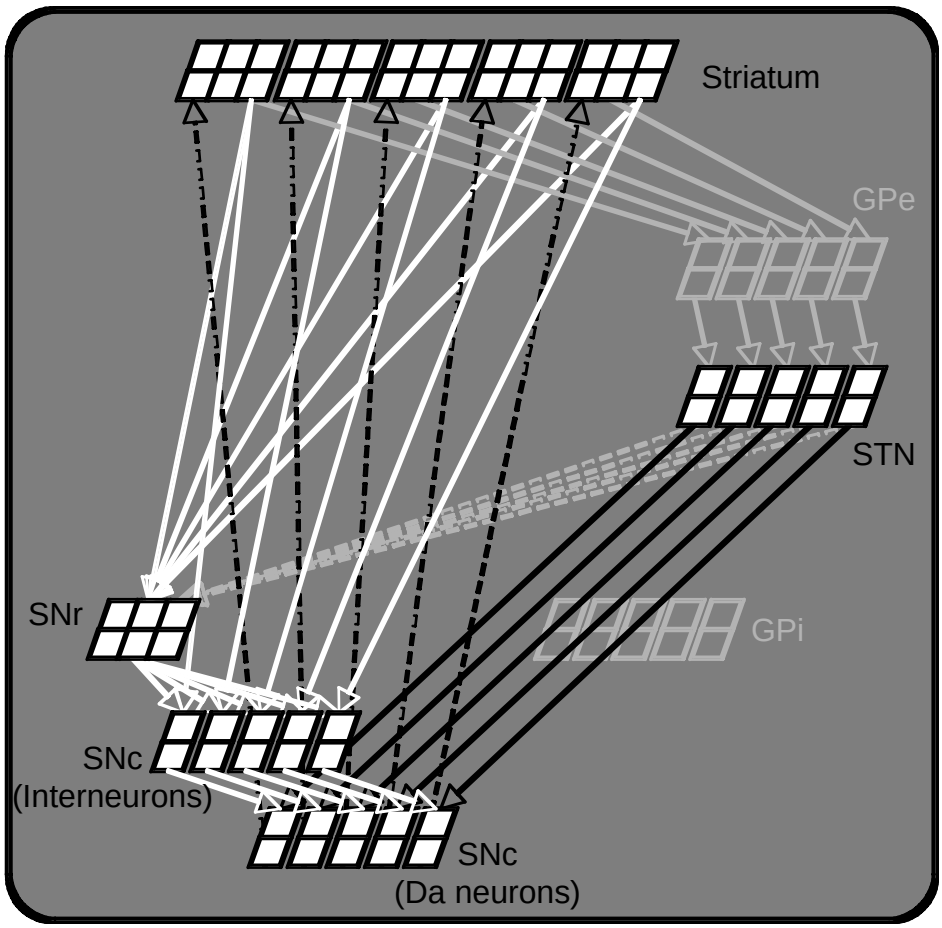

Figure 7: Example task

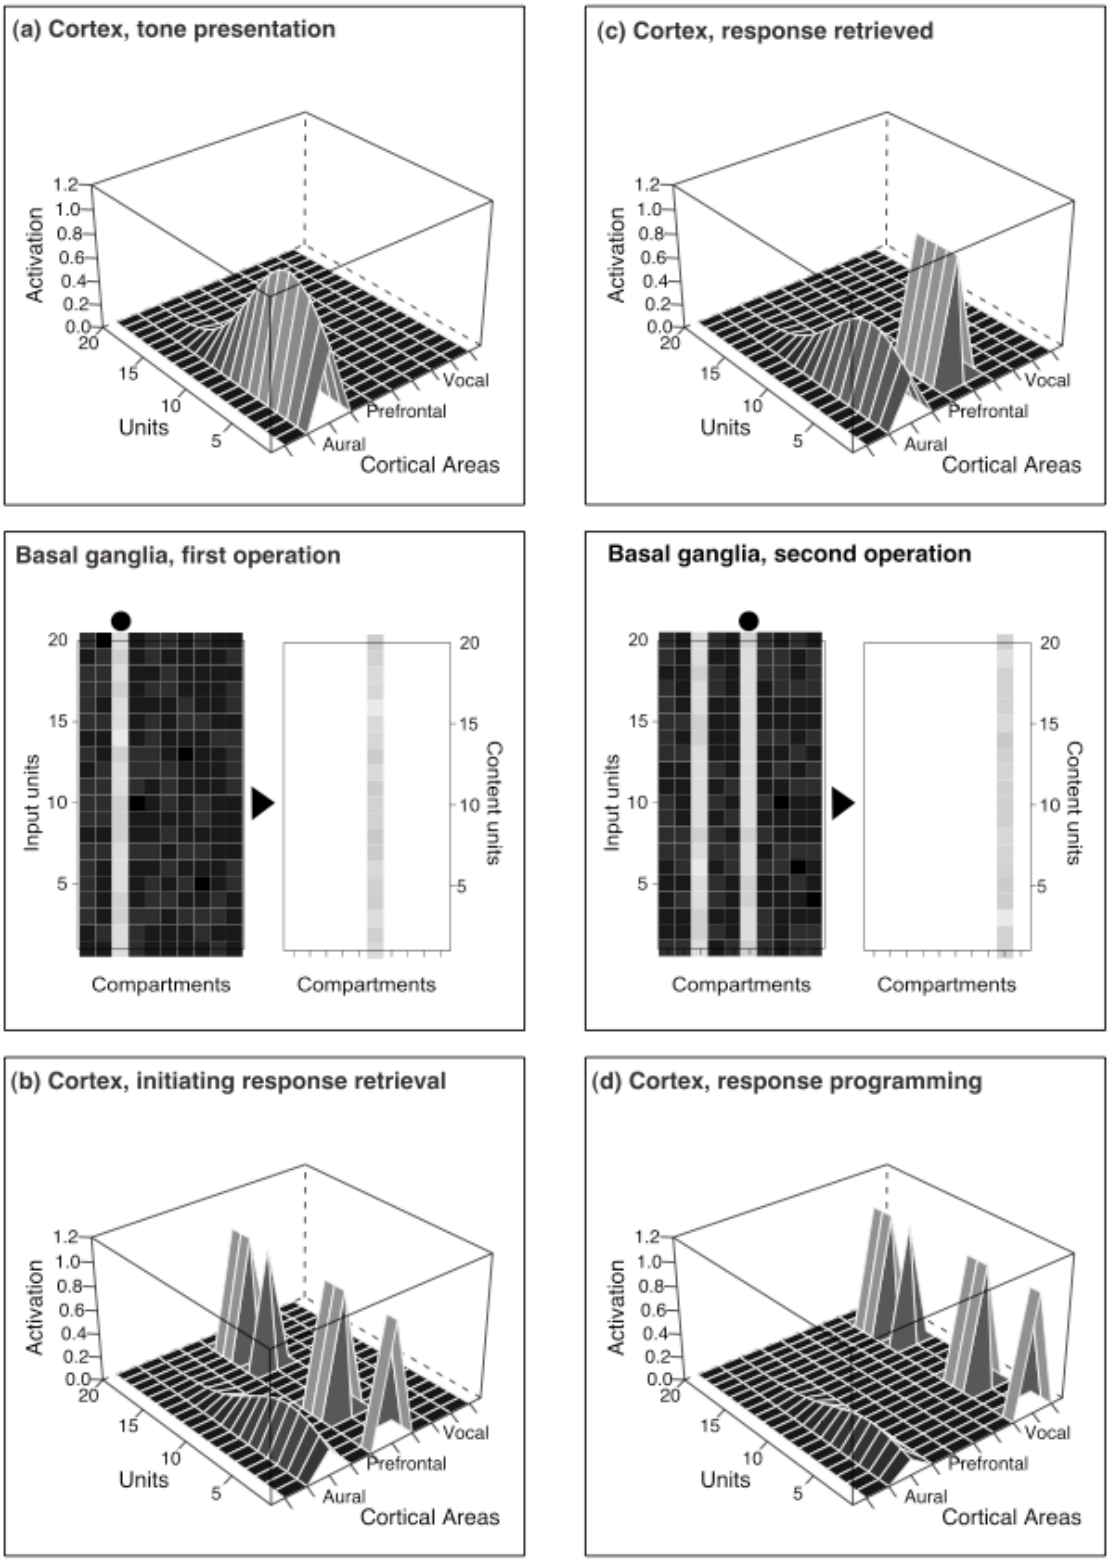

Figure 8: Model performance

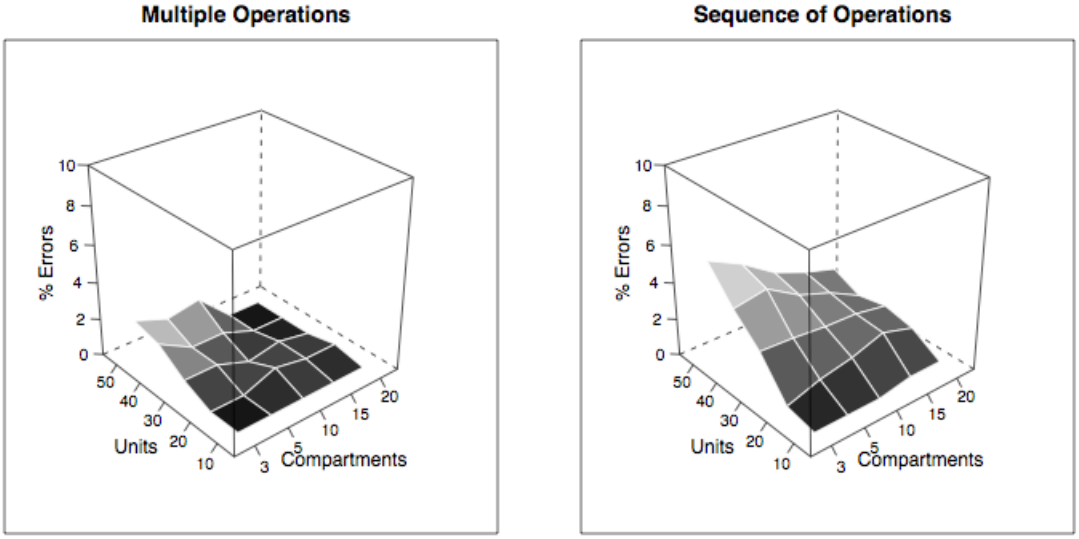

Figure 9: Dopamine-mediate learning of a new operation

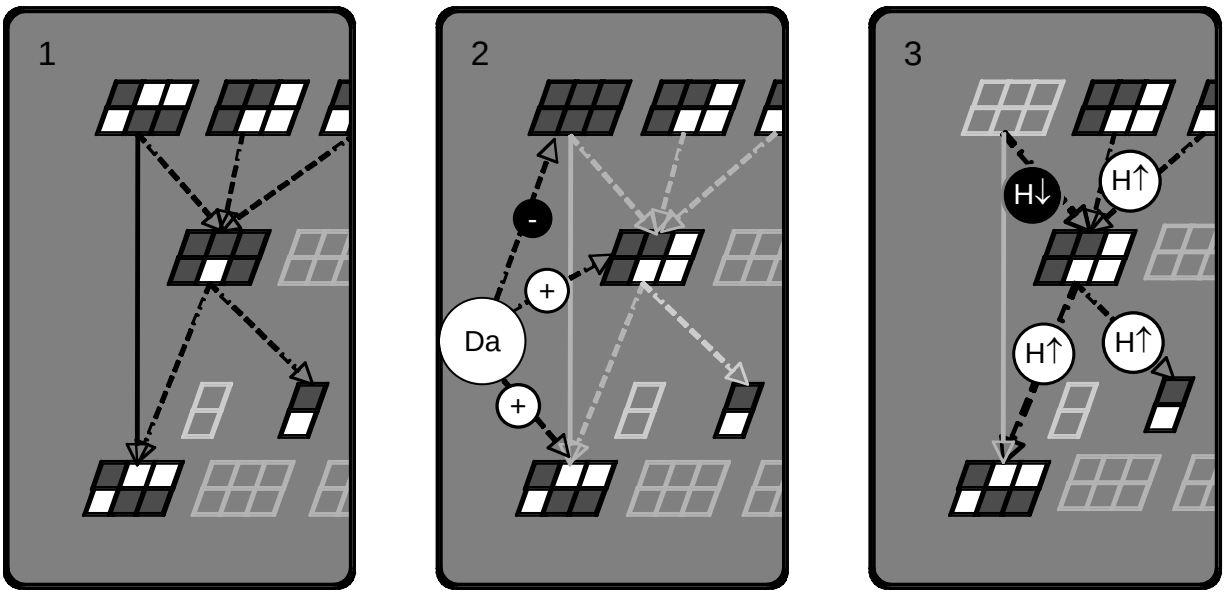

Figure 10: Task (after learning)

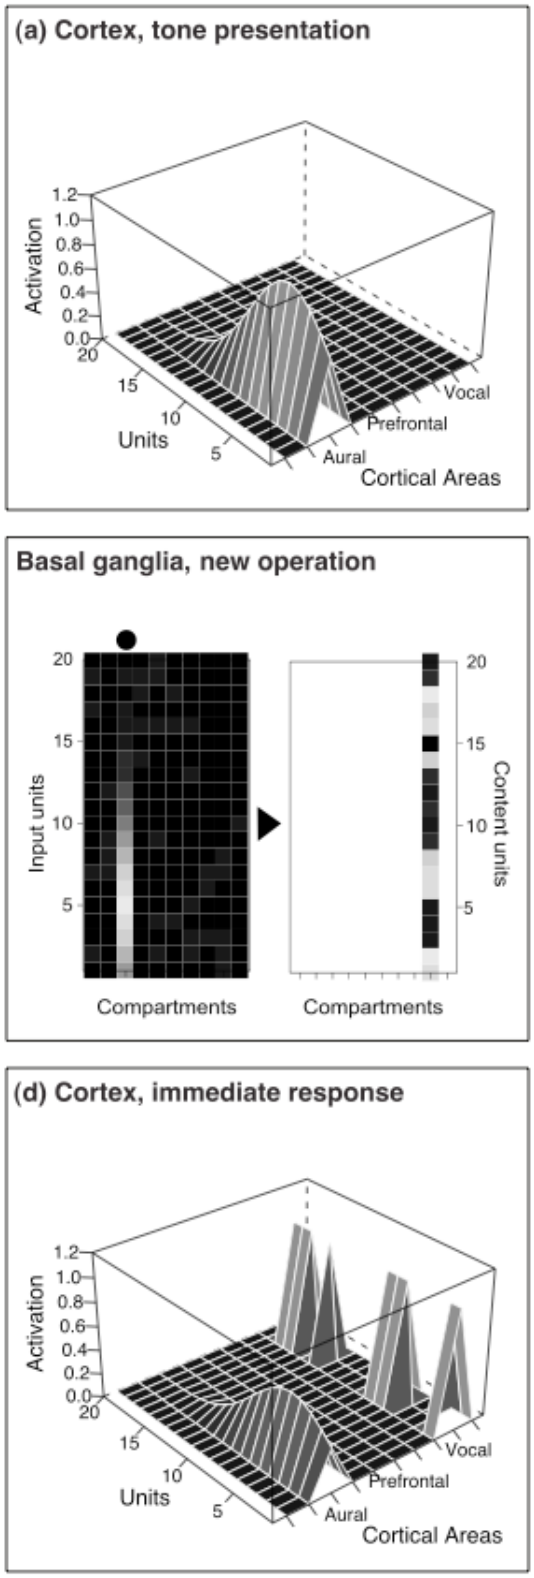

Figure 11: Effects of PD and HD