ELSEVIER

Progress in Neurobiology 71 (2003) 439–473

progress in Neurobiology

www.elsevier.com/locate/pneurobio

Information processing, dimensionality reduction and reinforcement learning in the basal ganglia

Izhar Bar-Gad[a,b,*], Genela Morris[a,b], Hagai Bergman[a,b,c]

[a] *Center for Neural Computation, The Hebrew Univeristy, Jerusalem, Israel*
[b] *Department of Physiology, Hadassah Medical School, The Hebrew University, P.O. Box 12272, Jerusalem, Israel*
[c] *Eric Roland Center for Neurodegenerative Diseases, The Hebrew University, Jerusalem, Israel*

Received 19 May 2003; accepted 1 December 2003

Abstract

Modeling of the basal ganglia has played a major role in our understanding of this elusive group of nuclei. Models of the basal ganglia have undergone evolutionary and revolutionary changes over the last 20 years, as new research in the fields of anatomy, physiology and biochemistry of these nuclei has yielded new information. Early models dealt with a single pathway through the nuclei and focused on the nature of the processing performed within it, convergence of information versus parallel processing of information. Later, the Albin–DeLong "box-and-arrow" model characterized the inter-nuclei interaction as multiple pathways while maintaining a simplistic scalar representation of the nuclei themselves. This model made a breakthrough by providing key insights into the behavior of these nuclei in hypo- and hyper-kinetic movement disorders. The next generation of models elaborated the intra-nuclei interactions and focused on the role of the basal ganglia in action selection and sequence generation which form the most current consensus regarding basal ganglia function in both normal and pathological conditions. However, new findings challenge these models and point to a different neural network approach to information processing in the basal ganglia. Here, we take an in-depth look at the reinforcement driven dimensionality reduction (RDDR) model which postulates that the basal ganglia compress cortical information according to a reinforcement signal using optimal extraction methods. The model provides new insights and experimental predictions on the computational capacity of the basal ganglia and their role in health and disease.

Contents

Abbreviations: GABA, gamma aminobutyric acid; GPe, globus pallidus external segment; GP, globus pallidus; GPi, globus pallidus internal segment; ICA, independent component analysis; MPTP, 1-methyl-4-phenyl-1,2,3,6-tetrahydropyridine; MSN, medium spiny neuron; PCA, principal component analysis; PD, Parkinson's disease; RDDR, reinforcement driven dimensionality reduction; SNc, substantia nigra pars compacta; SNr, substantia nigra pars reticulata; STN, sub-thalamic nucleus; TAN, tonically active neuron; VTA, ventral tegmental area

* Corresponding author. Tel.: +972-2-675-7388; fax: +972-2-643-9736.
E-mail address: izharb@alice.nc.huji.ac.il (I. Bar-Gad).

0301-0082/$ – see front matter © 2003 Elsevier Ltd. All rights reserved.
doi:10.1016/j.pneurobio.2003.12.001

1. Introduction

The basal ganglia are involved in normal processing of motor, associative and limbic information (Mink, 1996) and play a major role in some of the most common movement disorders such as Parkinson's disease (PD) (Denny-Brown, 1962; Boraud et al., 2002; Wichmann and DeLong, 2003). In the last 20 years extensive progress has been made in our understanding of these nuclei. The accumulation of experimental data has paved the way for new models of basal

ganglia functions. These models have, in turn, guided additional anatomical, physiological, biochemical and clinical research.

In this review, we shall first describe the core features of the experimental data underlying basal ganglia models. We will then describe the evolutions and revolutions of these models. Next, the review sets the groundwork for a new model of the basal ganglia: the reinforcement driven dimensionality reduction (RDDR) model. This model incorporates two general concepts from the field of machine learning: dimensionality reduction and reinforcement learning. The description of the basic RDDR model is presented in light of the known experimental facts and theoretical frameworks. Lastly, we turn to enhancements of this basic model that increase its scope in both the experimental and theoretical domains.

Overall, the basal ganglia are comprised of complicated structures with extensive and complex connectivity. In order to "see the forest and not just the trees" any model must ignore huge amounts of important information about the studied areas. The decision is never easy, is always very subjective, and clearly might be considered wrong to other researchers working on the basal ganglia (Parent and Cicchetti, 1998). It is thus worth stressing at the outset that future generations of basal ganglia models will doubtless use more and maybe even different features of current and future knowledge. This review hopes to avoid polemics by primarily providing information on experimental data and the theoretical background that inspired our understanding and hypotheses that serve as the basis for most current models of the basal ganglia. We hope that as previous models of the basal ganglia have done, the RDDR model will shed additional light on the mysterious and obscure functions of the basal ganglia in health and disease. These insights should lead to experimental predictions, which, when proved or disproved, can then form the basis for future (and better) models of the basal ganglia.

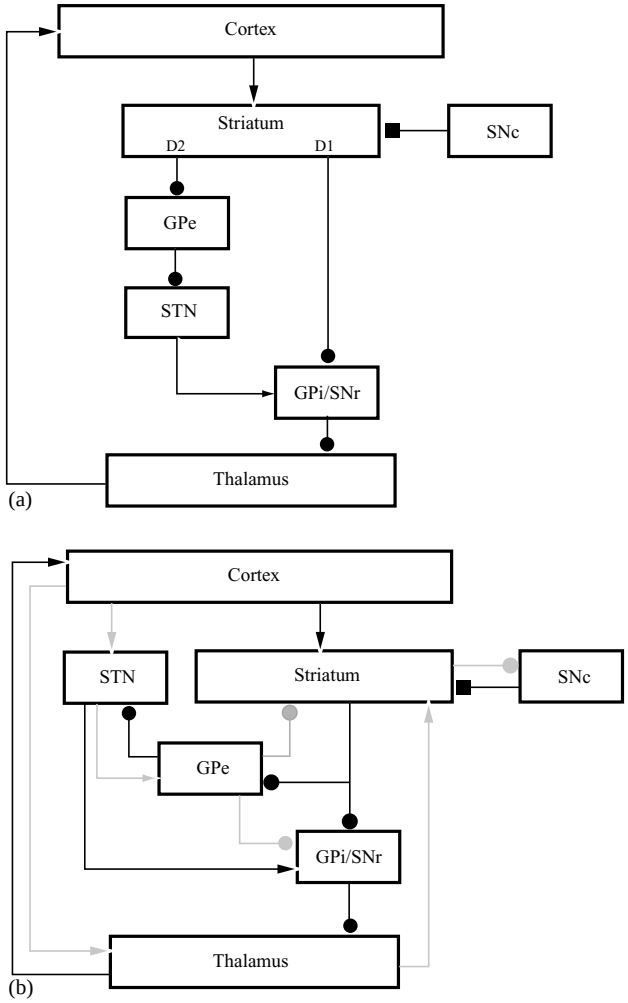

Fig. 1. The cortico-basal ganglia network. The box and arrow network of the different pathways of the basal ganglia. (a) The early Albin–DeLong network. (b) The up-to-date network. The early network is in black and later additions are in gray. Glutamatergic synapses are denoted by arrows, GABAergic synapses by circles and dopaminergic synapses by squares.

2. The biological basis of basal ganglia models

The basal ganglia participate in complex behaviors that require coordination between cognition, motivation and movements. This role is tightly linked to the anatomical position of the basal ganglia as a central part of a neuronal loop connecting most cortical areas with the frontal cortex. The basal ganglia are comprised of many nuclei with complex interactions among the neurons within each nucleus and among the different nuclei. This section describes the major issues of basal ganglia circuitry relevant to past and present models of the basal ganglia. More general and comprehensive reviews regarding basal ganglia anatomy, physiology and biochemistry can be found elsewhere (Parent and Hazrati, 1995; Gerfen and Wilson, 1996; Bolam et al., 2000). The biological background of this review is based primarily on the organization of the primate and human basal ganglia. However, data and insights are also presented from rodent work.

2.1. The basal ganglia nuclei

The striatum, the globus pallidus, the substantia nigra and the subthalamic nucleus (STN) are generally considered to be the main components of the basal ganglia. These components are further divided into sub-nuclei. The striatum is comprised of the caudate, putamen and ventral striatum (nucleus accumbens). The pallidal complex is comprised of the external (GPe) and internal (GPi) segments and the ventral pallidum. Finally, the substantia nigra is divided into the pars compacta (SNc), and the pars reticulata (SNr) (Fig. 1).

2.1.1. Striatum

The striatum is the primary input structure of the basal ganglia. It is divided into the dorsal and ventral striatum.

The dorsal striatum is further divided to form the caudate nucleus and the putamen. The ventral striatum is the ventral extension of the striatum that includes the nucleus accumbens, the medial and ventral portions of the caudate and putamen, and the striatal cells of the olfactory tubercle and anterior perforated substance (Gerfen and Wilson, 1996). It receives massive and topographic glutamatergic (excitatory) input from most parts of the cerebral cortex (Parent and Hazrati, 1995) and from the multiple nuclei in the thalamus (McFarland and Haber, 2000). The third major striatal input is from the midbrain (the SNc and VTA) dopaminergic cells (Haber et al., 2000).

Projection neurons make up the vast majority of the striatal neurons and are referred to as the medium spiny neurons (MSNs) of the striatum (Gerfen and Wilson, 1996). These medium sized (12–20 μm) cells have 25–30 dendritic branches and are covered by a large number of spines (Difiglia et al., 1976; Wilson and Groves, 1980). The branches radiate in all directions from the cell body to fill a spherical volume of 0.3–0.5 mm^3 (Gerfen and Wilson, 1996). Intracellular studies reveal that the membrane potential of these cells has up and down states (Wilson and Kawaguchi, 1996; Stern et al., 1997; Plenz, 2003). Evidence for the occurrence of such two-state membrane potential transitions come from numerical simulations and electrophysiological recordings in awake monkeys (Kitano et al., 2002). The transitions between the up and down states of the MSNs are controlled to a large extent by extrinsic connections (Wilson, 1993; Wickens and Wilson, 1998). This leads them to shift between the down state characterized by a zero firing rate to the short bursts of high firing rates (20–40 spikes/s) that characterize the up state (Crutcher and DeLong, 1984b; Hikosaka et al., 1989; Lee and Assad, 2003). The MSNs use gamma aminobutyric acid (GABA) (Bolam et al., 1985), typically considered to be an inhibitory neurotransmitter (Tremblay and Filion, 1989), as the primary transmitter and project to both segments of the globus pallidus and to the SNr (Parent and Hazrati, 1995).

The striatum contains several types of interneurons (Kawaguchi et al., 1995; Wilson, 1998; Haber and Gdowski, 2003), however we will discuss only two of them. The aspiny type II neurons are the largest of the striatal cells and represent about 1–2% of the total striatal population. These are the cholinergic interneurons, which are identifiable by their characteristic, spontaneous firing activity (3–10 Hz) (Kimura et al., 1984; Hikosaka et al., 1989; Wilson et al., 1990; Aosaki et al., 1995) and are therefore referred to as tonically active neurons (TANs). The TANs' spontaneous firing patterns are by and large a function of the intrinsic membrane properties of these cells (Bennett and Wilson, 1999; Bennett et al., 2000). These properties are modified by synaptic inputs that are influential in the temporal regulation of spike sequences (Aosaki et al., 1994; Raz et al., 1996). In addition, the TANs require only a relatively small number of extrinsic synaptic inputs to directly influence their spiking pattern (Bennett and Wilson, 1999). Extracellular recordings have revealed that TANs encode information related to reinforcement or incentive behavior (see Section 2.6). Another type of striatal interneuron is the parvalbumin-positive GABAergic fast spiking interneuron (Bolam et al., 2000) which represent about 3–5% of the striatal neuronal population. These cells have gap junctions, indicating that they are coupled in a continuous network (Koos and Tepper, 1999; Bolam et al., 2000). Thus, although relatively few in number; their inhibition on medium spiny projection neurons is likely to be particularly effective. Unfortunately, the GABAeregic interneurons cannot be identified by standard extra-cellular recording methods and data regarding their modulation during behavior are still lacking.

Histochemical and tracing studies have demonstrated discontinuities in transmitter-related molecules and afferent terminal distribution patterns in the striatum (Kunzle, 1975; Goldman and Nauta, 1977; Yeterian and VanHoesen, 1978; Graybiel and Ragsdale, 1980). Acetylcholinesterase (AChE)-poor striatal regions (termed striosomes) are surrounded by a densely stained "matrix" (Graybiel and Ragsdale, 1978). The significance of this compartmental organization in the striatum is still debated (Gerfen, 1989). In any case, within the striatum, the shape and extent of dendrites of most types of striatal neurons is restricted by the size and shape of the striatal compartment in which they are found (Gerfen, 1985; Malach and Graybiel, 1986; Penny et al., 1988). In contrast, the dendrites of the cholinergic interneurons readily cross striosome/matrix boundaries. Thus, these interneurons may help bridge between striatal compartments (Graybiel et al., 1994).

2.1.2. STN

The second major input structure of the basal ganglia is the STN which receives glutamatergic projections from the frontal cortex and cortical somato-motor areas (Monakow et al., 1978; Nambu et al., 1996). The STN is much smaller (in volume and number of cells) than the striatum with a ratio of 1:60 (Yelnik, 2002) in the primate to 1:200 (Oorschot, 1996) in the rodent. However, it plays a major role in basal ganglia activity in both normal and pathological conditions (Carpenter et al., 1950; Kitai and Kita, 1986; Bergman et al., 1990; Nambu et al., 2002b). As with most basal ganglia structures, the STN is largely composed of projection neurons and contains only a few small interneurons (Rafols and Fox, 1976; Yelnik and Percheron, 1979; Chang et al., 1983). The projection cells are tonically active (20–30 spikes/s) and fire short bursts during movement (Matsumura et al., 1992; Wichmann et al., 1994). STN glutamatergic (excitatory) projection neurons project to both segments of the globus pallidus and to the SNr (Parent and Hazrati, 1993; Smith et al., 1998a).

2.1.3. GPe

Classically, the GPe was considered to be a relay station within the basal ganglia, receiving input from the striatum and projecting to the STN (Albin et al., 1989; Alexander and

Crutcher, 1990). The accumulation of new data demonstrates that in addition to the classical connections there are STN inputs to the GPe and back projections to the parvalbumin-positive GABAergic interneurons of the striatum (Kita et al., 1999; Bolam et al., 2000). Thus, the GPe is reciprocally connected to both input structures of the basal ganglia, i.e. to the STN (Carpenter and Strominger, 1967) and the striatum. In addition, the GPe projects directly to the SNr (Sato et al., 2000) and its projections create dense GABAergic synaptic contacts with the soma of GPi neurons (Hazrati et al., 1990; Shink and Smith, 1995; Smith et al., 1998a). Therefore, in this review, we will emphasize the role of the GPe (or the globus pallidus (GP)—the rodent equivalent of the GPe) as part of the basal ganglia's intrinsic circuitry.

The vast majority of neurons observed in the GPe are large projection neurons with thick, smooth (spine free) and long (up to 1000 μm) dendrites (Fox et al., 1974; Difiglia et al., 1982; Park et al., 1982; Francois et al., 1984). Intracellular physiological studies also identify two to three types of neurons in the globus pallidus, with a single predominant type (Nakanishi et al., 1990; Nambu and Llinas, 1997; Cooper and Stanford, 2000). Extracellular recordings of pallidal spiking activity have revealed that the majority (>85%) of GPe neurons have a high frequency (50–70 spikes/s) discharge with pauses (DeLong, 1972). As with the striatum, the major transmitter of the pallidal projection neurons has been shown to be GABA (Oertel and Mugnaini, 1984).

2.1.4. GPi

The GPi, which receives GABAergic inputs from the striatum and GPe and glutamatergic input from the STN, is referred to as the output station of the basal ganglia. The GABAergic outputs from the GPi project to thalamic relay nuclei (the ventro-lateral thalamic nucleus; the ventral anterior thalamic nucleus, the medio-dorsal (MD) nucleus, and the lateral habenular nucleus) and to the intralaminar (non-specific) thalamic nuclei, the centromedian and parafascicular nucleus (Kuo and Carpenter, 1973; Kim et al., 1976; Parent et al., 2001). Most of the anatomical, physiological and biochemical features of the neurons of the GPi resemble those of the GPe. (Fox et al., 1974; Difiglia et al., 1982; Park et al., 1982; Francois et al., 1984; Nakanishi et al., 1990). However, unlike the GPe neurons, almost all of the GPi neurons have a high frequency (60–80 spikes/s) discharge with no pauses (DeLong, 1972). The entopeduncular nucleus (EP) is the rodent equivalent of the GPi.

2.1.5. SNr

The SNr is often considered a caudomedial extension of the GPi based on morphology, connectivity, biochemistry, and physiology (Schwyn and Fox, 1974; Kitai, 1981; Yelnik et al., 1987; Francois et al., 1987). The SNr, like the globus pallidus, consists primarily of large projection neurons with long, thick dendrites, which are almost completely ensheathed with synaptic contacts from the striatum (Yelnik et al., 1987; Francois et al., 1987). However, other evidence points to some differences between the SNr and the GPi. The SNr does not seem to be derived from the same developmental anlage as either the external or internal pallidal segments (Merchand et al., 1986). In addition, the dopaminergic cells of the ventral tier of the SNc have dendrites that penetrate deep into the SNr. Somato-dentritic release of dopamine (Jaffe et al., 1998), as well as GABA effects on SNc neurons therefore suggest that the SNr relations with the SNc dopaminergic neurons are more direct and complicated than those of the GPi (Celada et al., 1999). Finally there are the SNr-superior colliculus projections, which have a major effect on eye and attention-shift movements (Hikosaka and Wurtz, 1983; Redgrave et al., 1992; Handel and Glimcher, 1999; Jiang et al., 2003). Nevertheless, due to their general similarity of function, we will treat the GPi/SNr as a single unified output stage of the basal ganglia.

2.1.6. SNc

The SNc and other midbrain dopaminergic neurons receive input from many structures including the striatum (Heimer et al., 1982; Haber et al., 2000), the STN and the limbic system (Haber and Gdowski, 2003). The midbrain dopamine system is divided into the dorsal tier and ventral tier cells (Haber et al., 2000). The dorsal tier cells include the VTA (A10 in rodents), the retrorubral group (A8), and the dorsal cells of the SNc (A9). The ventral tier cells include the densocellular group and the cell columns that extend deep into the pars reticulata. Nevertheless, the responses of DA neurons recorded over different territories of SNc and VTA are not significantly different (Schultz, 1998) and therefore in this review we will treat the midbrain DA neurons as a homogenous neural structure, reporting the same information to its efferent structures.

Electrophysiological studies reveal two firing modes: either single spikes or in bursts (Grace and Bunney, 1984ab). In the behaving primate, the dopaminergic cells tend to fire spontaneously at a low rate (4–10 spikes/s) with elevations and suppressions of their firing rate which are related to mismatch between behavior and prediction (Schultz, 1998) (Section 2.6). Dopaminergic projections from the SNc and the VTA terminate onto the spines as well as the dendritic shafts of the MSNs of the striatum (Freund et al., 1984; Smith et al., 1994; Hanley and Bolam, 1997). Electron microscopic studies indicate that cortical terminals are often found on the heads of spines and dopaminergic terminals are found nearby, but on the dendritic shafts (Dube et al., 1988; Smith et al., 1994). The relationship between dopamine and thalamic terminals in the striatum is less clear and is not discussed in this review.

2.2. The basal ganglia pathways

Many areas in the nervous system are characterized by their reciprocal connections. A good rule of thumb is that if there are connections from area A to B, there will also be anatomical connections from B to A. This is the situation

in the cortex (Abeles, 1991; Scannell et al., 1995; Sporns et al., 2002) and in the thalamo-cortical loops (Sherman and Guillery, 2001; McFarland and Haber, 2002). However, the classical picture of connectivity in the basal ganglia stands in sharp contrast to this reciprocal connectivity rule. The main pathways of the basal ganglia form a feed-forward network with unidirectional connections from the cortex to the input nuclei of the basal ganglia (striatum and STN), from the striatum and the STN to the GPe, and from these three structures to the output nuclei of the basal ganglia (GPi and SNr). Finally, the GPi and the SNr provide unidirectional projections to their target nuclei in the thalamus. The feed-forward structure changes upon its termination at the reciprocally connected network of the thalamus and the cortex. The pathways finally create a partially closed loop by projecting to the thalamus and frontal cortex that send input back to the input nuclei of the basal ganglia.

2.2.1. Multiple feed-forward pathways

There are two main types of MSNs in the striatum: one that co-contains substance P and GABA and projects primarily to the GPi and SNr; and one that co-contains enkephalin and GABA and projects primarily to the GPe (Gerfen et al., 1990). MSNs that contain high mRNA expression levels for substance P also co-contain mRNA for the dopamine D_1 receptors, and cells that contain mRNA for enkephalin also co-contain high mRNA expression levels for the D_2 receptor subtype (Gerfen et al., 1990; Aubert et al., 2000). The separation of receptor subtypes within different pathways, along with their different pharmacological actions and co-localized neuropeptides, has been particularly important in defining the functional framework of the "direct" and "indirect" pathways (Albin et al., 1989; Alexander and Crutcher, 1990; Gerfen et al., 1990) (Section 3.2). In this framework, the "direct" pathway leads from the striatum directly to the GPi, whereas in the "indirect" pathway the information flows from the striatum to the GPe, onwards to the STN and from there to the GPi (Fig. 1a). However, in primates, the separation of the peptide co-transmitters into the two pallidal segments is not complete. Rather, the GPe contains substance P-positive immunoreactivity along the medial boundary, and enkephalin immunoreactivity is found in the medial portion of the GPi (Haber and Elde, 1981; Reiner et al., 1999). In addition, recent studies have reported subpopulations of MSNs that co-express different subtypes of dopamine receptors (Surmeier et al., 1996; Nicola et al., 2000), suggesting that the direct/indirect pathways are not as segregated as once thought. We will therefore postpone discussion of the direct/indirect concept to the advanced topics section of this review (Section 7.6). The functional role of substance P and the enkephalin is still in doubt (Steiner and Gerfen, 1998); we will therefore not examine their neuromodulating effects in this review. In addition to the cortico-striatal-pallidal pathways, studies in rodents (Kita, 1992; Ryan and Clark, 1992) and primates (Nambu et al., 2000, 2002b) emphasize the functional significance of the "hyper-direct" fast excitatory pathway from the cortex to the STN whose effect on the output nuclei is typically faster than the effect of the striatal "direct" and "indirect" pathways.

2.2.2. Feedback pathways

Exceptions are the rule in any biological system. The main exception to the feed-forward structure is the dual role of the GPe. Classically, the GPe was considered to be a relay station in the indirect pathway, receiving information from the striatum and sending it to the STN (Fig. 1a). However, new research has shown that the GPe also receives input from the STN and in addition sends output to the striatum (Fig. 1b).The parvalbumin-positive, GABA interneurons of the striatum which receive powerful input from the cortex (Lapper et al., 1992; Kita, 1993), are also the target of a significant back-projection from the GPe (Bolam et al., 2000). In addition, the role of the reciprocal connections between the GPe and STN has been emphasized in recent physiological (Plenz and Kitai, 1999) and computational (Bevan et al., 2002; Terman et al., 2002) studies. In conclusion, the GPe maintain a high degree of bidirectional information flow with the input nuclei of the basal ganglia (see also Section 2.1.3). However, since the physiological role of these projections is still unknown we will not discuss them further in this review.

A different reciprocal loop connects the striatum and the midbrain dopaminergic neurons. In addition to the dopaminergic innervation of the striatum (Freund et al., 1984; Smith et al., 1994; Hanley and Bolam, 1997) (see also Section 2.1.6), the striatum projects massively back to the substantia nigra (Szabo, 1980; Parent et al., 1983). Some studies indicate that these projections arise mainly from the striosomes of the striatum (Gerfen, 1992). This reciprocal pathway serves as a major component for the generation of the reinforcement signal (Houk et al., 1995).

2.2.3. Completing the loop

The thalamus has long been thought to convey spinal and subcortical information to the cortex, forming the last link in the cortico-basal ganglia circuitry (Albin et al., 1989; DeLong, 1990). Indeed, models of basal ganglia function view the primary role of the thalamus as a relay of information processed in the basal ganglia from the GPi/SNr level to the cortex. The thalamic nuclear groups that are associated primarily with this function are the ventral anterior (VA) and ventral lateral (VL) nuclei and the medio-dorsal thalamic nucleus (Haber and Gdowski, 2003). The organization of the thalamus (reticular nucleus and specific and non-specific nuclei) and its reciprocal connections with the cortex (Sherman and Guillery, 2001) are beyond the scope of this review. However, the complex reciprocal connections between the thalamus and the cortex require some explanation. Injections of bidirectional tracers into thalamic and frontal cortical sites also show that in comparison to "feed-forward" thalamo-cortical projections, cortico-thalamic "back" projections to the ventral anterior, ventral lateral and medio-dorsal nuclei are more widespread.

These findings demonstrate both the reciprocal and non-reciprocal components of the thalamo-cortico-thalamic relay (McFarland and Haber, 2002).

Anatomical and physiological studies indicate that there are projections from the intralaminar nuclei of the thalamus (centro-medial/parafasicular) to the striatum and the STN (Orieux et al., 2000; Haber and Gdowski, 2003). Recent studies in primates show a substantial projection from the VA and VL nuclei to the striatum, suggesting that the basal ganglia loop is closed both at the level of the frontal cortex and the striatum (McFarland and Haber, 2001).

2.3. Decrease in the neuronal population along the cortico-striatal-pallidal pathway

The main axis of the basal ganglia in terms of number of neurons is the pathway from the cortex to the striatum and from there to the GPi and SNr. A striking feature of this pathway is the funnel structure observed when looking at the number of neurons along each step of the pathway. The numbers themselves vary greatly depending on the animal and the research methodology. However, the same ratio seems to persist across different studies. In terms of number of neurons, the STN is very small relative to the striatum (see also Section 2.1.2) and therefore is not discussed here.

The first convergence takes place from the cortex to the striatum. The number of rat cortico-striatal neurons has been determined to be 17×10^6 converging onto 1.7×10^6 striatal neurons (Zheng and Wilson, 2002) which yields a convergence ratio of 10. A second, and larger reduction in the number of neurons is between the striatum and the GPi and SNr where the factor is in the range of 10^2–10^3. Research on the rat (Oorschot, 1996) shows a factor of 95 [$N_{\text{Striatum}}/(N_{\text{GPi}} + N_{\text{SNr}}) = 2790 \times 10^3/(3.2 \times 10^3 + 26.3 \times 10^3)$]. The numbers in other species need to be deduced from the many studies that have examined the different nuclei separately. It is important to note that pooling such data might lead to distortions. In the macaque, the data were integrated by Percheron et al. (1987) and indicated an even higher ratio of 571 [$12,000 \times 10^3/(9 \times 10^3 + 12 \times 10^3)$]. Research in humans, conducted by several groups (Thorner et al., 1975; Schroder et al., 1975; Percheron et al., 1987) reports a convergence ratio of 347 [$110,000 \times 10^3/(160 \times 10^3 + 157 \times 10^3)$].

An open question is the number of thalamic neurons receiving projections from the GPi/SNr, and the number of (frontal) cortical neurons connected to these thalamic neurons. Although indirect evidence supports the notion of an increase in the number of neurons in these stages (Arecchi et al., 1997; Francois et al., 2002), and therefore a bottle-neck structure of the cortex-basal ganglia-cortex loop, we are not aware of any quantitative study of these aspects.

2.4. Inhibitory collaterals—anatomy and physiology

Axon collaterals of the medium spiny neurons of the striatum terminate within the striatum, onto both striatal interneurons and other medium spiny cells (Bolam et al., 2000). Terminals of the MSNs form symmetrical synaptic contacts and contain GABA (Bolam et al., 1985; Kita, 1993; Smith et al., 1998b). Such GABAergic synapses are considered inhibitory and their connections to the nearby projection neurons would theoretically indicate lateral inhibition (Wickens, 1993; Wickens and Oorschot, 2000). However, early physiological studies failed to find the signature of such lateral interactions in the striatum and concluded that they were weak or non-existent (Jaeger et al., 1994). More recent studies, using averaging methods (to enhance the signal to noise ratio) have revealed inhibitory synaptic potentials between striatal projection neurons both in slice and organotypic tissue culture preparation (Tunstall et al., 2002; Czubayko and Plenz, 2002). The inhibitory potentials were found in less than a third of the tested pairs, and in most cases were unidirectional and not reciprocal.

The projection neurons of the globus pallidus give rise to thick, sparsely spined, poorly branching dendrites (Fox et al., 1974; Iwahori and Mizuno, 1981; Difiglia et al., 1982; Yelnik et al., 1984). These dendrites are very long, sometimes creating dendritic radii in excess of 2 mm in their principal plane with the very distal portions often branching elaborately to form complex dendritic endings (Difiglia et al., 1982; Francois et al., 1984; Difiglia and Rafols, 1988). The physiological significance of these specialized dendrites is not understood, although it has been suggested that they may contribute to local synapses between neighboring pallidal neurons (Yelnik et al., 1984; Francois et al., 1984; Shink and Smith, 1995). Another feature of many large pallidal and SNr neurons is the fine, beaded, generally unbranched axonal-like processes originating at irregular intervals from pallidal dendrites. These extend for moderate distances (mean, 80 μm) and, in some instances, appear to contact the soma or dendrites of adjacent neurons (Francois et al., 1984; Yelnik et al., 1997). Evidence of local arborization has been found in multiple animals (Kita and Kitai, 1994; Yelnik et al., 1997; Nambu and Llinas, 1997; Bevan et al., 1998; Parent et al., 2000; Sato et al., 2000). Sato et al. (2000) have shown that one-third of primate pallidal neurons have several collaterals that arborize within their entire somatodendritic domain and even beyond it. As with the complex endings of pallidal dendrites (Francois et al., 1984), these thin processes appear to be more frequent within the GPe than in the GPi. In a slice preparation of the GP, spontaneous IPSP can be observed (Cooper and Stanford, 2000). The authors suggest that these IPSP arise from active GP cells in the slice. However, GP to GP synaptic connectivity has only been observed in 1 out of 40 recorded pairs, and their spontaneous and glutamate evoked spiking activity was not correlated (Stanford, 2003).

Cross-correlation studies of spiking activity of pairs of neurons recorded simultaneously in behaving primates can reveal whether these neurons receive common inputs and whether they directly affect each other's activity (Perkel et al., 1967; Eggermont, 1990). Common inputs lead to

a higher level of coincident action potentials, resulting in a double-sided peak in the cross-correlation function, whereas direct synaptic connections lead to one-sided peaks or troughs in the cross-correlation functions. Flat cross-correlograms, on the other hand, indicate an absence of direct and indirect interactions between the neurons. Therefore, cross-correlation studies are essential to determine the different possible modes of functions in the basal ganglia (Bergman et al., 1998). Multiple electrode recordings (Baker et al., 1999) and spike sorting methods (Lewicki, 1998) that enable discrimination between two neighboring neurons whose electrical activity were recorded by a single electrode are the main methods used to achieve this goal. Recordings of striatal pairs in behaving monkeys (Jaeger et al., 1995) and recordings of striatal pairs in anesthetized rats (Stern et al., 1998) have shown no significant correlations on a short timescale. Less than 10% of the pallidal cross-correlograms (calculated at different behavioral epochs) displayed significant correlations in the normal monkey (Nini et al., 1995; Raz et al., 2000; Heimer et al., 2002). A study of the firing patterns of neighboring neurons shows that they have the same characteristics as remote neurons within the globus pallidus, namely uncorrelated firing (Bar-Gad et al., 2003).

In conclusion, there is extensive anatomical evidence for lateral inhibitory connections in the striatum. A similar picture, although probably less extensive, can be found in the globus pallidus. Nevertheless, intracellular and extra-cellular physiological studies have revealed only weak and sparse functional lateral connectivity within these structures. This lack of functional strong lateral connectivity is a major impetus for our model of the basal ganglia and will be discussed in depth below.

2.5. Basal ganglia connectivity

The identified pathways between the nuclei provide only part of the information needed to understand the complex connectivity within the basal ganglia. Understanding the pattern of connectivity between nuclei requires a fuller description involving a number of aspects. The first is the amount of convergence that occurs within the nuclei (Bergman et al., 1998); namely, does all the information merge together or are there parallel pathways segregating parts of the information? The second aspect is the sparseness of the organization, i.e. how much of the total input each neuron sees. These two aspects are actually intertwined and form a multitude of other characteristics of connectivity such as the topographic nature of the organization, differentiation between regions, and levels of common input. The computational aspects of the relationship between sparseness and organization will be discussed later (Section 7.5).

2.5.1. Parallel versus funnel-like connectivity

There are two extreme views regarding the amount of anatomical or information sharing in the basal ganglia (Percheron and Filion, 1991). The first view holds that neurons in the output stage of the basal ganglia receive a highly convergent input leading to information funneling (Bolam et al., 1993; Percheron et al., 1994), whereas the second view describes the basal ganglia as segregated parallel circuits with minimal interactions between the parallel pathways (Alexander et al., 1986; Alexander and Crutcher, 1990; Parent and Hazrati, 1993; Middleton and Strick, 2000).

Although the extent of the cortico-striatal projections is large, the main terminal region is localized in the most adjacent position to its cortical input (Kemp and Powell, 1970; Selemon and Goldman Rakic, 1985). In general, cortico-striatal projections terminate in a functional topographic manner (Takada et al., 1998; Tokuno et al., 1999) and can be roughly divided into the limbic, cognitive and motor domains (Parent and Hazrati, 1995). Thus, cortico-striatal projections terminate along the ventral to dorsal axis of the striatum. The dorso-lateral striatum (putamen) receives cortical input from sensory-motor areas, the central striatum (caudate) receives input from associative cortical areas, and the ventro-medial striatum receives input from limbic areas. For example, there is no substantial overlap between cortical inputs from the rostral (high-order cognitive control of movement) and caudal (motor execution) cingulate motor in the striatum (Takada et al., 2001). Anterograde labeling of the striato-pallidal projections in monkeys also supports the general concept of segregated basal ganglia pathways. The striato-pallidal fibers are organized in multiple elongated bands parallel to the external borders of the pallidum (Parent, 1990). Striato-pallidal axons originating from neurons located at the caudate nucleus (associative striatal territory) or at the post-commissural putamen (sensory-motor territory) remain well segregated from one another at the GPe and GPi levels (Parent, 1990). Furthermore, anterograde double labeling studies revealed that striatal axons emerging from adjacent groups of striatal neurons are segregated at the pallidal levels (Hazrati and Parent, 1992). The striato-nigral projections are organized in multiple, distributed but segregated plexuses like the striato-pallidal projections (Hazrati and Parent, 1992). Anatomical studies of trans-neuronal virus transport also support the notion of a parallel macro-organization of the basal ganglia. Thus, different frontal cortical areas project most densely to distinct striatal regions (input channels) and each of these cortical areas is influenced by projection from distinct GPi regions (output channels) (Hoover and Strick, 1993; Middleton and Strick, 2002).

Early electrophysiological studies also support (and were even a major impetus for) the concept of segregated parallel organization of the basal ganglia (Alexander et al., 1986). Single unit and micro-stimulation studies of striatum (Crutcher and DeLong, 1984a; Alexander and DeLong, 1985) and pallidum (DeLong et al., 1985) revealed somato-topic organization of these structures. Moreover, the lack of pallidal units with clear relations to both arm and leg movements (DeLong et al., 1985; Alexander and

Crutcher, 1990) suggest that the segregation may even apply to specific body parts. Electrophysiological stimulation studies of MI and SMA (Nambu et al., 2002a) reveal that most striatal neurons respond exclusively to MI stimulation or to SMA stimulation and that MI and SMA responding neurons are distributed predominantly in different portions of the putamen. Similar studies have also shown segregation of pallidal neurons receiving inputs from different cortical fields (Yoshida et al., 1993).

On the other hand, the supporters of the "funneling" model emphasize the fact that in all gross tracer and virus studies the borders between the different territories are not sharp, and extensive overlap and interdigitation exist in the striatum (Kunzle, 1975; Yeterian and VanHoesen, 1978; Selemon and Goldman Rakic, 1985; Ramanathan et al., 2002), GPi (Flaherty and Graybiel, 1993; Bolam et al., 1993) and STN (Kolomiets et al., 2001). Moreover, the wide dendritic arborization of pallidal neurons oriented at right angles to the incoming striatal axons (Percheron et al., 1984), the collateral system in the striatum and even in the pallidum (Parent et al., 2000), the divergent STN projection to the GP, and finally the substantial reduction in the number of neurons from the striatum to the globus pallidus suggest extensive convergence or funneling at the pallidal level.

Careful analysis of the results of the physiological studies has shown that even in these studies, the picture is not entirely black or white. The micro-stimulations studies cited above (Nambu et al., 2002a) report that about 20% of the recorded striatal neurons responded concurrently to stimulation in both the MI and SMA. Similarly, a considerable number of pallidal neurons were inhibited by the stimulation of more than one cortical area (the prefrontal, premotor, supplementary motor and arcuate premotor areas, and the motor cortex) (Yoshida et al., 1993). Funneling of inputs from remote striatal neurons to a focal area in the pallidum was also demonstrated by orthodromic stimulation of the striatum (Kimura et al., 1996). Our recent studies of monkeys involved in a complex behavioral task also revealed a significant fraction of neurons with responses to multiple behavioral events, including events which classically belong to the limbic, cognitive and motor domains.

2.5.2. *Sparse connectivity*

The connectivity between the nuclei along the main axis of the basal ganglia is extremely sparse. Out of an order of 10^8 cortico-striatal neurons in the primate (based on a ratio of 10 between cortico-striatal and striatal neurons (Zheng and Wilson, 2002) and an order of 10^7 striatal neurons in the primate (Percheron et al., 1987)), a single striatal neuron receives on the order of 10^4 cortical synapses (Wilson et al., 1983; Ingham et al., 1993). Thus, assuming that each synapse originates from a different cortical neuron, a single striatal neuron receives input from at most approximately 0.01% of the cortical neurons innervating the striatum. The same holds true in the GPi which is innervated by around 10^7 striatal neurons (Percheron et al., 1994), whereas each pallidal cell has only about 10^4–10^5 synapses (Percheron et al., 1994) originating from the striatum. Therefore, a typical pallidal neuron receives at most input from 0.1 to 1% of the striatal projection neurons. The connectivity between the neurons within the same nucleus via the GABAergic collaterals is also limited in its scope. Striatal neurons have about 10^3 collateral synapses (Kawaguchi et al., 1995) leading to a maximal probability of 0.01% for a synapse with another striatal projection neuron. Pallidal neurons form a smaller number of collateral synapses, probably on the order of 10^2 synapses (Parent et al., 2000) and a connection ratio of 0.1% within the nucleus. It is crucial to remember that even these low connectivity factors between and within the nuclei represent the highest probability estimate. Any occurrence of multiple synapses between neuron pairs would decrease the probability of a connection with other neurons even further (Kincaid et al., 1998; Zheng and Wilson, 2002).

However, the connectivity between the neurons of different nuclei and within the nuclei is not uniform. Specific pairs of neurons tend to have increased probability for a synapse depending on topographic and other organizational properties. Thus, these neurons tend to be "nearby", either functionally or spatially (Wilson, 2000). Cortico-striatal axons tend to create groups of synapses within a limited area (Kincaid et al., 1998), and the distance between the synapses follows an exponential distribution (Kincaid et al., 1998). In addition, the (at least partial) topographic organization of the striatum is derived from the fact that neurons participating in the same function tend to converge to the same area (Yeterian and VanHoesen, 1978; Selemon and Goldman Rakic, 1985; Flaherty and Graybiel, 1991). Thus, the probability of striatal neurons within the same topographic compartment to receive inputs from the same neuron or neighboring neurons within the cortex is significantly elevated. The pallidum also displays a topographic organization (DeLong et al., 1985), and the arborization of the striatal axons tends to follow a specific pattern (Hazrati and Parent, 1992). Striatal MSNs have a spherical shape to their dendritic and local axonic trees (Wilson, 2000), limiting the range of their connectivity to the diameter of 400 μm (Wilson and Groves, 1980) and thus leading to lateral connectivity limited to neighboring neurons. In addition, the dendrites of the MSNs do not cross the border between patch/matrix compartments (Malach and Graybiel, 1986). In the pallidum, the collaterals are sparse, and the size of the dendritic tree is about 1000 μm. However, as in the striatum, the collaterals are spatially organized, innervating spatially nearby neurons (Kita and Kitai, 1994; Yelnik et al., 1997; Nambu and Llinas, 1997; Bevan et al., 1998; Parent et al., 2000; Sato et al., 2000). An interesting observation by J. Yelnik (private communication) is that "as pallidal neurons are very few in number in comparison to spiny striatal neurons, a rather poor local arborization could be sufficient to interconnect an equal proportion of neighboring neurons in both structures.".

Overall, the striatum seems to be divided into domains (Wilson, 2000) with a higher level of connectivity and zero

or very low connectivity between domains, possibly mediated by interneurons (Kawaguchi et al., 1995). The pallidum also has a topographic division, however it is less clear-cut (Hamada et al., 1990). Thus, nearby striatal and pallidal neurons tend to share inputs representing the same functions in their input structures.

2.6. Reinforcement signals in the basal ganglia

Adaptive behavior requires the evaluation of environmental stimuli with respect to their behavioral significance (Robbins and Everitt, 1996). Neuromodulators such as dopamine and acetylcholine neurons acting in the striatum emit signals that can perform such an evaluation. These substances affect the plasticity of cortico-striatal transmission (Calabresi et al., 2000; Reynolds et al., 2001) and are therefore good candidates to serve as teachers of the basal ganglia system.

2.6.1. The dopaminergic signal

Midbrain dopamine neurons play a key role in the acquisition of new behaviors. While early studies emphasized the role of dopamine in reward-related (hedonic) behavior (Fibiger and Phillips, 1988; Wise and Rompre, 1989) the recent seminal studies by Schultz and colleagues reveal that the primary function of dopamine neurons is to report the mismatch between the animal's prediction and reality (Hollerman and Schultz, 1998; Schultz, 1998; Waelti et al., 2001; Fiorillo et al., 2003) (but see (Redgrave et al., 1999; Horvitz, 2000)). Striatal responses to dopamine are heterogeneous and are probably determined by multiple variables including complex interactions between several receptors, the state (up or down) of the striatal neurons and the firing pattern (tonic versus burst firing) of the dopaminergic neurons (Di Chiara et al., 1994; Arbuthnott and Wickens, 1996; Cepeda and Levine, 1998; Wickens and Oorschot, 2000; Calabresi et al., 2000). In this review, we will emphasize the recent findings showing that dopamine controls LTP/LTD in the cortico-striatal pathway (Wickens et al., 1996; Centonze et al., 1999; Kerr and Wickens, 2001).

2.6.2. The cholinergic signal

Extra-cellular physiological recordings in the striatum of awake, behaving monkeys have revealed pauses in the spontaneous TAN firing, often flanked by excitation in conjunction with rewarded (but also aversive) events (Kimura et al., 1984; Crutcher and DeLong, 1984b; Liles, 1985; Apicella et al., 1991; Ravel et al., 2003). TANs are therefore thought to be involved in the detection of stimuli that have inherent motivational significance (Kimura et al., 1984; Aosaki et al., 1994; Apicella et al., 1997; Ravel et al., 2001). Our recent study of dopaminergic and TAN neuronal populations in monkeys performing a probabilistic instrumental conditioning task reveals that while the response of the dopamine neurons reflects a mismatch between expectation and outcome, the TAN's response was indifferent to the predictability or reward value of the events. The spiking activity of the TANs is highly temporally correlated (Raz et al., 1996; Kimura et al., 2003). The different responses and correlation patterns of dopaminergic neurons of the SNc and TANs suggest that the two systems do not mirror one another. The data are in line with the working hypothesis that the cholinergic signal informs the basal ganglia when to learn, the dopaminergic signal tells them how to learn, and the nature of the previous cortico-striatal activity defines what will be learned (Section 4). Since acetylecholine reduces the responsiveness of striatal projection neurons to cortical inputs by fixing their up/down state (Akins et al., 1990), the pause in firing the TANs, exhibited in response to behaviorally significant events, may serve to ensure that the striatal neurons "get the message" in time, enabling the dopaminergic reinforcement signal to strengthen the synaptic efficacy between co-activated cortical and striatal neurons.

2.6.3. Other reinforcement signals

The other main brainstem input to the striatum is serotonergic and arises from the dorsal and medial raphe nuclei (Lavoie and Parent, 1990). Unlike the distribution of dopaminergic terminals in the striatum, serotonergic terminals are not evenly distributed throughout the striatum. Rather, they are more abundant in the ventral striatum (Lavoie and Parent, 1990). There are very few physiological studies (Sawyer et al., 1985) on the role of these neuromodulating systems in the functions of the basal ganglia, and we will devote no further attention to them in the present review.

2.7. Pathophysiology of the basal ganglia

A relationship between the substantia nigra and Parkinson's disease was first postulated at the turn of the 19th century (Brissaud, 1895). This was later confirmed by the demonstration of progressive cell loss in the substantia nigra of Parkinson's patients (Hassler, 1939). The connection between the neurotransmitter dopamine and the substantia nigra was made by the discovery that dopamine was depleted in PD patients (Ehringer and Hornykiewicz, 1960). The development of the primate 1-methyl-4-phenyl-1,2,3,6-tetrahydropyridine (MPTP) model of Parkinsonism in the early 1980s (Burns et al., 1983; Wilms et al., 1999) enabled great strides forward in our understanding of the pathophysiology of Parkinson's disease. The MPTP model serves as the key model for basal ganglia disorders in this review. In addition, insights are incorporated from other animal models such as the 6-OHDA rodent model and from human studies.

Several primate studies have shown that MPTP has a differential effect on firing rates in the different basal ganglia nuclei. The average firing rate of GPi and STN neurons increases from ∼80 to ∼100 spikes/s (Miller and DeLong, 1987; Filion and Tremblay, 1991; Boraud et al., 1998) and from ∼20 to ∼25 spikes/s (Miller and DeLong, 1987; Bergman et al., 1994), respectively. The rate of GPe neurons

decreases from ∼70 to ∼50 spikes/s (Miller and DeLong, 1987; Filion and Tremblay, 1991; Boraud et al., 1998) (but see (Boraud et al., 2002) for unaltered firing rates). The finding of an abnormal tonic firing rate in the basal ganglia of MPTP monkeys was confirmed in human PD patients (Merello et al., 1999; Levy et al., 2001) and by the alleviation of all Parkinsonian symptoms by ablation or inactivation of the GPi or STN (Bergman et al., 1990; Benabid et al., 1994).

In addition, firing patterns in the basal ganglia circuit of primates are dramatically altered following MPTP treatment. There is an increase in the percentage of neurons that discharge in irregular or oscillatory bursts (Miller and DeLong, 1987; Filion and Tremblay, 1991; Bergman et al., 1994; Boraud et al., 1998). Studies in human PD patients have reported cells whose discharge is modulated in the tremor frequency range in GPi and in the STN (Levy et al., 2002). However, the correlation between tremor and neuronal oscillations is at best intermittent and is dynamic in nature (Hurtado et al., 1999; Lemstra et al., 1999; Raz et al., 2000).

In normal behaving animals most of the cross-correlograms between pallidal neurons are flat, indicating that the neuronal pair is functionally independent (Nini et al., 1995; Raz et al., 2000). However, physiological studies in the globus pallidus of MPTP-treated monkeys demonstrate that the cross-correlograms become peaked and oscillatory (Raz et al., 2001; Heimer et al., 2002). Similar findings of increased synchronization within primate brains following MPTP treatment have been found in primary motor cortex (Goldberg et al., 2002), among striatal TANs and between the TANs and pallidal neurons (Raz et al., 1996, 2001).

In conclusion, physiological studies of the basal ganglia of MPTP treated primates and PD human patients have revealed significant changes in the firing rate, pattern and synchronization of these structures. The amelioration of most Parkinsonian symptoms by inactivation or deep-brain stimulation of several basal ganglia targets, further suggests that these or other changes in basal ganglia activity (following striatal dopamine depletion) are the main cause of the these symptoms.

3. Evolution and revolutions of basal ganglia models

Over the last 20 years, information processing models of the basal ganglia (Beiser et al., 1997; Gillies and Arbuthnott, 2000) in health and disease have undergone revolutionary and evolutionary changes. Early models of these nuclei dealt with the main pathway; namely, the pathway from the cortex sequentially through the striatum to the globus pallidus. These single pathway models were primarily concerned with the functional organization of this loop and the issue of parallel processing (Alexander et al., 1986) versus convergence (Percheron et al., 1984) of information along the pathway. The incorporation of other known pathways within the basal ganglia into the models shifted the focus to the delicate interplay between these multiple pathways. Descriptions of the pathways and their interaction using the box and arrow models led to multiple insights into the function of the basal ganglia in health and disease (Albin et al., 1989; DeLong, 1990). The number and complexity of elements in the pathways (“boxes”) and their interactions (“arrows”) grew to encompass new anatomical and physiological data (Wichmann and DeLong, 1996). This evolution of the box and arrow models led in turn to the opening of the boxes, describing the internal nuclei structure, and modeling the interplay within the nuclei in both the spatial and temporal domains (Hikosaka et al., 1993). This generation of models focused on action selection properties of the basal ganglia using more elaborate neural networks emphasizing the computational properties of the nuclei (Wickens, 1997). Recent work expands the basal ganglia models to incorporate information flow within the larger framework of the information loop from the cortex through the basal ganglia and back to the cortex, creating a partially closed loop functioning in sequential processing (Berns and Sejnowski, 1998).

3.1. Single pathway models

Early models of the basal ganglia emphasized the central pathway of information: the feed-forward flow from the cortex, through the striatum to the GPi. Despite the general consensus regarding the direction of the information flow and the associated nuclei, there was (and still is) controversy concerning the interaction of information from multiple sources flowing through the basal ganglia. The two extreme views are that there is either convergence of all the disparate information from many domains or a parallel flow of information through segregated loops.

Percheron et al. (1984) used anatomical data to demonstrate the convergence of multiple striatal sources onto the same pallidal neurons. Each pallidal neuron can have access to a very large number of striatal axons passing through its dendritic tree. This anatomical structure can lead to the convergence of sensorimotor, associative and limbic information and therefore can serve as a substrate for integration of dispersed cortical information. In addition, this convergence is highest in neurons of the GPi, which is the output part of the basal ganglia. The convergence is not only the result of the reduction in the number of neurons (“cardinal convergence”) but primarily the outcome of the shape of the emitting axonal and receiving dendritic arborization (“reception convergence”) (Percheron and Filion, 1991). Later electrophysiological data (Yoshida et al., 1993) supported the anatomical data, showing that although the inputs seem to segregate into two parallel loops, there is still high convergence from multiple sources in the cortex to the globus pallidus within each loop.

In contrast, in the parallel segregated loops model, information derived from different domains is believed to be segregated within separate loops and no interaction occurs

between the different parallel pathways. In this model, the different loops arise from different cortical areas, maintain a separation throughout their flow through the basal ganglia and thalamus and finally end up in different areas in the frontal cortex. The original division was into a "motor" loop and an "associative" loop (DeLong and Georgopoulos, 1981). The division was extended to five functionally distinct parallel loops that originate from different, functionally-related cortical areas, from which a convergent process occurs within each segregated loop separately (Alexander et al., 1986). The division into five loops was further granulated by a subdivision of the motor loop into three separate and seemingly segregated pathways connecting different pallidal areas with distinct cortical areas (Hoover and Strick, 1993). This type of subdivision led to the extension of the notion of parallel processing to processing within the different loops and not only to the relation between the different loops (Alexander and Crutcher, 1990). The division into sub-channels within the loops was suggested as the result of the somatotopic organization and the topographical projection in the different nuclei. The separation into multiple loops restricts information convergence and integration to be at most within the narrow sub-channels. Recent publications tend to describe a large number of loops of information flowing through the basal ganglia, e.g. 10 loops have been reported (Middleton and Strick, 2000) with distinct anatomical and physiological characteristics.

These two extreme views, a converging funnel that integrates all the cortical information, or a parallel pathway structure that keeps the information segregated throughout the cortico-basal ganglia-cortical loop, were merged in some of the later models to create more complex functional architectures. One such architecture is constructed of open interconnected loops (Joel and Weiner, 1994). This architecture, which relies on divergence from the striatum to the SNr and GPi, assumes that although the output from a single frontal cortical area is sent to a single input area in the basal ganglia, it receives the input from at least two basal ganglia output nuclei. This causes at least some of the loops through the basal ganglia to return to areas other than the ones that served as their input (i.e. an open loop). This theory deviates from the idea of closed loops which keep information segregated, and rather presents a complex interconnection and information flow between different cortical areas due to the multiple targets of the returning information. A different method of integrating parallel loops was proposed in the ascending spirals model (Haber et al., 2000). In this model, the integrating factor of the loops is not the frontal cortex but rather the striato-nigral system. The model is based on findings that the nigro-striatal loop is not a topographically reciprocal closed loop. Rather, nigral (dopaminergic) projections also extend to striatal areas that did not innervate them, thereby passing information between different striatal areas. The information flow within the loops is thus modulated by the activity of other loops.

3.2. Multiple pathway models

The accumulation of data regarding the complexity of basal ganglia connectivity shifted the focus of models to the interactions between the different nuclei and pathways. Mapping the anatomical components and their chemical interaction formed the basis for the box and arrow model of the basal ganglia (Albin et al., 1989; DeLong, 1990) (Fig. 1). In this model, the various nuclei are modeled as simplified units ("boxes"), which can increase or decrease their overall firing rate depending on the rate change in afferent nuclei and on the type of their connections ("arrows"). The connections are either positive using the neurotransmitter glutamate or negative using the neurotransmitter GABA. The cortico-basal ganglia-cortex loop is perceived as a feedback loop with two competing elements: a positive feedback mechanism mediated by the direct pathway (striatum-GPi) while the indirect pathway (striatum-GPe-STN-GPi) mediates negative feedback. Thus, the basal ganglia control the overall inhibition on the thalamus, which is relayed to the cortex. The relative activity of the two pathways is controlled by the dopaminergic signal of the SNc, which has opposing effects on the two pathways: increasing the activity of the direct pathway through D1 receptors while decreasing the activity of the indirect pathway through D2 receptors. This model has been especially successful in explaining the hypo-kinetic and hyper-kinetic movement disorders associated with the basal ganglia (DeLong, 1990). Reduced activity in the direct pathway and enhanced activity in the indirect pathway (due to dopamine depletion) leads to elevated activity in the output structures of the basal ganglia. This leads to the inhibition of the cortex through the thalamic relay, resulting in hypo-kinetic disorders (e.g. Parkinson's disease). On the other hand, enhanced activity in the direct pathway and reduced activity in the indirect pathway lead to the inhibition of the output structures, which in turn disinhibit (i.e. activate) the cortex, resulting in hyper-kinetic disorders (e.g. Huntington's disease and hemibalismus). The organization of the circuits themselves may incorporate either the parallel loops or the convergent structure and accordingly is concerned with the direct/indirect pathway balance within each loop or within the whole structure. Over the years multiple additions have been made to incorporate new anatomical data such as the hyper-direct (cortico-STN-GPi) pathway and the GPe to GPi pathway (Wichmann and DeLong, 1996; Chesselet and Delfs, 1996). The various patches made to the original model over the years have managed to explain some of the new anatomical and physiological data, but have also increased its complexity. Nevertheless, the box and arrow model still serves as the main model for understanding the information processing of the basal ganglia, despite the debate over its shortcomings and viability (Wichmann and DeLong, 1996; Chesselet and Delfs, 1996; Parent and Cicchetti, 1998). Although the box and arrow model has been a valuable tool in understanding the hypo-kinetic and hyper-kinetic pathologies associated with

the basal ganglia, it provides little insight into the computation performed by the healthy basal ganglia. Analysis of the cortico-basal ganglia-cortex loop through gross representations of the complete nuclei and observation of the firing rate alone is bound to overlook the complex firing patterns, the computation performed within each of the nuclei, and the interaction between the neurons within the same nuclei.

3.3. Action selection

The shortcomings of the box and arrow model motivated a switch from system level models to more detailed network level models aimed at understanding the normal function as well as pathological states of the basal ganglia. Currently, the main network level models of the basal ganglia are action selection models. Their common emphasis is on the role of the basal ganglia, within the cortico-basal ganglia-cortical loop, in choosing one or more actions out of a multitude of such actions presented to the basal ganglia by the cortex (Mink, 1996). The models vary in the nature of the actions selected, which range in their definition from low-level "simple" motor actions to high-level "complex" behavioral schemes. The models are also divided into two major categories depending on whether the selection mechanism is intra-nuclei or inter-nuclei selection. Intra-nuclei selection is achieved using lateral (or recurrent) inhibition within the nuclei whereas inter-nuclei selection uses feed-forward competition between the different pathways. A few models integrate the two selection mechanisms to achieve an enhanced selection process.

Models of intra-nuclei action selection utilize anatomical data showing lateral GABAergic connections between the medium spiny neurons of the striatum (Wickens, 1993, 1997; Beiser and Houk, 1998; Plenz, 2003). In these models, a winner-take-all mechanism, which forms the basis of many action selection methods, is implemented using a neural network that converges to a single winner. The dense network of inhibitory connections between MSNs is assumed to inhibit neighboring neurons, thereby maintaining the activity of only one single neuron. The selection process takes place within the striatum and the chosen action is then conveyed via the different pathways to the output layer of the basal ganglia. These models are also used as the actor part in the actor-critic model (Houk et al., 1995; Baldassarre, 2002) in which the matrix part of the striatum utilizes the reward signal from the SNc to select the "correct" action or higher level plan. Although most models have focused on the selection of a single action, some work has been done on "soft" selection in which multiple actions may be chosen in parallel. A model that implements such an approach is the winner-shares-all model (Fukai and Tanaka, 1997). This model, which is a variant of the winner-take-all mechanism, depends on the relationship between lateral and self-inhibition of the neuron to enable more than one neuron to become active. This mechanism enables the transition from a selection model to a more dynamic system that filters inputs based on a changing threshold dependent on the overall input.

Inter-nuclei selection models vary in the ways their selection mechanisms operate and in the suggested interplay between the different nuclei and pathways of the basal ganglia. One of these models describes the disinhibition of the thalamo-cortical networks by the basal ganglia in two basic ways: simultaneous or sequential activation of opposite pathways (Hikosaka et al., 1993). In simultaneous activation, the direct and indirect pathways co-activate the output nuclei simultaneously, which results in a sharper effect and spatial focusing of the output targets. In sequential activation (also called temporal scaling), the indirect pathway maintains a non-selective inhibition, which is followed by a selective disinhibition performed by the direct pathway. According to the focused selection model (Mink, 1996) the basal ganglia receives input from multiple motor pattern generators (MPGs) and enables normal motor function by releasing (or disinhibiting) a single pattern while inhibiting the other patterns. Using this focused selection, competing MPGs are prevented from working simultaneously and disrupting normal motor function. Unlike the earlier box and arrow model, this model does not describe the different basal ganglia nuclei as uniform entities, but rather explores the internal relationships between representations within the same nuclei. In the winner-lose-all model (Berns and Sejnowski, 1996), the indirect pathway enhances background activity of the GPi, leading to inhibition of all the actions whereas the inhibition of a single program using the direct pathway (hence the term winner-lose-all) leads to the disinhibition of a single cortical action. Competing actions are blocked by using temporal differences between the different pathways. The model assumes the existence of "units" which correspond to neuron pools. These "units" are part of segregated streams of information leading from the striatum through the globus pallidus to the thalamus. A model presented by Gurney et al. (2001) is also based on the notion of a focusing signal sent by the striatum, as compared to the diffused excitation of the STN. This achieves the off-center on-surround pattern of activation. However, in this model the GPe is added as a control loop that helps stabilize the selection pathway and enhances selectivity. This differs from the role of the GPe in the Berns and Sejnowski model in which it acts to create the temporal delay critical for the actual selection process. In addition, Gurney's model enables both "hard" and "soft" selection (single versus several selected actions). Finally, some models attempt to put the action selection networks of the basal ganglia within a larger framework. This includes the role of the thalamus and the thalamo-cortical connection as part of the action selection network (Humphries and Gurney, 2002).

3.4. Sequence generation

Another approach to modeling the basal ganglia associates the basal ganglia with the role of selecting or generating action sequences rather than single actions. Sequence

generation and detection models of the basal ganglia can be broken down into two main categories. The first assumes an internal basal ganglia mechanism for maintaining memories of actions or states that make up the sequence. In the second, the basal ganglia are "memory-less" components of the sequential action mechanism. In this category, the state is maintained by the cortex and reflects to the striatum via the cortico-striatal connection.

One model using an internal basal ganglia state builds on an earlier action selection model (Berns and Sejnowski, 1996) by adding the selection and encoding of sequences (Berns and Sejnowski, 1998). In this model, sequence learning and generation is achieved using short-term memory encoded by the reciprocal connectivity between the GPe and the STN. The loop contains traces of previous activity and effectively works as a short-term memory in the millisecond time scale. Longer-term memory is associated with either the connection of the output neurons back to the input neurons (a timescale of hundreds of milliseconds) or with memory storage within the prefrontal cortex (a timescale of seconds).

Another model for reproducing and discriminating between sequences (Dominey, 1995) breaks down the sequencing task into sub-networks. The prefrontal cortex encodes the history of the input (i.e. the internal state of the animal). The cortico-striatal connections act as an associative memory component involved in linking the history with the desired output. The firing of the striatal neurons represents the motor output. The associations are determined by the reinforcement signal received from the SNc, which modifies the cortico-striatal synapses. The model resembles the action selection models except for the fact that the striatum identifies sequences instead of single actions. Along the same lines, Beiser and Houk (1998) presented a model for encoding sensory information using the cortico-basal ganglia-cortical loop. In this model, the striatum classifies cortical inputs which contain the contexts of the action sequences.

Alternatively, in a model presented by Fukai (1999), the actual action sequences are stored in the frontal cortex using oscillations. In this model, the role of the basal ganglia is to decode the sequences and their components. One group of neurons in the striatum is responsible for choosing the initial part of the sequence. This is done as part of a temporal winner-take-all mechanism implemented by lateral inhibition. The other striatal group retains the currently executed action. The excitatory signal to the GPi by the STN is suggested to be an external signal to the basal ganglia, signaling movement transitions planned in the cortical buffer. The actual action selection is conceptualized as part of the thalamic relay with a possibility for additional selectivity in the striatum and globus pallidus.

4. Reinforcement learning

Inspired by animal behavior studies, classical machine learning theories divide the term "learning" to two separate types. In the first type, the learning element, the agent, is accompanied by an all-knowing element, the teacher, which informs it on the correct action to be taken in each situation. This form of learning is termed "supervised learning", and is typically comprised of a set of examples (tasks) and their respective answers (actions). When applying this to neural networks, the back-propagation architecture is often implemented. There, an example is presented and the network responds according to its previous knowledge. Then, the correct answer is provided and is compared to the network's output. This comparison is used for learning: the difference between each unit's output and the desired output indicates which connections should be changed, layer by layer, in reverse fashion (Lippmann, 1987). The second type is known as "unsupervised learning" because it lacks the teacher element. In this case, only the inputs are used and are classified according to the network dynamics. Such learning is usually applied to tasks aimed at discovering regularities in input statistics. These systems are self-organizing (Linsker, 1988; Kohonen, 1995), and do so by adjusting weights (synaptic efficacies), typically according to some version of Hebb's local learning rule (Hebb, 1949; Churchland and Sejnowski, 1992).

Application of these two learning modes to modeling of brain function encounters some feasibility problems. The most obvious relates to supervised learning, since introducing an all-knowing teacher to a network is biologically unrealistic (Churchland and Sejnowski, 1992). At the other extreme, classical unsupervised learning (e.g. the original Hopfield network (Hopfield, 1982) or principal component analysis (PCA) network (Oja, 1982)), is useful for a restricted, albeit undisputedly important, set of carefully chosen problems. However, this learning mode will never, by definition, discover critical information that does not correspond to the statistical structure of the input (Churchland and Sejnowski, 1992).

4.1. General principles of reinforcement learning

A combination of both these learning approaches, usually placed under the broad heading of supervised learning, is reinforcement learning. The field of reinforcement learning generally deals with situations in which an agent, with an explicit goal, acts upon the environment. The agent's actions change the state of the environment, which in turn provides feedback (reward or punishment) on its actions. In this scheme, external reward functions as an evaluative signal, indicating the degree of appropriateness of network performance. Thus, on the one hand, reinforcement learning is a form of supervised learning, because the network receives and uses feedback information from the environment. However, this information is a scalar value, and is therefore evaluative, rather than instructive (Hertz et al., 1994). The evaluative signal is given either as a single bit of information: right or wrong, or, in the continuous case, as a value describing the degree of correctness. The correct answer itself

remains unknown to the actor (unlike supervised learning). Reinforcement learning is therefore sometimes referred to as "weakly supervised learning" (Pennartz et al., 2000).

In a nutshell, reinforcement learning systems include a number of concepts in addition to the agent and the environment (Sutton and Barto, 1998): a policy describes the agent's actions given various environmental states, i.e. a set of stimulus-response rules or associations. A reward function maps each state (or state-action pair) onto a single number, a reward, indicating the immediate desirability of that state. A value function specifies the long-term desirability of the state, taking into consideration the states that are likely to follow, and the rewards available in them. Ultimately, almost all reinforcement learning methods are structured around estimating value functions; in other words, they attempt to estimate how good it is to be in a given state, or to perform a given action in a given state, with respect to a certain policy. Finally, some reinforcement learning systems involve an internal model of the environment (Suri, 2002).

The basic idea behind various reinforcement learning methods stems from the notion that if an action improves the condition of the actor, the tendency to produce that action is strengthened, according to a rule known as the Thorndike law of effect (Thorndike, 1911). Learning according to this rule is based on the "generate and test" scheme (Barto, 1995). In this scheme, alternatives are generated, tested, and ultimately behavior is directed toward the better alternative (or sequence of alternative actions). The aim of a reinforcement learning system is to maximize the expected return, which is a function of the subsequent reward sequence. If time is divided into discrete episodes, each composed of a finite number of states, their sum is sufficient. In infinite scenarios, a discounted series of rewards should be maximized. Since the reinforcement signal gives no hint as to what the correct answer should be, but only how close the last answer (or set of answers) was, a reinforcement learning network must include a source of randomness to be able to explore the set of possible outputs until an optimal one is found. This notion embodies a key element in a good reinforcement learning strategy: exploration versus exploitation: the agent has to exploit what it already knows in order to obtain a reward, but it also has to explore in order to achieve better performance in the future (Sutton and Barto, 1998).

In animal and human learning as well as in machine learning, the reward is often delayed relative to the behavior that brought it about. Thus, the actor must perform a sequence of actions, (or go through a sequence of states) before it receives information regarding the correctness of the response. Therefore, a frequent difficulty in reinforcement learning is that of temporal credit assignment. The actor has to be able to assign credit and blame individually to each action in the sequence. Thus, temporal credit assignment is critical for biological reinforcement learning algorithms.

4.2. Estimating the value function and policy

A system chooses a policy by evaluating the value of each of its alternatives. The value of an alternative, or more generally, a state, is defined as the sum of rewards received when starting in that state and following a fixed policy in all the future states. Given an accurate value function, extracting the optimal policy reduces to a trivial task. A general solution for this would be repeated iterations of improving the estimate of the state value according to a given policy, followed by improving the policy in view of the change in the value estimate (Sutton and Barto, 1998). Thus, the first objective of a reinforcement learning algorithm is to find an estimation of the optimal value function for each state such that the error of this estimation is zero.

The goal of value approximation is achievable if three basic ideas of reinforcement learning are followed (Barto, 1994). First, throughout learning the agent's policy must remain fixed. Second, for approximation to be possible, one must assume that environmental situations tend to recur, so that approximation is based on recollection of identical (or similar) situations. And finally, variability in state/outcome combinations can be overcome by probability theory. Thus, the value of a state can, in fact, be replaced by the expected sum of future rewards, i.e. the (weighted) average over all possible future scenarios.

4.2.1. Dynamic programming

Assuming the learning system possesses full knowledge of the dynamics of the environment, algorithms known as dynamic programming can be applied. In these algorithms, approximation of the value function can naturally be achieved using the value iteration algorithm. This method assumes that situations recur, and that their consequences are constant. One can find the optimal value function by performing sweeps throughout the state space, updating the value approximation of each state according to the current error of approximation. These iterations are continued until no change is required (i.e. the approximation error is zero). Next, the policy improvement step is performed, in which the current policy is changed if an action is found such that it will improve the state value. This step is repeated until no substantial improvement can be made (Sutton and Barto, 1998).

This section has dealt with the simplistic scenario of a deterministic process. Clearly, not all problems can be classified as such. When a given action in a given state gives rise to a set of potential successor states (with a known probability distribution function) the process is termed non-deterministic. In such cases, the value iteration will seek the maximum expected value. This expected value is calculated across the values of all possible subsequent states following a given action. Dynamic programming is only possible given full knowledge of the system's dynamics, and therefore is poorly suited for modeling biological systems.

4.2.2. Monte Carlo methods

Dynamic programming methods have two obvious disadvantages: first, they require a complete model of the environment, which is not always available. Second, even when such a model exists, it is still very costly to calculate the expected returns, particularly in large state spaces. Monte Carlo (MC) methods overcome both problems. In MC methods, the processes of value estimation and policy improvement are based on sample experiences, rather than on the (unknown) full probability distribution. Basically, algorithms using MC methods average complete rewards observed following visits to each state. If many iterations are performed (i.e. many different episodes in which the state was encountered in any step), this average will converge to the expected cumulative future reward, which is exactly the value of the state (Sutton and Barto, 1998).

Learning in Monte Carlo methods is not truly on-line: the learning system must wait for the completion of the episode, rather than updating step by step. Therefore, it is only partially incremental. The changes occur between episodes (or iterations of the entire task sequence) and not step by step (Sutton and Barto, 1998).

4.2.3. Temporal difference methods

Temporal difference (TD) models are a class of models that deal with the task of predicting future returns for a course of a pre-defined sequence of states. This corresponds to a classical conditioning paradigm, where the agent has to predict the reward associated with a preceding event. TD methods are the most widely used reinforcement learning methods today, partly because they integrate some of the benefits of Monte Carlo methods with those of dynamic programming. They use the model free approach as do the Monte Carlo methods, i.e. they learn directly from experience. However, unlike Monte Carlo methods, they update estimates based, in part, on previously learned estimates (like dynamic programming). In each step, the state value is compared with the subsequent (immediate) reward plus the estimated value of the next state. The result of this comparison, called the TD error, is used to update the state value. Thus, unlike Monte Carlo, TD is naturally implemented on-line, since it is truly incremental: whereas Monte Carlo method has to wait until the final state of the episode to update a state value, in TD methods only a single step has to pass (Sutton and Barto, 1998). Although the TD model computes predictive signals rather than selecting an optimal action (Suri, 2002), there is an extension to decision processes which is discussed in an upcoming section (Section 4.3.1).

4.2.4. TD(λ)-bridging Monte Carlo and temporal difference methods

TD and MC methods are actually two extremes on a continuum. They differ in the number of steps within an episode that are examined for reward, as opposed to those that are estimated using future state estimates. As mentioned, TD learning waits a single step, using an estimate of the subsequent state. MC learning does not make use of any other state estimates. In fact, any combination of actual reward/state estimates can be used (Sutton and Barto, 1998). TD(λ) refers to such a combination: the rewards and estimates of each of the future states are averaged according to a decay factor ($0 < \lambda < 1$). Alternatively, it can be looked upon in reverse fashion: each past state has a decaying memory trace associated with it. The global TD error signal triggers proportional updates to all recently visited states, according to λ: TD(0) is the classical temporal difference method, and TD(1) corresponds to Monte Carlo.

4.2.5. Balancing exploration and exploitation

Both Monte Carlo and temporal difference methods raise the problem of sufficient exploration. These methods will only work if all states are encountered, at least after a large number of iterations. In addition, particularly in model free methods, it would be useful to estimate action values, rather than state values. The Q-learning algorithm (Watkins and Dayan, 1992; Jaakkola et al., 1994) provides a solution to both problems: it assigns values (called *Q*-values) to state/action pairs. The *Q*-value is the sum of (discounted γ) reinforcements received following the action and the given policy in the future. Thus, an optimal *Q*-value is the sum of reinforcements received following the action and the use of optimal policy in the future. In practice, the expected *Q*-value is estimated by a single random sample of a successor state. Depending on the exact algorithm used, Q-learning can be regarded as an off-policy method (one in which while behaving under one policy a different policy is estimated (Sutton and Barto, 1998)) for TD learning or for MC learning.

4.3. Reinforcement learning architectures

Although the concept of reinforcement learning is free of constraints relating to computational or physiological architecture, the implementation of theories of animal classical and instrumental conditioning by experts in artificial intelligence and control theory have produced a number of computationally powerful learning architectures.

4.3.1. Actor/critic architecture

The actor/critic architecture provides reinforcement learning with a solution to the temporal credit assignment problem (Tesauro, 1994; Barto, 1995). In this architecture, a separate unit (critic) receives all the information from the environment that the agent (actor) receives, i.e. the inputs and reinforcement, along with information regarding the actor's output. The critic is intended to predict the reinforcement produced by the environment. Once this model unit is trained, it can be used to calculate the values of each state: since the actor should try to maximize the cumulative reinforcement, rather than the immediate reward, a weighted average of the future reinforcement can be estimated, and fed back, step by step, to the actor. A network (or algorithm)

that is used to predict reinforcement, and feed the difference between its predictions and true reinforcement to the actor is called the adaptive critic (Sutton, 1988), and the entire system is termed actor/critic. The temporal credit assignment problem is solved, since based on the model's predictions of future reinforcement, it provides the actor with immediate evaluative feedback. Thus, an actor that learns with an objective to maximize immediate effective reinforcement, actually acts according to the strategic objective of maximizing some function of future reward (Barto, 1995). This also makes the actor/critic architecture particularly suited to learning action sequences (Suri, 2002). The term "adaptive" emphasizes the fact that the critic adapts (according to its expectations) upcoming rewards to immediate evaluative feedback. Obviously, it is desirable for the critic to improve with experience (Sutton, 1988; Barto et al., 1989). The adaptive critic uses the TD error (see above) to update its weights: if the critic's predictions in adjacent time steps have changed, the adaptive critic should change.

In fact, the output the critic provides to the actor (the effective reinforcement) is most often the same as the TD error as well (Barto, 1995). This is reasonable if one considers the actor's learning objective: it has been repeatedly shown that reward-dependent learning in animals and humans hinges on the degree of the unpredictability of the reward (Rescorla and Wagner, 1972). If a response to a sensory stimulus (an action) has the expected consequences (TD error = 0), then that response tendency should remain unchanged. However, if the consequences are either better (TD error is positive), or worse (TD error is negative) than expected, the response should be strengthened or weakened accordingly.

The actor/critic architecture has been suggested as the basic mechanism by which the basal ganglia function (Houk et al., 1995). According to this model the reinforcement signal is generated by the dopaminergic neurons as a result of input from a primary reinforcement source (such as the lateral hypothalamus) and input from the MSNs within the striosomes. The reciprocal loop between the striosomal MSNs and the dopaminergic neurons, which together are called the "striosomal module", actually performs as an adaptive critic which calculates the error signal for the actor. On the other hand the MSNs in the matrix pass the information to the frontal cortex through the globus pallidus and thalamus, thus forming the "matrix module", which functions as the actor part of the actor/critic architecture.

4.3.2. Internal models

A full representation of the environment by the actor can be very useful for planning (Sutton and Barto, 1998; Suri, 2002). Situations where optimal performance requires calculation of several steps ahead and evaluation of the values of complex hypothetical future states of the environment could be much assisted by a computational unit which serves as a predictor. This model will not only predict future reinforcement, but whole state-spaces. Internal model approaches simulate future moves, and use hypothetical future outcomes to select the best move. This approach has been used to model the dopamine neuron activity and anticipatory neural activity in striatum and cortex for correct action selections in situations that require planning (Suri et al., 2001).

5. Dimensionality reduction

Dimensionality reduction describes the process of projecting inputs from a high dimensional data space to a lower dimensional space (Haykin, 1999). Dimensionality reduction in the nervous system can be depicted as compression of the information encoded by a large neuronal population to a significantly smaller number of neurons. Efficient reduction is achieved when all or most of the information contained within the original space is preserved. For example, population coding has been frequently associated with neural coding (Georgopoulos et al., 1986). This coding involves combining the activity of a population of neurons to encode a specific variable (such as the angle of movement of a limb). Extraction of the underlying variable may replace the need to keep or process the full and detailed activity of the whole population. Thus, we can transfer the data from a large dimension equal to the number of neurons in the encoding population (where the size of the population is the superficial dimension of the data) to a single neuron (which makes up the intrinsic dimension of the data) encoding the relevant variable (angle) by its firing rate. Another example can be drawn from the sensory world (Barlow, 1989, 1992b). In the visual system the input, made up of a huge number of activities of retinal cells, is transformed into the activity of a simple cell representing a line in the cortical V1 area. The idea of transforming the correlated information of a large population into a minimal number of active cells which Barlow (Barlow, 1992a) termed "cardinal neurons", was suggested to serve as a major mechanism for information encoding in the brain.

5.1. Motivation

One of the major tasks that the brain faces is enabling the organism to make the transition between successive states. This transition can be simplified into the transformation between the organism's current state and the next action to be performed. The computation of the upcoming action requires information from different modalities stored in many diverse areas of the brain that process sensory, motor, limbic and cognitive information. However, multiple problems arise during the transfer of all this information to the areas involved in planning such as the frontal cortex. These problems, which stem from the need to transfer huge quantities of information, are surprisingly similar to the ones encountered in a completely different domain, that of mass communication while trying to connect geographically dispersed information sources. Connecting n sources with m destinations may become a huge problem that hinders the flow of information. This problem may be broken down into four main sub-problems:

1. Number of communication lines: $n \cdot m$ lines of communication must exist between the source and the destination.
2. Number of communication endpoints: n connection points must be created on each destination, leading to a total of $n \cdot m$ connections.
3. Communication flexibility: any changes in the source or destination points such as a new point, removal of a point or a change in connectivity demand physical changes in both the communication lines and endpoints.
4. Communication modulation: any modulation of the information flow such as filtering, enhancement, etc. must be performed separately for each connection.

The same problems are even more severe in the domain of the central nervous system, which is naturally limited in both its physical size and the resources available to create and maintain the system. For example, in the human cortex, an order of 10^9–10^{10} neurons (Shepherd, 1998) encoding information regarding the current state must be connected to about the same order of magnitude of neurons in the frontal cortex (Abeles, 1991; Braitenberg and Schuz, 1991). Linking the frontal cortex via 10^9 incoming axons and creating 10^9 synapses on each neuron does not seem feasible. A typical neuron in the cortex has only 10^3–10^4 synapses (Abeles, 1991), which constitute all the information it receives. In addition, no central modulation is possible in such a dispersed system, making any central control mechanism of the information flow inoperative. We propose that the central nervous system uses a solution similar to the one used by communication systems; namely a central switch that receives all the information from the large number of source elements. The switch encodes the information by prioritizing and compressing the data and transferring it via a limited number of channels to the destination.

In addition to the connectivity problem, the distributed encoding of the brain leads inevitably to a major computational problem. The "curse of dimensionality" (Belman, 1961), is a term used for a group of computational problems arising from data residing within high dimensions. The number of samples needed to describe a certain space increases exponentially. A neural network requires resources proportional to the hyper-volume of the input space, which increases exponentially with its dimension. In the biological world this means that the neurons handling the planning and selection of actions would require an unrealistically large amount of examples to learn the correct mapping from the current state to a specific action. A decrease in the number of the needed inputs to achieve useful mapping from states to actions could be achieved by mapping the input into a lower dimensional space containing the key features of the input.

5.2. Different reduction methods

A large number of dimensionality reduction methods exist, differing first and foremost in terms of the reason for the reduction. In addition, dimensionality reduction performed for the same purpose may use different methods depending on the constraints set on the learning and processing system.

5.2.1. Data and dimensionality reduction

Reduction can be achieved in two basic ways: by reducing the dimensionality of the data, or by reducing the data itself. During a reduction of dimensionality, the original data are embedded within a lower dimensionality space. However, during a data reduction process, the data are associated with a limited number of patterns that represent them. Whereas dimensionality reduction methods such as PCA typically try to maintain a certain information criterion regarding the data, data reduction methods such as vector quantization (VQ) and clustering attempt to remove the variability of the input and transform it into one of a given set of possible outputs. This review focuses on dimensionality reduction methods that enable minimization of information loss, rather than methods which involve encoding the data into specific patterns while losing most of the information.

5.2.2. Supervised, unsupervised and reinforcement based dimensionality reduction

The type of dimensionality reduction is derived from the method of learning the "correct" reduction. In supervised learning, a teacher provides the optimal reduction during the learning phase. Thus, supervised algorithms learn to imitate the results achieved by such a teacher. These methods are usually not biologically plausible due to the lack of a teaching signal. At the other extreme, unsupervised learning algorithms do not use any teacher. Unsupervised learning algorithms rely on the properties of the data set itself to learn the optimal reduction. It is important to note that supervised and unsupervised methods for dimensionality reduction may achieve the same result. However, the way of arriving at that result will differ significantly; for example a PCA algorithm may be achieved via a supervised neural network (Bourlard and Kamp, 1988) or an unsupervised neural network (Foldiak, 1990; Kung and Diamantaras, 1990).

5.2.3. Local and global dimensionality reduction

A major difference between reduction methods is the way they perceive the interactions between inputs. Global dimensionality reduction methods (such as classical PCA) assume that all of the inputs share a global correlation scheme. Thus, all of the inputs can be reduced to a single smaller subspace. Local dimensionality reduction, on the other hand, assumes that there are local interactions between subsets of the data leading to a different optimal reduction for each of the subsets (Chakrabarti and Mehrotra, 2000). Such methods enable efficient reduction when the original input space contains several very different sub-structures. The local structures are reduced in a separate manner, which is locally optimal regardless of global criteria for the whole input space.

5.2.4. Linear and non-linear dimensionality reduction

Dimensionality reduction methods have traditionally been divided into linear and non-linear methods. Linear methods attempt to find a sub-space to represent the linear combination of the inputs. The more general non-linear methods attempt to find a sub-space that represents a non-linear combination of the inputs (Oja et al., 1995; Malthouse et al., 1995). Linear methods typically have the advantage of simplicity and easier implementation. However, their success depends entirely on their underlying assumption that there exists a good representation of the input as linear combinations. When this is not the case, such as input derived from a sphere, the reduction will fail completely. On the other hand, non-linear methods tend to be more complex but enable a compact representation of more complex distributions of inputs.

5.3. Principal component analysis (PCA)

Principal component analysis is typically an unsupervised, linear and global dimensionality reduction method. However, there are variants that use supervised (Bourlard and Kamp, 1988), non-linear (Oja et al., 1995; Malthouse et al., 1995) or local (Chakrabarti and Mehrotra, 2000) dimensionality reduction methods. Although PCA is by no means the most powerful or the most general method for dimensionality reduction, the rest of the review will focus on this method. The reasons for choosing PCA are its simplicity, the fact that it has been researched extensively, and that it has a biologically plausible implementation using neural networks.

PCA is a mathematical process, also known as the discrete Karhunen–Loéve transform (KLT) and singular value decomposition (SVD) of the covariance matrix. PCA projects an n dimensional input onto m ($m < n$) orthogonal axes containing the maximal variance of the original input. PCA is primarily useful due to its information preservation properties. For a given reduction to m dimensions, PCA is the linear transformation leading to the minimal reconstruction error of the output from the input (Gerbrands, 1981) and is also the transformation that preserves the maximal amount of Shannon information (Linsker, 1988).

Over the last 20 years, neural network models for performing PCA have evolved considerably. The earliest work, published by Oja (1982), was based on a learning rule for a linear single-unit network. This network extracted only the first principal component of the input data. The learning rule of the neuron was based on a modified Hebb's rule, which accordingly was termed the "normalized Hebbian" learning rule. In the standard Hebbian rule (Hebb, 1949), if the input and the output neurons have positively correlated firing rates the strength of their connection (weight) increases; alternatively, if they tend to fire in a negatively correlated manner, the strength of their connection decreases. This standard Hebbian rule tends to lead to unbounded growth in the strength of the connection between the neurons. For this reason, a decay factor is incorporated which is reduced from the value of the connection leading to a normalized Hebbian learning rule. The biological plausibility of such a normalized Hebbian rule has been demonstrated (Friston et al., 1993). This learning process leads to the convergence of the weights to the first principal component and maintains the maximal variance of the input space by utilizing a completely local algorithm.

The first neural network models for extracting more than a single component were the generalized Hebbian algorithm (GHA) (Sanger, 1989) and the stochastic gradient algorithm (SGA) (Oja and Karhunen, 1985). These models generalized Oja's rule to a multi-unit output stage. In these models, there are n input neurons, and m output neurons ($m < n$) and the learning rule leads to the extraction of the data's m first principal components. The extraction is performed using only feed-forward weights between the inputs and the outputs through introduction of a global quantity which is not locally available to the individual neurons. This global quantity significantly reduces the plausibility of such a neural network models for biological systems.

Foldiak extended Oja's model by symmetrically connecting the output neurons to each other by inhibitory connections (Foldiak, 1989; Foldiak, 1990). The output of the neuron is obtained by subtraction of the activity transferred by the lateral weights from the activity derived from the feed-forward network. The feed-forward weights are trained by a Hebbian rule similar to Oja's learning rule, and the inhibitory lateral weights are trained with an anti-Hebbian learning rule. Anti-Hebbian learning rules lead to the reduction of the connection if the input and output neurons fire together, and the enhancement of the connection if they tend to fire in an opposite manner. Thus, inhibitory lateral connections become stronger (more negative) when the firing of the two neurons is correlated. As in Sanger's model, this model extracts the input data's m principal components. The lateral weights in this model replace the global quantity and act as a local mechanism for transferring information between the neurons of the output layer, leading them to encode different components of the input. It is important to note that in this model the network extracts the principal subspace containing the same information, however the feed-forward weights do not contain the principal components themselves.

The model devised by Kung and Diamantaras (and an earlier model (Rubner and Tavan, 1989; Rubner and Schulten, 1990)) uses hierarchical lateral connectivity between the output layer neurons (Kung and Diamantaras, 1990; Diamantaras and Kung, 1996) in which neuron i is connected to neurons $i + 1$ through n. In this model, which is called adaptive principle component extraction (APEX), the feed-forward weights are trained by a Hebbian rule similar to the one used by Oja, and the inhibitory lateral weights are trained by an anti-Hebbian learning rule. The network converges to a solution in which the lateral weights are zero and the feed-forward weights encode the first m principal components. The weights of the ith output neuron are equal to the ith eigenvector of the input and the

activity of the *i*th neurons is equal to the projection of the input on the *i*th principal component. Most importantly, the model presented by Foldiak and the one presented by Kung and Dimantaras do not use any global information, hence maintaining their biological plausibility.

6. Basic reinforcement driven dimensionality reduction model of the basal ganglia

The basic reinforcement driven dimensionality reduction model (Bar-Gad et al., 2000) is an extension of a neural-network based model for performing principal component analysis. Two different unsupervised neural networks serve as the basis for the model. The first network is based on work by Foldiak (1989) using a symmetric neural network, i.e. every pair of neurons within the same layer is linked by reciprocal connections. The second network utilizes the APEX network developed by Kung and Diamantaras (1990), which uses an asymmetric neural network, i.e. for every pair of neurons in the same layer only one neuron affects the other. This type of connectivity is also known as hierarchical in some implementations. Both networks use all-to-all connectivity between the input and the output layers and utilize the same basic mechanisms for the neuron's activity and learning. The neuron activity function in both networks is linear, i.e. the output is the weighted sum of the feed-forward inputs minus the weighted sum of the lateral inputs. The learning rules are extensions of the normalized Hebbian rule (Oja, 1982) for the feed-forward weights and extensions of anti-Hebbian rule for the lateral network.

The RDDR model adds a reinforcement factor to the unsupervised learning of the two networks. This reinforcement factor is multiplied by the input and the output of the neuron to create a multi-Hebbian learning algorithm. In a multi-Hebbian learning rule the reinforcement signal regulates the amount of change in the weight of a synaptic connection for a given input/output pair (Reynolds and Wickens, 2002). This rule is inspired by the complex cortico-striatal-dopaminergic synapse triad (see also Section 2.5.1). Most of the cortical input is received in a synapse located on the spines (Freund et al., 1984; Gerfen, 1988) of the medium spiny neurons (MSN) which are the projection neurons of the striatum. The dopaminergic (reinforcement) signal is relayed from the SNc to the same neurons and creates a synapse on the neck of the spine (Freund et al., 1984; Bouyer et al., 1984). Physiological studies have shown that the dopaminergic system can control the learning process of the cortico-striatal synapse (Calabresi et al., 1997; Kerr and Wickens, 2001).

6.1. The model

The basic RDDR network is constructed of two layers (input and output) with a modulating scalar value representing the reinforcement (prediction error) signal. The number of neurons in the input layer of the RDDR model is larger than the number of neurons in the output layer; in other words the input space of the network has a higher superficial dimension than its output space (Fig. 2a). However, the input may actually lie within a significantly lower sub-space of the input space that makes up the intrinsic dimension of the input data. For example, a typical input is the lines within a matrix (Rumelhart and Zipser, 1985; Foldiak, 1990). The input space is made up of a matrix of size $n \cdot n$ (making up a superficial dimension of n^2). However, the input only contains a combination of vertical and horizontal lines within this matrix. Each element of the input matrix is generated

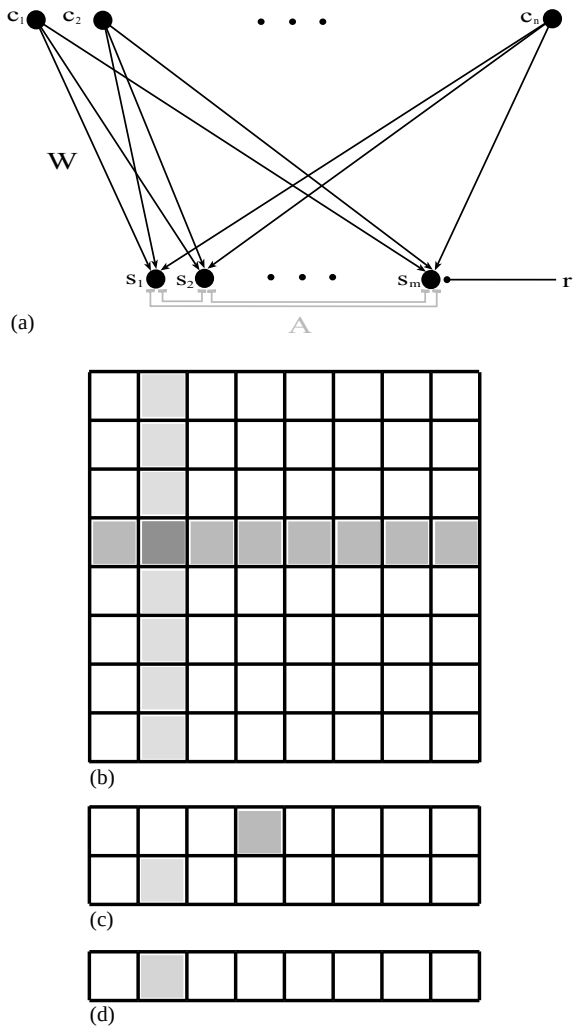

Fig. 2. The basic RDDR model. (a) Structure of the basic RDDR network, C: input (cortical) neurons; S: output (striatal) neurons; W: input/output (cortico-striatal) feed-forward weights; A: output layer (striato-striatal) lateral weights, r: reinforcement (SNc) signal. (b) Example of an input to the system from an input space with a superficial dimension of 64. (c) Reduction of the example into the intrinsic dimension of 16 in which each line (horizontal or vertical) of the input is represented by the activity of a single element of the output. (d) In the case of reinforcement given only to vertical lines, reduction of the example can be made into the smaller dimension of eight output elements.

from a linear transformation (i.e. line summation) of a $2n$ dimensional source vector, whose components are independent. This definition of the input matrix leads to an intrinsic dimension of $2n$ ($\ll n^2$, for large n) (Fig. 2b). In an optimally encoding network, assuming that the output is encoded by m neurons, the network can encode all of the information when the number of neurons is larger than the intrinsic dimension of the data (i.e. $m = 2n$) otherwise the network will encode the m dimensional subspace which contains the largest portion of the input information (Fig. 2c). To study selective reinforcement in the RDDR model, a differential reward signal is given for different patterns (Fig. 2d).

The activity of the output neurons (s) is a linear weighted sum of its feed-forward inputs ($w \cdot c$) and lateral ($a \cdot s$) inputs:

$$s_i = \sum_{j=1}^{n^2} w_{ij} c_j + \sum_{j=1}^{m} a_{ij} s_j. \quad (1)$$

The learning rule for the feed-forward weights (w) is a normalized multi-Hebbian rule (Kung and Diamantaras, 1990), combining feed-forward and reinforcement (r) signals with a learning rate constant (η).

$$\Delta w_{ij} = \eta r [s_i c_j - s_i^2 w_{ij}]. \quad (2)$$

The learning rule for the lateral weights (a) is a normalized anti-Hebbian rule (Kung and Diamantaras, 1990), except for the self-connection which is set to 0.

$$\Delta a_{ij} = -\eta [s_i s_j + s_i^2 a_{ij}], \quad a_{ii} = 0. \quad (3)$$

When using an asymmetric organization, all $a_{ij} = 0$ when $i < j$. For simplicity the same learning rate constant (η) is used for both the feed-forward and lateral weights. To measure the information loss of the network due to the RDDR process, the m dimensional representation of the input patterns in the output layer is expanded back to an n^2 dimensional space to create the reconstructed, decompressed pattern. The reconstruction error is the mean squared difference between the original and reconstructed elements over all input patterns.

6.2. Results

The network displays a dynamic pattern of activity which, given inputs from a certain distribution, goes through a learning process leading to stabilization. Several critical parameters display important changes during the learning process.

6.2.1. Correlation

The input to the network (Foldiak, 1990) contains significant correlations since all neurons on the same row and columns of the matrix tend to change their rate together, simulating the correlated input from the neighboring areas in the cortex (Eggermont, 1990). The correlated input that converges into the network is expected to produce correlations within the output layer. However, the dynamics of the lateral inhibitory connections lead to a more complex change in the correlation (Fig. 3a). Initially, the output neurons display correlated activity since multiple output neurons encode the same aspects of the input. However, over time the lateral inhibitory network causes an orthogonalization of the activity of the output neurons, leading to an uncorrelated firing pattern (Foldiak, 1989). The uncorrelated state is maintained by the pattern of efficacies of the feed-forward weights. This uncorrelated firing pattern is very different from the predictions of the action selection models. In such models the correlation between any two neurons in the output stage is predicted to be highly negative (i.e. the increased firing rate of one cell leads to a decrease in the rate of all of its neighbors) and to remain negative in the steady state of the network. The experimental data regarding spiking activity in the striatum (Jaeger et al., 1994; Stern et al., 1998) and the globus pallidus (Nini et al., 1995; Raz et al., 2000; Heimer et al., 2002; Stanford, 2003; Bar-Gad et al., 2003) indicate weak or non-existent correlations, thus resembling the predictions of the RDDR model.

6.2.2. Lateral connectivity

The lateral connectivity of the network is inhibitory in nature and assumes an anti-Hebbian learning rule. In such

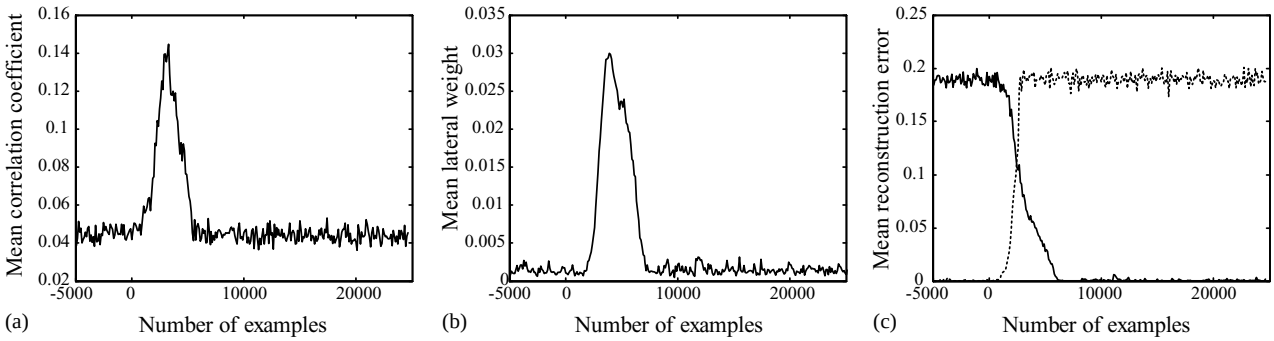

Fig. 3. Dynamic changes in properties of the RDDR network. Changes occurring in the network following a shift in the rewarded set of inputs at time 0. (a) Mean correlation between neurons of the output layer; (b) mean value of the lateral weights between neurons of the output layer; (c) mean reconstruction error for the newly rewarded patterns (solid line) and previously rewarded patterns (dotted line). Simulation of 64 input neurons, 8 output neurons, only-positive feed-forward weights, and only-negative lateral weights.

a learning rule, if both neurons fire together the efficacy of their connection decreases (i.e. becomes more negative or more inhibitory) whereas if they do not fire together the efficacy increases. The lateral weights that model the typical GABAergic collaterals of the basal ganglia serve therefore as a means of decorrelating the output of different neurons, leading them to encode different properties of the input. The beginning of the network learning process leads to a correlated activity of neurons in the output layer. This correlated activity results in the negative values of the lateral weights. This process leads to the decorrelation of the output, which eventually causes the reduction of the values of the lateral weights to very low values (Fig. 3b). At this stage of steady-state, the encoding of the principal components is performed by the feed-forward weights (Fig. 4). The lateral weights, therefore, play a crucial role in setting the efficacies of the feed-forward weights but not in the steady state encoding of the ongoing input. This is very different from the role of the lateral weights in action selection models which use a winner take (or lose) all mechanism (Wickens, 1993, 1997; Berns and Sejnowski, 1996). In such models, the lateral weights are significant in the ongoing processing of the input during steady state. Physiological evidence shows that the lateral weights in the striatum (Jaeger et al., 1994; Tunstall et al., 2002; Czubayko and Plenz, 2002) and the globus pallidus (Stanford, 2003) display weak and asymmetric inhibition of the neighboring neurons closely resembling the expectations of the RDDR model.

6.2.3. Information encoding

Information encoding of the network is a key measure of its ability to maintain the information encoded in its input as it is being passed to a significantly smaller number of neurons in the output stage. In the initial state, the encoding of the information is random, leading to an inability to reconstruct the input. Thus, most of the information is lost and cannot be recovered from the activity of the output neurons. The learning process improves the encoding of the information significantly as measured by the decrease in reconstruction error (Fig. 3c). In cases of an input space with a lower intrinsic dimension than the number of neurons in the output stage, the information is encoded fully in the activity of the output neurons and no information is lost. Hence, full reconstruction is possible from the activity of the network output. In cases where the intrinsic dimension of the input space is larger than the number of neurons in the output stage, the output neurons encode the dimensions of the input which contain the maximal information. Thus, optimal (although not full) reconstruction may be achieved. This encoding is very different from the data reduction scheme suggested by the action selection models. In such selection models, only one chosen action is encoded while all other aspects are lost. This leads to a very low information capacity of the network, which loses most aspects of the input. Thus, assuming for simplicity m binary output neurons, an action selection network would be able to encode m patterns whereas the RDDR network would be able to encode 2^m patterns.

6.2.4. Reinforcement signal

The reinforcement signal constitutes a control signal that modulates the Hebbian learning rule of the feed-forward network, i.e. it changes the structure of the input space as it is seen by the network. It increases the variability of dimensions associated with reward activity while diminishing the relative variability of dimensions which are not. This leads to a network that does not simply encode the maximal variability of its input space but rather encodes the variability of the reward-distorted space. Thus, information regarding the reward-related activity is maintained in optimal form while the non-rewarded information is not encoded and therefore

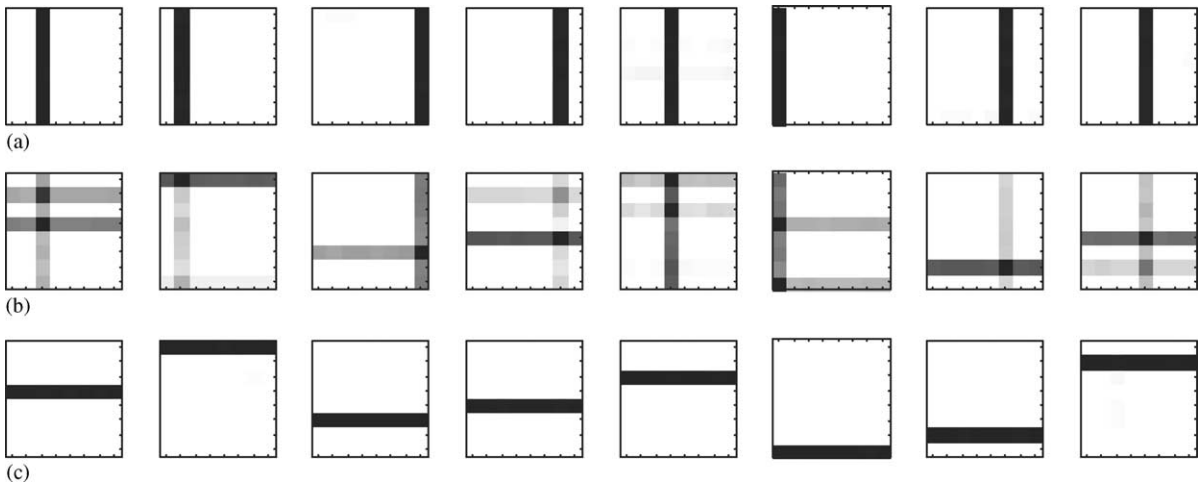

Fig. 4. Changes in the encoding of the feed-forward weights. The network is presented with both horizontal and vertical lines. Until time 0, vertical lines are associated with the reinforcement signal. Following time 0, the horizontal lines are associated with the reinforcement signal. The value of the eight feed-forward weights is shade coded in 8 × 8 matrices for times: (a) −5000 examples; (b) +2000 examples; (c) +10,000 examples. Simulation of 64 input neurons, 8 output neurons, only-positive feed-forward weights, and only-negative lateral weights.

cannot be reconstructed (Fig. 3c). Using the reinforcement signal, the network ceases to encode information based solely on its statistics in an unsupervised manner. Rather, it neglects unimportant information and better encodes information that is deemed important by the reinforcement signal.

6.2.5. Pathologies

Diseases of the basal ganglia are associated with some of the most severe neurological disorders such as Parkinson's disease. These diseases are characterized by severe motor dysfunctions and additional cognitive and limbic problems (Sethi, 2002). PD is characterized by death of the midbrain dopaminergic neurons (Ehringer and Hornykiewicz, 1960). Normally, there is a steady state level of dopamine in the striatum maintained by these cells which may be increased due to better than predicted events or decreased due to disappointing events (Schultz et al., 1997; Schultz, 1998; Fiorillo et al., 2003). During PD, the normal levels of dopamine release decrease due to neuronal death in the SNc leading to lower steady state levels in the striatum (Bernheimer et al., 1973; Schultz, 1982; Hornykiewicz and Kish, 1987). In the RDDR model, this is equivalent to a negative reinforcement signal for all the inputs given to the network. Negative reinforcement signals for all the inputs cause a breakdown in the encoding of the system. This in turn leads to a very low information capacity of the network and increased correlations between neurons of the output level. The major line of defense against PD is dopamine replacement therapy—either by the dopamine precursor L-Dopa or by post-synaptic dopamine agonists. The non-specific (in time and space) effects of such dopamine replacement therapy are best understood in the framework of background dopamine activity representing a match between prediction and reality, and therefore preservation of the current network. The major side effect of these treatments is the development of hyper-kinetic disorder—Levodopa-induced dyskinesias (LID). The RDDR model provides a natural explanation for the development of LID. Dopamine replacement treatment leads to a pulsatile level of dopamine in the plasma (and probably in the striatum) (Shoulson et al., 1975; Nutt et al., 2000). The levels of dopamine cease to correlate with the actual performance of the subject but rather only with the time of medication. Thus, the network continuously goes through changes in the encoding of the information which is mistakenly considered as reinforcing. Such mistaken encoding of the information may lead to the activation of the "wrong" muscles during movement leading to dyskinesias.

7. Advanced reinforcement driven dimensionality reduction models of the basal ganglia

The basic RDDR model is primarily a conceptual representation of the basal ganglia in gross lines. The model's simplicity is important for purposes of exploring its main ideas and predictions within the specific biological system. The basic model is simplistic in its assumption of a two layer network in which all neurons are connected to all others, both within each layer and between layers. The basic model also makes the simplistic assumption that neurons are linear and the modifiable weights between them are not constrained to any value. Refining these assumptions of the basic model can yield additional insights into the function of the basal ganglia. The expanded RDDR model is able to incorporate additional aspects of the experimental data. These additions to the model are either on the network structure level or on the single neuron level. On the single neuron level, the additional features include constraints on the polarity of the connections between the neurons and non-linear activation functions. On the network structure level, the additional features include multiple layered networks, changing the pure feed-forward structure into a partially closed loop, the addition of sparse and ordered connectivity between and within the layers and multiple pathways between the layers.

7.1. Constrained weights

The basic RDDR model utilizes non-limited weights, as found in network-based PCA models (Foldiak, 1989; Kung and Diamantaras, 1990). These weights, used for both the feed-forward and lateral connections, may reach unconstrained negative and positive values and can even switch between positive and negative values. In general, this is problematic due to fact that in most cases a single transmitter is either excitatory or inhibitory. The cortex, thalamus and STN projection neurons utilize glutamate, which is considered excitatory, whereas the projection neurons of the striatum and globus pallidus utilize GABA which is considered inhibitory. There are a few exceptions in which a neurotransmitter may act as both excitatory and inhibitory, such as GABA (Wagner et al., 1997; Chavas and Marty, 2003) and glutamate (Katayama et al., 2003). However, these exceptions are rare, and results from the striatum (Tunstall et al., 2002; Czubayko and Plenz, 2002) and our preliminary in-vitro studies of the globus pallidus (Rav-Acha, Bergman and Yarom, unpublished results) failed to find non-inhibitory effects of GABA in these structures.

Limiting the weights of the connection between the neurons to a single sign is equivalent to the non-negative constraint applied in non-negative matrix factorization (NMF) (Lee and Seung, 1999). A constraint to a single sign, either positive or negative, means that the relationship between the features encoded by each neuron can only be additive. Additive encoding means that no subtraction is possible between the outputs of different neurons. This in turn leads to an encoding of parts of the input space by each output neuron. Encoding parts of the space is a form of local encoding or sparse encoding that replaces the global encoding performed by PCA. The results obtained for non-negative matrix factorization are extremely useful in their relationship to the RDDR model since they enable, in addition to the obvious weight constraint, a coding which better suits the

topographic innervation of the various layers and the sparseness of the encoding (see also Section 7.6). Other research has shown that such a non-negative decomposition may be achieved using lateral inhibition (Plumbley, 2001). This study shows that using the non-negative constraint, a network may perform independent component analysis (ICA) separating the input not only into its decorrelated components (like PCA) but into its independent components. Independent components differ from decorrelated components by the fact that the minimization includes higher order and not only second order statistics. Limiting the weights to a single sign is also a form of non-linearity leading to the benefits achieved by non-linear encoding (Section 7.2).

However, despite the advantages of single sign constraints, it is crucial to remember that in real biological environments, such constraints may not be "hard" constraints. Some input neurons may convey information opposite to that conveyed by other input neurons. For example, neurons in the motor cortex encode directions, and one neuron might encode a specific direction whereas another might encode the opposite (Georgopoulos et al., 1986)). When both input neurons, conveying the feature and its opposite, are connected to the same output neuron using the same neurotransmitter (and therefore the same constraint), they provide supplementary information. In such cases the two inputs function as the two symmetric parts of the information. This forces a situation in which at any given time during the learning process, only one of the input neurons effectively activates the output neuron. The total effect is the removal of the single sign constraint because of the complementary effect of the opposing input neurons.

7.2. Non-linear elements

The basic RDDR model performs PCA by using linear neurons. Neurons with a linear response function transform the sum of their weighted inputs to firing rates, in a linear manner by simple summation. However, neurons in the basal ganglia, like all neurons in the brain, are limited in their upper and lower extremes of firing rates. At such rates they become saturated and cannot increase (or decrease) their rates. For example, a lower extreme of the firing rate may be zero since neurons cannot switch to negative firing rates. In addition to saturation, neurons in the different nuclei of the basal ganglia have very different response functions which are not necessarily linear, even between the saturation extremes. On the one hand, neurons in the globus pallidus have a linear response (Nakanishi et al., 1990, 1991; Kita and Kitai, 1991; Kita, 1992) and a broad range of firing rates (DeLong, 1972; Miller and DeLong, 1987). On the other hand, the MSNs of the striatum are non-linear with two basic subthreshold states. In the down state, MSN neurons have zero or a very low firing rate while a very high firing rate can be reached in the up state (Nisenbaum and Wilson, 1995; Wilson, 1995; Wilson and Kawaguchi, 1996; Stern et al., 1997; Bennet and Wilson, 2000). The intermediate dynamic range of these neurons is small, leading mostly to one of the extremes states.

The addition of non-linearity to the network enables an increase in the scope of the computation performed by the network. Linear versions of PCA may find only linear interactions between the input, and the optimal connections formed are based solely on second order statistics (covariance of the activity). In cases of non-linear interactions between the input elements and higher order statistics of interaction, a linear network will not be able to reduce the dimensionality of the input effectively and maintain the information within it. Oja et al. (1991) has shown that using the same constrained Hebbian learning rules employed in PCA networks with a non-linear transfer function (typically a sigmoid) enables learning the structure of complex non-linear inputs embedded within a noisy environment. The addition of non-linearity elements (Jutten and Herault, 1991; Karhunen and Joutsensalo, 1994) makes possible the extraction of higher order statistics and the formation of independent components at the output layer, thus enabling the separation of the mixture represented by the inputs into its independent underlying causes.

An important feature of extremely non-linear computation is generalization. Linear encoding transforms different inputs into unique outputs. On the other hand, the extreme non-linear cases of 0/1 transformation or the sigmoid transfer function transforms a group of inputs into a single output. The input is thus divided into groups (or clusters) according to the transformation output. One such non-linear network was implemented by Carlson (1990). In his network, the activity of the neurons depends on the neuron's threshold and the transition width. The width defines the neuron's non-linearity: when the width is 0, activity is highly non-linear (binary), and when it is large, activity is linear. Both the threshold and the width are adjusted during learning. The output neurons are hierarchically organized (like the APEX linear network and others (Kung and Diamantaras, 1990; Rubner and Schulten, 1990)). The outcome of this algorithm in the non-linear case is a partition of the input into clusters, whereas in the linear case it achieves a continuous linear representation.

7.3. Multiple layers

The basic RDDR model is a two layer (input and output) network. However, in reality the basal ganglia form a multi-layered network: the input from the cortex is received within the first (basal ganglia input) layer—the striatum and STN, and from there it is passed to the next (basal ganglia output) layer—the GPi and SNr. The benefits of such a multi-layer feed-forward network stem from two aspects: the ability to perform the same processing sequentially and the ability to perform a different type of processing by each layer.

Sequential processing, performed by a multi-layer (or hierarchical) dimensionality reduction network, has two main advantages arising from the ability to perform the same

computation in several stages (Barlow, 1992b):

- A single neuron cannot directly access a large portion of the input due to the aforementioned limits on the number of synapses. A multi-layer structure enables access to a larger portion of the information within the input in a gradual convergence process.
- The huge size of the input space impedes efficient learning of its representation due to its sparseness. The problem is therefore broken down into multiple layers, where the first layers learn a limited number of input patterns and subsequent layers learn larger portions of the input space.

In the basal ganglia, each of the two layers gains access to up to 10^3–10^4 input neurons (Zheng and Wilson, 2002; Yelnik, 2002). Theoretically, the combination of the two layers enables the neurons of the output layer (the pallidal neurons) to gain access to a large part of the cortical input (up to 10^7 cortical neurons). This number may be further increased because of the partially closed structure of the cortico-basal ganglia loop (Section 7.4) and the multiple feed-forward pathways (Section 7.6).

The second aspect of multi-layer dimensionality reduction is the ability to perform a different type of processing within each layer. Multi-layer networks enable the combination of linear and non-linear layers within a single network. Work by Oja (1991) has shown a network containing a non-linear (encoding) layer followed by a linear (bottle-neck) layer performing non-linear and linear PCA, respectively. This network is able to perform any continuous mapping from the input space to an output layer which is smaller in size. This network is surprisingly similar to the multi-layer structure of the main feed-forward pathway in the basal ganglia (Fig. 5a). Cortical and thalamic inputs reach the MSNs of the striatum, which are non-linear and have a transfer function much like the sigmoid used by Oja. From the striatum, the information is passed to the globus pallidus, which is characterized by neurons with a relatively linear transfer function. Oja's work also speculates on the role of two additional layers in decoding information processed by the initial encoding layers. These decoding layers fit nicely within the cortico-basal ganglia-cortico loop as an equivalent to the thalamus and frontal cortex. Further research (DeMers and Cottrell, 1993; Malthouse et al., 1995) on the practical and theoretical implications of such multi-layer non-linear networks has demonstrated their enhanced capabilities in complex feature extraction which cannot be performed by linear networks.

7.4. Partially closed loop

The basic RDDR model simulates a feed-forward network with an information flow from most of the cortex to the basal ganglia. However, this funnel structure is only part of a partially closed loop (Fig. 5b). The output of the basal ganglia reaches the ventral lateral, ventral anterior, mediodorsal and centromedian nuclei of the thalamus (Haber and McFarland, 2001). The neurons of the thalamus project back to the striatum (McFarland and Haber, 2001) and make up as much as 50% of the striatal input (Kemp and Powell, 1971; Bolam et al., 2000). In addition, the output of these thalamic nuclei reaches the prefrontal, premotor, supplementary motor and motor cortex (McFarland and Haber, 2002; Middleton and Strick, 2002) that form a large part of the cortical input to the striatum (Parent and Hazrati, 1995). This structure forms a

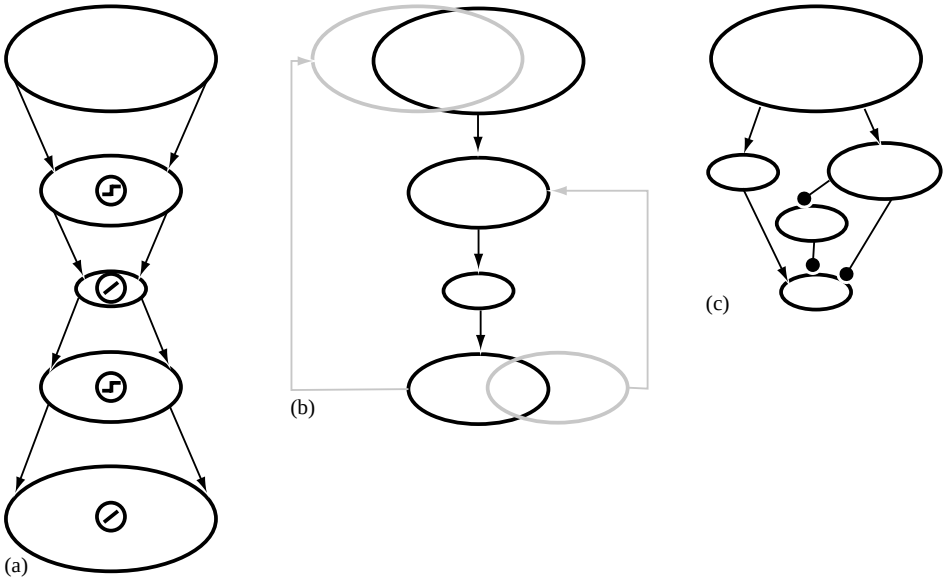

Fig. 5. Enhancements of the RDDR network. Enhancements of the basic network structure: (a) non-linear multi-layer network, the globus pallidus and cortex are represented by layers containing linear neurons while the striatum and thalamus are represented by layers containing non-linear neurons; (b) recurrent network creating a partially closed loop using the feedback to the striatum from the layers representing the cortex and the thalamus; (c) network with multiple pathways representing the direct, indirect and hyper-direct pathways of the basal ganglia.

partially closed loop in which information that is transmitted from the basal ganglia is fed back as its input together with information from other sources (Joel and Weiner, 1994).

The temporal delay of the neural loop structure enables merging of information presented at different times. As a result, the partially closed loop appears to be an optimal substrate for sequential learning as well as sequence detection and generation. Indeed, earlier research has shown the involvement of neurons within different nuclei of the basal ganglia in sequential activity in animals (Kermadi and Joseph, 1995; Nakahara et al., 2001). Several models of action selection have been expanded to include the selection of sequential actions (Berns and Sejnowski, 1998; Beiser and Houk, 1998) (see also Section 3.4).

The RDDR model provides an additional perspective into information processing with the addition of a partially closed loop. In such a structure, the RDDR network receives not only the input representing a single point in time but also additional information about the prior inputs to the network. By creating an input to the network which is partially dependent on its output, the RDDR network performs a reduction which is dependent not only on properties of a single temporal event but also on prior encoded information. The loop structure can lead to preferential representation of sequences that lead to a reward signal. The network will therefore reduce information at a certain time in a way which depends on the prior information fed to the network. A given input may have a very different representation depending on the previous inputs provided to the network. Encoding sequences rather than single temporal events is highly beneficial in cases in which most of the information follows a specific order (such as natural actions in normal environments). Typical inputs do not occur at a single point in time but rather follow other inputs which are temporally correlated to them.

The second aspect of the partially closed loop and its relationship to the RDDR model is spatial; namely its ability to distribute information and lead to its integration. The limited spatial scope of both the feed-forward and lateral connections and their inherent connectivity sparseness means that only a small portion of the information reaches each neuron (Section 7.5). The feedback information supplied by the thalamus and frontal cortex via the loop structure increases the amount of information available to each neuron within the first layer (striatum), thus enabling improved dimensionality reduction.

The participation of a partially closed loop structure in temporal and spatial information integration may complement the integration performed by the multiple layers of the basal ganglia (Section 7.3). Multilayer integration is crucial in identifying and processing increasingly complex features (Barlow, 1992b). A partially closed loop enables a reduction in the number of layers needed for the encoding of complex features by feeding the partially processed information again through the same layers. This process leads to the formation of highly compressed information using the limited number of layers within the network repetitively.

7.5. Sparse connectivity

The basic RDDR model assumes complete connectivity between and within the different layers. Complete connectivity implies an anatomical connection between each neuron and all the neurons in the preceding and subsequent layers, and between all neurons within the same layer. However, the actual anatomy and physiology of the basal ganglia are far from exhibiting such a complete connectivity pattern. In particular, the connectivity of the different nuclei of the basal ganglia is known to be incomplete, and neurons are not connected to most of the other neurons. This aspect is crucial to understanding the behavior of neural networks and is usually called sparse connectivity.

Sparse connectivity may appear anywhere between two extreme forms. The first type is “ordered sparseness” in which the probability that neighboring neurons will receive similar connections is significantly higher than the chances of remote neurons. The second type is an “unordered sparseness” in which all neurons have the same probability of receiving similar connections. The spatial organization detected anatomically and physiologically suggests that both the striatum (Hoover and Strick, 1993; Kincaid et al., 1998; Zheng and Wilson, 2002) and the globus pallidus (Hoover and Strick, 1993) receive their inputs in a manner much closer to the “ordered sparseness” form. In addition the anatomical evidence shows that the collaterals within the layers are mostly local (Wilson and Groves, 1980; Parent et al., 2000) and therefore the lateral connectivity is also mostly ordered and limited in its spatial scope.

The ordered sparseness of the lateral connections leads to local domains that can become repetitive with global input. The same reduction could be performed independently by different groups of non-connected (or very sparsely connected) neurons leading to repetitions of the encoding. However, the ordered sparseness of the feed-forward connections may lead to the formation of distinct groups of neurons that are highly interconnected and perform a reduction on a subset of the input rather than on the overall input to the nuclei. This, in turn, can lead to the formation of multiple interconnected small funnels instead of a single large funnel. This division into multiple funnels might reduce the overall efficiency of the reduction. However, since most of the redundancy in the input tends to be local, as may be seen by the increased correlation of neighboring cortical neurons relative to remote ones, the reduction may still remain highly effective.

The type of sparseness displayed by the basal ganglia is closely tied to the issue of parallel processing versus convergence of information, a question which has prompted numerous debates in the past (Section 3.1). In ordered sparseness a group of neurons share similar input which is different from the one shared by the other neurons. Thereby they perform their reduction on a segregated channel of information. The connection of multiple groups with each group sharing a subset of the information leads to parallel

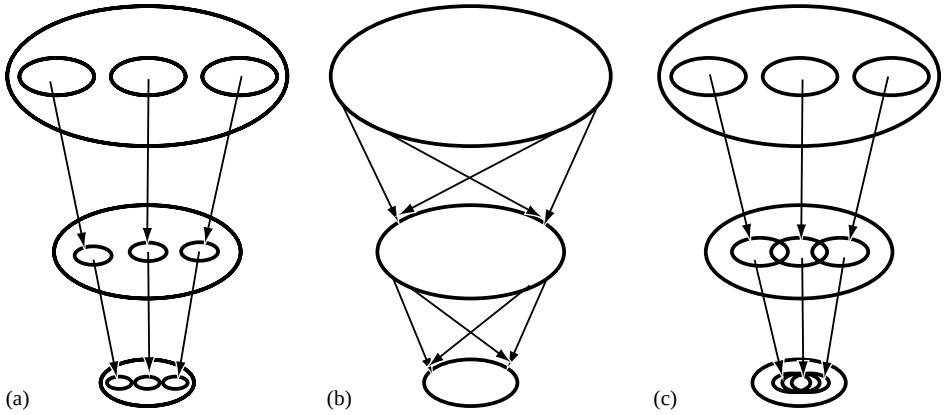

Fig. 6. Patterns of sparse connectivity in the basal ganglia networks. Organization of the feed-forward network in: (a) parallel segregated pathways reflecting ordered sparseness; (b) converging pathways reflecting unordered sparseness; (c) intermediate structure suggested by the known connectivity pattern of the basal ganglia.

processing of the information with no interaction (Fig. 6a). For example, in this extreme case every 10^4 cortical neurons, out of a total of 10^8, will project to a group of 10^3 striatal neurons, thus forming 10^4 separate funnels performing the reduction with complete connectivity within each funnel. At the other extreme lies completely unordered sparseness, in which the whole input structure projects equally to all the neurons in the output structure. Thus, all of the information is shared throughout the layer leading to the convergence of information (Fig. 6b). Using the same example, now all of the 10^8 cortico-striatal neurons extend to the 10^7 striatal neurons, thus creating a single funnel with a very sparse (probability of 10^{-4} for a connection) connectivity pattern. The pattern of connectivity between the layers of the basal ganglia is probably intermediate (Fig. 6c). This means that the sparseness is somewhat ordered, leading to multiple (although maybe not completely segregated) funnels of information, with high levels of connectivity.

Analyzing sparse (or partial) connectivity cannot be complete by looking at a single layer alone. The basal ganglia form a multi-layer (Section 7.3) partially closed loop (Section 7.4). This structure enables the integration of multiple smaller funnels into larger funnels encompassing greater portions of the overall information. Thus, despite the anatomical limits of the connectivity, it is possible to perform dimensionality reduction with various levels of locality and integration. In addition, the local structure of connectivity is well suited to the expected formation of feature extraction (or input reduction) by single sign weights (Section 7.1). In cases of single sign the input tends to be additive and therefore maintains localized aspects rather than the global formations suggested by PCA.

7.6. Multiple pathways

The basic RDDR model considers only a single pathway. This pathway, also called the main axis of the basal ganglia (Percheron et al., 1994), leads from the cortex through the striatum and from there directly to the output nuclei (the GPi and SNr). However, the flow of information through the basal ganglia uses additional pathways: the indirect pathway (Albin et al., 1989) from the striatum through the GPe and the STN to the output nuclei and the hyper-direct pathway (Kita, 1992; Ryan and Clark, 1992; Nambu et al., 2002b) from the cortex directly to the STN and from there to the output nuclei (see Section 2.2 and Fig. 1 for more details). Expanding the RDDR model to include multiple pathways results in several major improvements over the basic network:

- Balance of positive and negative inputs from the same sources.
- Improved integration of multiple inputs.
- Integration of various temporal delay lines.

The direct pathway from the striatum to the globus pallidus exerts a net inhibitory effect on the GPi/SNr. Any elevation in cortical firing increases the striatal firing, thereby decreasing the firing in the output nuclei (Tremblay and Filion, 1989). On the other hand, the two other pathways exert a net positive effect of the cortex on the output nuclei. In the hyper-direct pathway the cortex's increased firing leads to an increase in the STN, which in turn excites the GPi/SNr. In the indirect pathway, the cortex increases the rate of the striatum, which leads to a decrease in the GPe, resulting in an increase in the STN and finally to the excitation of the GPi/SNr (Nambu et al., 2000) (Fig. 5c). Utilizing the multiple pathways, the overall average rate can generally be maintained without shifting to extremely high or low rates (leading to non-linear activation of the output neurons). In addition, the same input may be received in both a net excitatory and a net inhibitory manner via two pathways, leading to unconstrained adaptation to the input. This may serve as a means of overcoming the constraints to single sign input set by the specific neurotransmitters (see Section 7.1).

The different pathways display different levels of integration of cortical information: the striatal direct output to the GPi/SNr presents a significantly more segregated and distinct input than the output of the various pathways going through the STN (Mink, 1996). These different levels of convergence enable the neurons of the output nuclei to access and perform reduction on a combination of information originating from different levels of specificity. The information ranges from specific for each modality (resulting from low convergence) and global information (resulting from high convergence).

Finally, the different pathways have different time delays: the hyper-direct pathway is the fastest, the direct pathway is slower and the information flow through the indirect pathway is the slowest (Nambu et al., 2002b). This enables the output nuclei to receive information from multiple times and to reduce the combined information. This mechanism, combined with the partially closed loop structure of the cortico-basal ganglia loop (Section 7.4) enables the reduction of temporal sequences and not just unitary events in time.

8. Conclusion

New experimental data are continuing to prompt new developments in models of the basal ganglia. Data from many fields, including anatomy, physiology, biochemistry, as well as clinical and even computational evidence are paving the way for the modeling of additional aspects of the function of these nuclei in health and disease. The RDDR model was developed to account for the large numerical reduction in the number of neurons along the cortex-striatum-GPi/SNr axis and the seemingly contradictory anatomical and physiological data regarding lateral connectivity within the basal ganglia nuclei. The model solves the contradiction by suggesting that the basal ganglia perform a compression of the information received from the cortex, based on a reinforcement signal. Thus, the funnel-like structure of the basal ganglia is used to remove the redundancy in the data space spanned by the activity of cortical neurons and to maintain the most important information (from a reward standpoint) to enable efficient planning of new actions.

This review summarizes the background leading to the development of the RDDR model as regards the experimental knowledge acquired in different research fields (Section 2) and earlier models of the cortico-basal ganglia loop (Section 3). This background, together with the theoretical basis of reinforcement learning (Section 4) and dimensionality reduction (Section 5) algorithms suggest that the RDDR model can operate as a natural extension to our assumptions regarding basal ganglia function in health and disease. This experimental and theoretical background also serves as the driving force for the extension of the basic RDDR model (Section 6) to more comprehensive and detailed models (Section 7). The basic RDDR model is a simple neural network representing the general concept of the basal ganglia as a central dimensionality reduction system which is modulated by a reinforcement signal. This simplified view is still able to provide insights into puzzling aspects of basal ganglia anatomy and physiology such as its vast lateral network resulting in no correlations and the seemingly little interaction between the neurons. The different enhancements to the RDDR model increase its scope and are designed to transpose the model from a conceptual one to one which is tightly linked to the biology of the basal ganglia. The various computational aspects of these different advances have shed new light on the possible role of different properties of these nuclei.

The RDDR lays the groundwork for several validation experiments. The first experiment involves the presumed dynamics of the network. During the learning phase the correlations between neurons (Arkadir et al., 2002) and the strength of the lateral weights are expected to increase significantly. Thus, during such a learning period, which might arise either from testing a young animal (Tepper and Trent, 1993) or an animal exposed to a completely new environment, these two parameters should change. Different (or novel) motor and sensory mappings are expected to completely alter the structure of the input space that the animal (and therefore the basal ganglia) experience. The second experiment concerns learning rules. Hebbian and anti-Hebbian learning rules have yet to be documented for this brain area, and so has their modulation by the reinforcement signal.

The key message of this review is the need to explore new experimental data using state-of-the-art analytic and computational techniques. New explorations of the basal ganglia and their interaction with the rest of the brain are crucial as an addition to the gradual evolution of existing concepts. New models such as the RDDR model can serve two crucial needs: as a general concept of the behavior of the basal ganglia in health and disease, and as a tool for the generation of specific experiments and the testing of their expected (or unexpected) results. These new experiments will, in turn, prompt the formulation of improved models, which will either represent an evolution of the RDDR model or a revolution leading to a new, different and better understanding of the basal ganglia.

Acknowledgements

This study was partly supported by a Center of Excellence grant administered by the Israel Science Foundation (ISF), the United States-Israel Binational Science Foundation, the German-Israeli Binational Foundation (GIF) and the BMBF German-Israeli collaboration in medical research. G.M. was supported by the Horowitz fellowship.

References

Abeles, M., 1991. Corticonics—neural circuits of the cerebral cortex. Cambridge University Press, Cambridge.

Akins, P.T., Surmeier, D.J., Kitai, S.T., 1990. Muscarinic modulation of a transient k^+ conductance in rat neostriatal neurons. Nature 344, 240–242.

Albin, R.L., Young, A.B., Penney, J.B., 1989. The functional anatomy of basal ganglia disorders. Trends Neurosci. 12, 366–375.

Alexander, G.E., Crutcher, M.D., 1990. Functional architecture of basal ganglia circuits: neural substrates of parallel processing. Trends Neurosci. 13, 266–271.

Alexander, G.E., DeLong, M.R., 1985. Microstimulation of the primate neostriatum. II. Somatotopic organization of striatal microexcitable zones and their relation to neuronal response properties. J. Neurophysiol. 53, 1417–1430.

Alexander, G.E., DeLong, M.R., Strick, P.L., 1986. Parallel organization of functionally segregated circuits linking basal ganglia and cortex. Ann. Rev. Neurosci. 9, 357–381.

Aosaki, T., Tsubokawa, H., Ishida, A., Watanabe, K., Graybiel, A.M., Kimura, M., 1994. Responses of tonically active neurons in the primate's striatum undergo systematic changes during behavioral sensorimotor conditioning. J. Neurosci. 14, 3969–3984.

Aosaki, T., Kimura, M., Graybiel, A.M., 1995. Temporal and spatial characteristics of tonically active neurons of the primate's striatum. J. Neurophysiol. 73, 1234–1252.

Apicella, P., Scarnati, E., Schultz, W., 1991. Tonically discharging neurons of monkey striatum respond to preparatory and rewarding stimuli. Exp. Brain. Res. 84, 672–675.

Apicella, P., Legallet, E., Trouche, E., 1997. Responses of tonically discharging neurons in the monkey striatum to primary rewards delivered during different behavioral states. Exp. Brain. Res. 116, 456–466.

Arbuthnott, G.W., Wickens, J.R., 1996. Dopamine cells are neurones too!. Trends Neurosci. 19, 279–280.

Arecchi, B.P., Yelnik, J., Francois, C., Percheron, G., Tande, D., 1997. Three-dimensional morphology and distribution of pallidal axons projecting to both the lateral region of the thalamus and the central complex in primates. Brain Res. 754, 311–314.

Arkadir, D., Ben Shaul, Y., Morris, G., Maraton, S., Goldber, J.A., Bergman, H., 2002. False detection of dynamic changes in pallidal neuron interactions by the joint peri-stimulus histogram method. In: Nicholson, L.F.B., Faull, R.L.M. (Eds.), The Basal Ganglia VII, pp. 181–187.

Aubert, I., Ghorayeb, I., Normand, E., Bloch, B., 2000. Phenotypical characterization of the neurons expressing the D1 and D2 dopamine receptors in the monkey striatum. J. Comp. Neurol. 418, 22–32.

Baker, S.N., Philbin, N., Spinks, R., Pinches, E.M., Wolpert, D.M., MacManus, D.G., Pauluis, Q., Lemon, R.N., 1999. Multiple single unit recording in the cortex of monkeys using independently moveable microelectrodes. J. Neurosci. Methods 94, 5–17.

Baldassarre, G., 2002. A modular neural-network model of the basal ganglia's role in learning and selecting motor behaviours. J. Cogn. Syst. Res. 3, 5–13.

Bar-Gad, I., Havazelet-Heimer, G., Goldberg, J.A., Ruppin, E., Bergman, H., 2000. Reinforcement-driven dimensionality reduction—a model for information processing in the basal ganglia. J. Basic Clin. Physiol. Pharmacol. 11, 305–320.

Bar-Gad, I., Heimer, G., Ritov, Y., Bergman, H., 2003. Functional correlations between neighboring neurons in the primate Globus Pallidus are weak or nonexistent. J. Neurosci. 23, 4012–4016.

Barlow, H.B., 1989. Unsupervised learning. Neural Comput. 1, 295–311.

Barlow, H.B., 1992a. Single cells versus neuronal assemblies. In: Aertsen, A., Braitenberg, V.B. (Eds.), Information Processing in the Cortex: Experiments and Theory. Springer-Verlag, pp. 169–174.

Barlow, H.B., 1992b. The biological role of the neocortex. In: Aertsen, A., Braitenberg, V.B. (Eds.), Information Processing in the Cortex: Experiments and Theory. Springer-Verlag, pp. 54–80.

Barto, A.G., 1994. Reinforcement learning control. Curr. Opin. Neurobiol. 4, 888–893.

Barto, A.G., 1995. Adaptive critics and the basal ganglia. In: Houk, J.C., Davis, J.L., Beiser, D.G. (Eds.), Models of Information Processing in the Basal Ganglia. MIT Press, Cambridge, pp. 215–232.

Barto, A.G., Sutton, R.S., Watkins, C.J.C.H., 1989. Learning and sequential decision making. In: Gabriel, M., Moore, J.W. (Eds.), Learning and Computational Neuroscience, MIT Press.

Beiser, D.G., Houk, J.C., 1998. Model of cortical-basal ganglionic processing: encoding the serial order of sensory events. J. Neurophysiol. 79, 3168–3188.

Beiser, D.G., Hua, S.E., Houk, J.C., 1997. Network models of the basal ganglia. Curr. Opin. Neurobiol. 7, 185–190.

Belman, R.E., 1961. Adaptive Control Processes. Princeton University Press.

Benabid, A.L., Pollak, P., Gross, C., Hoffmann, D., Benazzouz, A., Gao, D.M., Laurent, A., Gentil, M., Perret, J., 1994. Acute and long-term effects of subthalamic nucleus stimulation in Parkinson's disease. Stereotact. Funct. Neurosurg. 62, 76–84.

Bennett, B.D., Wilson, C.J., 1999. Spontaneous activity of neostriatal cholinergic interneurons in vitro. J. Neurosci. 19, 5586–5596.

Bennet, B.D., Wilson, C.J., 2000. Synaptology and physiology of neostriatal neurons. In: Miller, R., Wickens, J.R. (Eds.), Brain Dynamics and the Strital Complex. Australia, pp. 111–140.

Bennett, B.D., Callaway, J.C., Wilson, C.J., 2000. Intrinsic membrane properties underlying spontaneous tonic firing in neostriatal cholinergic interneurons. J. Neurosci. 20, 8493–8503.

Bergman, H., Wichmann, T., DeLong, M.R., 1990. Reversal of experimental parkinsonism by lesions of the subthalamic nucleus. Science 249, 1436–1438.

Bergman, H., Wichmann, T., Karmon, B., DeLong, M.R., 1994. The primate subthalamic nucleus. II. Neuronal activity in the MPTP model of parkinsonism. J. Neurophysiol. 72, 507–520.

Bergman, H., Feingold, A., Nini, A., Raz, A., Slovin, H., Abeles, M., Vaadia, E., 1998. Physiological aspects of information processing in the basal ganglia of normal and parkinsonian primates. Trends Neurosci. 21, 32–38.

Bernheimer, H., Birkmayer, W., Hornykiewicz, O., Jellinger, K., Seitelberger, F., 1973. Brain dopamine and the syndromes of Parkinson and Huntington. Clinical, morphological and neurochemical correlations. J. Neurol. Sci. 20, 415–455.

Berns, G.S., Sejnowski, T.J., 1996. How the basal ganglia make decisions. In: Damasio, A., Damasio, H., Christen Y. (Eds.), Neurobiology of Decision Making. Springer-Verlag, Berlin, pp. 101–114

Berns, G.S., Sejnowski, T.J., 1998. A computational model of how the basal ganglia produce sequences. J. Cogn. Neurosci. 10, 108–121.

Bevan, M.D., Booth, P.A., Eaton, S.A., Bolam, J.P., 1998. Selective innervation of neostriatal interneurons by a subclass of neuron in the globus pallidus of the rat. J. Neurosci. 18, 9438–9452.

Bevan, M.D., Magill, P.J., Terman, D., Bolam, J.P., Wilson, C.J., 2002. Move to the rhythm: oscillations in the subthalamic nucleus-external globus pallidus network. Trends Neurosci. 25, 525–531.

Bolam, J.P., Powell, J.F., Wu, J.Y., Smith, A.D., 1985. Glutamate decarboxylase-immunoreactive structures in the rat neostriatum: a correlated light and electron microscopic study including a combination of Golgi impregnation with immunocytochemistry. J. Comp. Neurol. 237, 1–20.

Bolam, J.P., Smith, Y., Ingham, C.A., von Krosigk, M., Smith, A.D., 1993. Convergence of synaptic terminals from the striatum and the globus pallidus onto single neurones in the substantia nigra and the entopeduncular nucleus. Prog. Brain Res. 99, 73–88.

Bolam, J.P., Hanley, J.J., Booth, P.-A.C., Bevan, M.D., 2000. Synaptic organisation of the basal ganglia. J. Anat. 196, 527–542.

Boraud, T., Bezard, E., Guehl, D., Bioulac, B., Gross, C., 1998. Effects of L-DOPA on neuronal activity of the globus pallidus externalis (GPe) and globus pallidus internalis (GPi) in the MPTP-treated monkey. Brain Res 787, 157–160.

Boraud, T., Bezard, E., Bioulac, B., Gross, C.E., 2002. From single extracellular unit recording in experimental and human Parkinsonism

to the development of a functional concept of the role played by the basal ganglia in motor control. Prog. Neurobiol. 66, 265–283.
Bourlard, H., Kamp, Y., 1988. Auto-association by multilayer perceptrons and singular value decomposition. Biol. Cybern. 59, 291–294.
Bouyer, J.J., Park, D.H., Joh, T.H., Pickel, V.M., 1984. Chemical and structural analysis of the relation between cortical inputs and tyrosine hydroxylase-containing terminals in rat neostriatum. Brain Res. 302, 267–275.
Braitenberg, V.B., Schuz, A., 1991. Anatomy of the Cortex: Statistics and Geometry. Springer-Verlag.
Brissaud, E., 1895. Pathogenie et symptomes de la maladie de Parkinson. In: Leçons sur les maladies nerveuses. Paris, pp. 469–487.
Burns, R.S., Chiueh, C.C., Markey, S.P., Ebert, M.H., Jacobowitz, D.M., Kopin, I.J., 1983. A primate model of parkinsonism: selective destruction of dopaminergic neurons in the pars compacta of the substantia nigra by *N*-methyl-4-phenyl-1,2,3,6-tetrahydropyridine. Proc. Natl. Acad. Sci. U.S.A. 80, 4546–4550.
Calabresi, P., Pisani, A., Centonze, D., Bernardi, G., 1997. Role of dopamine receptors in the short- and long-term regulation of corticostriatal transmission. Nihon Shinkei Seishin Yakurigaku Zasshi 17, 101–104.
Calabresi, P., Centonze, D., Gubellini, P., Marfia, G.A., Pisani, A., Sancesario, G., Bernardi, G., 2000. Synaptic transmission in the striatum: from plasticity to neurodegeneration. Prog. Neurobiol. 61, 231–265.
Carlson, A., 1990. Anti-Hebbian learning in a non-linear neural network. Biol. Cybern. 64, 171–176.
Carpenter, M.B., Strominger, N.L., 1967. Efferent fibers of the subthalamic nucleus in the monkey. A comparison of the efferent projections of the subthalamic nucleus, substantia nigra and globus pallidus. Am. J. Anat. 121, 41–72.
Carpenter, M.B., Whittier, J.R., Mettler, F.A., 1950. Analysis of choreoid hyperkinesia in the rhesus monkey: surgical and pharmacological analysis of hyperkinesia resulting from lesions in the subthalamic nucleus of Luys. J. Comp. Neurol. 92, 293–332.
Celada, P., Paladini, C.A., Tepper, J.M., 1999. GABAergic control of rat substantia nigra dopaminergic neurons: role of globus pallidus and substantia nigra pars reticulata. Neuroscience 89, 813–825.
Centonze, D., Gubellini, P., Picconi, B., Calabresi, P., Giacomini, P., Bernardi, G., 1999. Unilateral dopamine denervation blocks corticostriatal LTP. J. Neurophysiol. 82, 3575–3579.
Cepeda, C., Levine, M.S., 1998. Dopamine and *N*-methyl-D-aspartate receptor interactions in the neostriatum. Dev. Neurosci. 20, 1–18.
Chakrabarti, K., Mehrotra, S., 2000. Local Dimensionality Reduction: A New Approach to Indexing High Dimensional Spaces. VLDB, pp. 89–100.
Chang, H.T., Kita, H., Kitai, S.T., 1983. The fine structure of the rat subthalamic nucleus: an electron microscopic study. J. Comp. Neurol. 221, 113–123.
Chavas, J., Marty, A., 2003. Coexistence of excitatory and inhibitory GABA synapses in the cerebellar interneuron network. J. Neurosci. 23, 2019–2031.
Chesselet, M.F., Delfs, J.M., 1996. Basal ganglia and movement disorders: an update. Trends Neurosci. 19, 417–422.
Churchland, P.S., Sejnowski, T.J., 1992. The Computational Brain. MIT Press, Cambridge, MA.
Cooper, A.J., Stanford, I.M., 2000. Electrophysiological and morphological characteristics of three subtypes of rat globus pallidus neurone in vitro. J. Physiol. 527 (Part 2), 291–304.
Crutcher, M.D., DeLong, M.R., 1984a. Single cell studies of the primate putamen. I. Functional organization. Exp. Brain. Res. 53, 233–243.
Crutcher, M.D., DeLong, M.R., 1984b. Single cell studies of the primate putamen. II. Relations to direction of movement and pattern of muscular activity. Exp. Brain. Res. 53, 244–258.
Czubayko, U., Plenz, D., 2002. Fast synaptic transmission between striatal spiny projection neurons. Proc. Natl. Acad. Sci. U.S.A. 99, 15764–15769.
DeLong, M.R., 1972. Activity of basal ganglia neurons during movement. Brain Res. 40, 127–135.
DeLong, M.R., 1990. Primate models of movement disorders of basal ganglia origin. Trends Neurosci. 13, 281–285.
DeLong, M.R., Georgopoulos, A.P., 1981. Motor functions of the basal ganglia. In: Brookhart, J.M., Mountcastle, V.B., Brooks, V.B., Geiger, S.R. (Eds.), Handbook of Physiology. The Nervous System. Motor Control, Section 1, vol. II, Part 2. American Physiological Society, Bethesda, pp. 1017–1061.
DeLong, M.R., Crutcher, M.D., Georgopoulos, A.P., 1985. Primate globus pallidus and subthalamic nucleus: functional organization. J. Neurophysiol. 53, 530–543.
DeMers, D., Cottrell, G., 1993. Nonlinear dimensionality reduction. Neural Inf. Process. Syst. 5, 580–587.
Denny-Brown, D., 1962. The Basal Ganglia and their Relation to Disorders of Movement. Oxford University Press, London.
Diamantaras, K.I., Kung, S.Y., 1996. Principal Component Neural Networks: Theory and Applications. Wiley, New York.
Di Chiara, G., Morelli, M., Consolo, S., 1994. Modulatory functions of neurotransmitters in the striatum: ACh/dopamine/NMDA interactions. Trends Neurosci. 17, 228–233.
Difiglia, M., Rafols, J.A., 1988. Synaptic organization of the globus pallidus. J. Elctron. Microsc. Tech. 10, 247–263.
Difiglia, M., Pasik, P., Pasik, T., 1976. A Golgi study of neuronal types in the neostriatum of monkeys. Brain Res. 114, 245–256.
Difiglia, M., Pasik, P., Pasik, T., 1982. A golgi and ultrastructural study of the monkey globus pallidus. J. Comp. Neurol. 212, 53–75.
Dominey, P.F., 1995. Complex sensory-motor sequence learning based on recurrent state representation and reinforcement learning. Biol. Cybern. 73, 265–274.
Dube, L., Smith, A.D., Bolam, J.P., 1988. Identification of synaptic terminals of thalamic or cortical origin in contact with distinct medium-size spiny neurons in the rat neostriatum. J. Comp. Neurol. 267, 455–471.
Eggermont, J.J., 1990. The Correlative Brain. Theory and Experiment in Neuronal Interaction. Springer-Verlag, Berlin.
Ehringer, H., Hornykiewicz, O., 1960. Verteilung von noradrenalin und dopamin (3-hydroxytyramin) im gehirn des menschen und ihr verhalten bei erkrankungen des extrapyramidalen systems. Klin Wschr 38, 1236–1239.
Fibiger, H.C., Phillips, A.G., 1988. Mesocorticolimbic dopamine systems and reward. Ann. N. Y. Acad. Sci. 537, 206–215.
Filion, M., Tremblay, L., 1991. Abnormal spontaneous activity of globus pallidus neurons in monkeys with MPTP-induced parkinsonism. Brain Res. 547, 142–151.
Fiorillo, C.D., Tobler, P.N., Schultz, W., 2003. Discrete coding of reward probability and uncertainty by dopamine neurons. Science 299, 1898–1902.
Flaherty, A.W., Graybiel, A.M., 1991. Corticostriatal transformations in the primate somatosensory system. Projections from physiologically mapped body-part representations. J. Neurophysiol. 66, 1249–1263.
Flaherty, A.W., Graybiel, A.M., 1993. Output architecture of the primate putamen. J. Neurosci. 13, 3222–3237.
Foldiak, P., 1989. Adaptive network for optimal feature extraction. Proc. Int. Joint Conf. Neural Netw. 1, 401–405.
Foldiak, P., 1990. Forming sparse representations by local anti-Hebbian learning. Biol. Cybern. 64, 165–170.
Fox, C.A., Andrade, A.N., Lu Qui, I.J., Rafols, J.A., 1974. The primate globus pallidus: a golgi and electron microscopic study. J. Hirnforsch. 15, 75–93.
Francois, C., Percheron, G., Yelnik, J., Heyner, S., 1984. A Golgi analysis of the primate globus pallidus. I. Inconstant processes of large neurons, other neuronal types, and afferent axons. J. Comp. Neurol. 227, 182–199.
Francois, C., Yelnik, J., Percheron, G., 1987. Golgi study of the primate substantia nigra. II. Spatial organization of dendritic arborizations in relation to the cytoarchitectonic boundaries and to the striatonigral bundle. J. Comp. Neurol. 265, 473–493.

Francois, C., Tande, D., Yelnik, J., Hirsch, E.C., 2002. Distribution and morphology of nigral axons projecting to the thalamus in primates. J. Comp. Neurol. 447, 249–260.

Freund, T.F., Powell, J.F., Smith, A.D., 1984. Tyrosine hydroxylase-immunoreactive boutons in synaptic contact with identified striatonigral neurons, with particular reference to dendritic spines. Neuroscience 13, 1189–1215.

Friston, K.J., Frith, C.D., Frackowiak, R.S., 1993. Principal component analysis learning algorithms: a neurobiological analysis. Proc. R. Soc. Lond. B Biol. Sci. 254, 47–54.

Fukai, T., 1999. Sequence generation in arbitrary temporal patterns from theta-nested gamma oscillations: a model of the basal ganglia-thalamo-cortical loops. Neural Netw. 12, 975–987.

Fukai, T., Tanaka, S., 1997. A simple neural network exhibiting selective activation of neuronal ensembles: from winner-take-all to winners-share-all. Neural Comput. 9, 77–97.

Georgopoulos, A.P., Schwartz, A.B., Kettner, R.E., 1986. Neuronal population coding of movement direction. Science 233, 1416–1419.

Gerbrands, J.J., 1981. On the relationships between SVD, KLT, and PCA. Pattern Recogn. 14, 375–381.

Gerfen, C.R., 1985. The neostriatal mosaic. I. Compartmental organization of projections from the striatum to the substantia nigra in the rat. J. Comp. Neurol. 236, 454–476.

Gerfen, C.R., 1988. Synaptic organization of the striatum. J. Electron Microsc. Tech. 10, 265–281.

Gerfen, C.R., 1989. The neostriatal mosaic: striatal patch-matrix organization is related to cortical lamination. Science 246, 385–388.

Gerfen, C.R., 1992. The neostriatal mosaic: multiple levels of compartmental organization. Trends Neurosci. 15, 133–139.

Gerfen, C.R., Wilson, C.J., 1996. The basal ganglia. In: Swanson, L.W., Bjorklund A., Hokfelt, T. (Eds.), Handbook of Chemical Neuroanatomy. Integrated Systems of the CNS, vol. 12, part III, Elsevier Science, pp. 371–468.

Gerfen, C.R., Engber, T.M., Mahan, L.C., Susel, Z., Chase, T.N., Monsma Jr., F.J., Sibley, D.R., 1990. D1 and D2 dopamine receptor-regulated gene expression of striatonigral and striatopallidal neurons. Science 250, 1429–1432.

Gillies, A., Arbuthnott, G., 2000. Computational models of the basal ganglia. Mov. Disord. 15, 762–770.

Goldberg, J.A., Boraud, T., Maraton, S., Haber, S.N., Vaadia, E., Bergman, H., 2002. Enhanced synchrony among primary motor cortex neurons in the 1-methyl-4-phenyl-1,2,3,6-tetrahydropyridine primate model of Parkinson's disease. J. Neurosci. 22, 4639–4653.

Goldman, P.S., Nauta, W.J.H., 1977. An intricately patterned prefronto-caudate projection in the rhesus monkey. J. Comp. Neurol. 171, 369–386.

Grace, A.A., Bunney, B.S., 1984a. The control of firing pattern in nigral dopamine neurons: burst firing. J. Neurosci. 4, 2877–2890.

Grace, A.A., Bunney, B.S., 1984b. The control of firing pattern in nigral dopamine neurons: single spike firing. J. Neurosci. 4, 2866–2876.

Graybiel, A.M., Ragsdale, C.W., 1978. Histochemically distinct compartments in the striatum of human being, monkey, and cat demonstrated by the acetylcholinesterase staining method. Proc. Natl. Acad. Sci. U.S.A. 75, 5723–5726.

Graybiel, A.M., Ragsdale, C.-W.J., 1980. Clumping of acetylcholinesterase activity in the developing striatum of the human fetus and young infant. Proc. Natl. Acad. Sci. U.S.A. 77, 1214–1218.

Graybiel, A.M., Aosaki, T., Flaherty, A.W., Kimura, M., 1994. The basal ganglia and adaptive motor control. Science 265, 1826–1831.

Gurney, K., Prescott, T.J., Redgrave, P., 2001. A computational model of action selection in the basal ganglia. I. A new functional anatomy. Biol. Cybern. 84, 401–410.

Haber, S., Elde, R., 1981. Correlation between met-enkephalin and substance P immunoreactivity in the primate globus pallidus. Neuroscience 6, 1291–1297.

Haber, S.N., McFarland, N.R., 2001. The place of the thalamus in frontal cortical-basal ganglia circuits. Neuroscientist 7, 315–324.

Haber, S.N., Gdowski, M.J., 2003. The basal ganglia. In: Paxinos, G., Mai, J. (Eds.), The Human Nervous System. Academic Press.

Haber, S.N., Fudge, J.L., McFarland, N.R., 2000. Striatonigrostriatal pathways in primates form an ascending spiral from the shell to the dorsolateral striatum. J. Neurosci. 20, 2369–2382.

Hamada, I., DeLong, M.R., Mano, N., 1990. Activity of identified wrist-related pallidal neurons during step and ramp wrist movements in the monkey. J. Neurophysiol. 64, 1892–1906.

Handel, A., Glimcher, P.W., 1999. Quantitative analysis of substantia nigra pars reticulata activity during a visually guided saccade task. J. Neurophysiol. 82, 3458–3475.

Hanley, J.J., Bolam, J.P., 1997. Synaptology of the nigrostriatal projection in relation to the compartmental organization of the neostriatum in the rat. Neuroscience 81, 353–370.

Hassler, R., 1939. Zur pathologischen anatomie des senilen und des parkinsonistischen tremor. Journal fur psychologie und neurologie, pp. 13–15.

Haykin, S., 1999. Neural networks—a comprehensive foundation. Prentice Hall, Upper Saddle River, NJ.

Hazrati, L.N., Parent, A., 1992. The striatopallidal projection displays a high degree of anatomical specificity in the primate. Brain Res. 592, 213–227.

Hazrati, L.N., Parent, A., Mitchell, S., Haber, S.N., 1990. Evidence for interconnections between the two segments of the globus pallidus in primates: a PHA-L anterograde tracing study. Brain Res. 533, 171–175.

Hebb, D.O., 1949. The Organization of Behavior. Wiley, New York.

Heimer, L., Switzer, R.D., Van Hoesen, G.W., 1982. Ventral striatum and ventral pallidum components of the motor system? Trends Neurosci. 5, 83–87.

Heimer, G., Bar-Gad, I., Goldberg, J.A., Bergman, H., 2002. Dopamine replacement therapy reverses abnormal synchronization of pallidal neurons in the 1-methyl-4-phenyl-1,2,3,6-tetrahydropyridine primate model of parkinsonism. J. Neurosci. 22, 7850–7855.

Hertz, J.A., Krogh, A.S., Palmer, R.G., 1994. Introduction to the Theory of Neural Computation. Addison-Wesley, Redwood City, CA.

Hikosaka, O., Wurtz, R.H., 1983. Visual and oculomotor functions of monkey substantia nigra pars reticulata. I. Relation of visual and auditory responses to saccades. J. Neurophysiol. 49, 1230–1253.

Hikosaka, O., Sakamoto, M., Usui, S., 1989. Functional properties of monkey caudate neurons. I. Activities related to saccadic eye movements. J. Neurophysiol. 61, 780–798.

Hikosaka, O., Matsumura, M., Kojima, J., Gardiner, T.W., 1993. Role of basal ganglia in initiation and suppression of saccadic eye movements. In: Mano N., Hamada, I. DeLong, M.R. (Eds.), Role of the Cerebellum and Basal Ganglia in Voluntary Movement. Excerpta Medica, Amsterdam, pp. 213–219.

Hollerman, J.R., Schultz, W., 1998. Dopamine neurons report an error in the temporal prediction of reward during learning. Nat. Neurosci. 1, 304–309.

Hoover, J.E., Strick, P.L., 1993. Multiple output channels in the basal ganglia. Science 259, 819–821.

Hopfield, J.J., 1982. Neural networks and physical systems with emergent collective computational abilities. Proc. Natl. Acad. Sci. U.S.A. 79, 2554–2558.

Hornykiewicz, O., Kish, S.J., 1987. Biochemical pathophysiology of Parkinson's disease. Adv. Neurol. 45, 19–34.

Horvitz, J.C., 2000. Mesolimbocortical and nigrostriatal dopamine responses to salient non-reward events. Neuroscience 96, 651–656.

Houk, J.C., Adams, J. L., Barto, A.G., 1995. A model of how the basal ganglia generate and use neural signals that predict reinforcement. In: Houk, J.C., Davis, J.L., Beiser, D.G. (Eds.), Models of Information Processing in the Basal Ganglia. MIT Press, pp. 249–270.

Humphries, M.D., Gurney, K.N., 2002. The role of intra-thalamic and thalamocortical circuits in action selection. Network 13, 131–156.

Hurtado, J.M., Gray, C.M., Tamas, L.B., Sigvardt, K.A., 1999. Dynamics of tremor-related oscillations in the human globus pallidus: a single case study. Proc. Natl. Acad. Sci. U.S.A. 96, 1674–1679.

Ingham, C.A., Hood, S.H., van-Maldegem, B., Weenink, A., Arbuthnott, G.W., 1993. Morphological changes in the rat neostriatum after unilateral 6-hydroxydopamine injections into the nigrostriatal pathway. Exp. Brain. Res. 93, 17–27.

Iwahori, N., Mizuno, N., 1981. A Golgi study on the globus pallidus of the mouse. J. Comp. Neurol. 197, 29–43.

Jaakkola, T., Jordan, M.I., Singh, S.P., 1994. On the convergence of stochastic iterative dynamic programing algorithms. Neural Comp. 6, 1185–1201.

Jaeger, D., Kita, H., Wilson, C.J., 1994. Surround inhibition among projection neurons is weak or nonexistent in the rat neostriatum. J. Neurophysiol. 72, 2555–2558.

Jaeger, D., Gilman, S., Aldridge, J.W., 1995. Neuronal activity in the striatum and pallidum of primates related to the execution of externally cued reaching movements. Brain Res. 694, 111–127.

Jaffe, E.H., Marty, A., Schulte, A., Chow, R.H., 1998. Extrasynaptic vesicular transmitter release from the somata of substantia nigra neurons in rat midbrain slices. J. Neurosci. 18, 3548–3553.

Jiang, H., Stein, B.E., McHaffie, J.G., 2003. Opposing basal ganglia processes shape midbrain visuomotor activity bilaterally. Nature 424, 982–986.

Joel, D., Weiner, I., 1994. The organization of the basal ganglia-thalamocortical circuits: open interconnected rather than closed segregated. Neuroscience 63, 363–379.

Jutten, C., Herault, J., 1991. Blind separation of sources. Part I. An adaptive algorithm based on neuromimitic architecture. Signal Process. 24, 1–10.

Karhunen, J., Joutsensalo, J., 1994. Representation and separation of signals using nonlinear PCA type learning. Neural Netw. 7, 113–127.

Katayama, J., Akaike, N., Nabekura, J., 2003. Characterization of pre- and post-synaptic metabotropic glutamate receptor-mediated inhibitory responses in substantia nigra dopamine neurons. Neurosci. Res. 45, 101–115.

Kawaguchi, Y., Wilson, C.J., Augood, S.J., Emson, P.C., 1995. Striatal interneurons: chemical, physiological and morphological characterization. Trends Neurosci. 18, 527–535.

Kemp, J.M., Powell, T.P.S., 1970. The cortico-striate projection in the monkey. Brain 93, 525–546.

Kemp, J.M., Powell, T.P., 1971. The termination of fibres from the cerebral cortex and thalamus upon dendritic spines in the caudate nucleus: a study with the Golgi method. Philos. Trans. R. Soc. Lond. B Biol. Sci. 262, 429–439.

Kermadi, I., Joseph, J.P., 1995. Activity in the caudate nucleus of monkey during spatial sequencing. J. Neurophysiol. 74, 911–933.

Kerr, J.N., Wickens, J.R., 2001. Dopamine D-1/D-5 receptor activation is required for long-term potentiation in the rat neostriatum in vitro. J. Neurophysiol. 85, 117–124.

Kim, R., Nakano, K., Jayaraman, A., Carpenter, M.B., 1976. Projections of the globus pallidus and adjacent structures: an autoradiographic study in the monkey. J. Comp. Neurol. 169, 263–290.

Kimura, M., Rajkowski, J., Evarts, E., 1984. Tonically discharging putamen neurons exhibit set-dependent responses. Proc. Natl. Acad. Sci. U.S.A. 81, 4998–5001.

Kimura, M., Kato, M., Shimazaki, H., Watanabe, K., Matsumoto, N., 1996. Neural information transferred from the putamen to the globus pallidus during learned movement in the monkey. J. Neurophysiol. 76, 3771–3786.

Kimura, M., Matsumoto, N., Okahashi, K., Ueda, Y., Satoh, T., Minamimoto, T., Sakamoto, M., Yamada, H., 2003. Goal-directed, serial and synchronous activation of neurons in the primate striatum. Neuroreport 14, 799–802.

Kincaid, A.E., Zheng, T., Wilson, C.J., 1998. Connectivity and convergence of single corticostriatal axons. J. Neurosci. 18, 4722–4731.

Kita, H., 1992. Responses of globus pallidus neurons to cortical stimulation: intracellular study in the rat. Brain Res. 589, 84–90.

Kita, H., 1993. GABAergic circuits of the striatum. Prog. Brain Res. 99, 51–72.

Kita, H., Kitai, S.T., 1991. Intracellular study of rat globus pallidus neurons: membrane properties and responses to neostriatal, subthalamic and nigral stimulation. Brain Res. 564, 296–305.

Kita, H., Kitai, S.T., 1994. The morphology of globus pallidus projection neurons in the rat: an intracellular staining study. Brain Res. 636, 308–319.

Kita, H., Tokuno, H., Nambu, A., 1999. Monkey globus pallidus external segment neurons projecting to the neostriatum. Neuroreport 10, 1467–1472.

Kitai, S.T., 1981. Electrophysiology of the corpus striatum and brain stem integrating systems. In: Brookhart, J.M., Mountcastle, V.B., Brooks, V.B., Geiger, S.R. (Eds.), Handbook of Physiology: The Nervous System. Motor Control, Section 1, vol. II, Part 2. American Physiological Society, Bethesda, MD, pp. 997–1015.

Kitai, S.T., Kita, H., 1986. Anatomy and physiology of the subthalamic nucleus: a driving force of the basal ganglia. In: Carpenter M.B., Jayaraman, A. (Eds.), The Basal Ganglia, vol. II. Plenum Press, New York, pp. 357–373.

Kitano, K., Cateau, H., Kaneda, K., Nambu, A., Takada, M., Fukai, T., 2002. Two-state membrane potential transitions of striatal spiny neurons as evidenced by numerical simulations and electrophysiological recordings in awake monkeys. J. Neurosci. 22, RC230.

Kohonen, T., 1995. Self Organizing Maps. Springer, Berlin.

Kolomiets, B.P., Deniau, J.M., Mailly, P., Menetrey, A., Glowinski, J., Thierry, A.M., 2001. Segregation and convergence of information flow through the cortico-subthalamic pathways. J. Neurosci. 21, 5764–5772.

Koos, T., Tepper, J.M., 1999. Inhibitory control of neostriatal projection neurons by GABAergic interneurons. Nat. Neurosci. 2, 467–472.

Kung, S.Y., Diamantaras, K.I., 1990. A neural network learning algorithm for adaptive principal component extraction (APEX). Proc. IEEE Int. Conf. Acoustics Speech Signal Process. 2, 861–864.

Kunzle, H., 1975. Bilateral projections from precentral motor cortex to the putamen and other parts of the basal ganglia. An autoradiographic study in Macaca fascicularis. Brain Res. 88, 195–209.

Kuo, J.S., Carpenter, M.B., 1973. Organization of pallidothalamic projections in the rhesus monkey. J. Comp. Neurol. 151, 201–236.

Lapper, S.R., Smith, Y., Sadikot, A.F., Parent, A., Bolam, J.P., 1992. Cortical input to parvalbumin-immunoreactive neurones in the putamen of the squirrel monkey. Brain Res. 580, 215–224.

Lavoie, B., Parent, A., 1990. Immunohistochemical study of the serotoninergic innervation of the basal ganglia in the squirrel monkey. J. Comp. Neurol. 299, 1–16.

Lee, D.D., Seung, H.S., 1999. Learning the parts of objects by non-negative matrix factorization. Nature 401, 788–791.

Lee, I.H., Assad, J.A., 2003. Putaminal activity for simple reactions or self-timed movements. J. Neurophysiol. 89, 2528–2537.

Lemstra, A.W., Verhagen, M.L., Lee, J.I., Dougherty, P.M., Lenz, F.A., 1999. Tremor-frequency (3–6 Hz). activity in the sensorimotor arm representation of the internal segment of the globus pallidus in patients with Parkinson's disease. Neurosci. Lett. 267, 129–132.

Levy, R., Dostrovsky, J.O., Lang, A.E., Sime, E., Hutchison, W.D., Lozano, A.M., 2001. Effects of apomorphine on subthalamic nucleus and globus pallidus internus neurons in patients with Parkinson's disease. J. Neurophysiol. 86, 249–260.

Levy, R., Hutchison, W.D., Lozano, A.M., Dostrovsky, J.O., 2002. Synchronized neuronal discharge in the basal ganglia of parkinsonian patients is limited to oscillatory activity. J. Neurosci. 22, 2855–2861.

Lewicki, M.S., 1998. A review of methods for spike sorting: the detection and classification of neural action potentials. Network 9, R53–R78.

Liles, S.L., 1985. Activity of neurons in putamen during active and passive movements of wrist. J. Neurophysiol. 53, 217–236.

Linsker, R., 1988. Self-organization in a perceptual network. IEEE Comput. 21, 105–117.

Lippmann, R.P., 1987. An introduction to computing with neural nets, IEEE ASSP Mag. 4–22.

Malach, R., Graybiel, A.M., 1986. Mosaic architecture of the somatic sensory-recipient sector of the cat's striatum. J. Neurosci. 6, 3436–3458.

Malthouse, E.C., Mah, R.H. S., Tamhane, A.C., 1995. Some theoretical results on nonlinear principal component analysis, Proc. Am. Control Conf. 744–748.

Matsumura, M., Kojima, J., Gardiner, T.W., Hikosaka, O., 1992. Visual and occulomotor functions of the monkey subthalamic nucelus. J. Neurophysiol. 67, 1615–1632.

McFarland, N.R., Haber, S.N., 2000. Convergent inputs from thalamic motor nuclei and frontal cortical areas to the dorsal striatum in the primate. J. Neurosci. 20, 3798–3813.

McFarland, N.R., Haber, S.N., 2001. Organization of thalamostriatal terminals from the ventral motor nuclei in the macaque. J. Comp. Neurol. 429, 321–336.

McFarland, N.R., Haber, S.N., 2002. Thalamic relay nuclei of the basal ganglia form both reciprocal and nonreciprocal cortical connections, linking multiple frontal cortical areas. J. Neurosci. 22, 8117–8132.

Merchand, R., Lajoie, L., Blanchet, C., 1986. Histogenesis at the level of the basal forebrain: the entopeduncular nucleus. Neuroscience 17, 591–607.

Merello, M., Balej, J., Delfino, M., Cammarota, A., Betti, O., Leiguarda, R., 1999. Apomorphine induces changes in GPi spontaneous outflow in patients with Parkinson's disease. Mov. Disord. 14, 45–49.

Middleton, F.A., Strick, P.L., 2000. Basal ganglia and cerebellar loops: motor and cognitive circuits. Brain Res. Brain Res. Rev. 31, 236–250.

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

Miller, W.C., DeLong, M. R., 1987. Altered tonic activity of neurons in the globus pallidus and subthalamic nucleus in the primate MPTP model of parkinsonism. In: Carpenter, M.B., Jayaraman, A. (Eds.), The Basal Ganglia, vol. II. Plenum Press, New York, pp. 415–427.

Mink, J.W., 1996. The basal ganglia: focused selection and inhibition of competing motor programs. Prog. Neurobiol. 50, 381–425.

Monakow, K.H., Akert, K., Kunzle, H., 1978. Projections of the precentral motor cortex and other cortical areas of the frontal lobe to the subthalamic nucleus in the monkey. Exp. Brain Res. 33, 395–403.

Nakahara, H., Doya, K., Hikosaka, O., 2001. Parallel cortico-basal ganglia mechanisms for acquisition and execution of visuomotor sequences—a computational approach. J. Cogn. Neurosci. 13, 626–647.

Nakanishi, H., Kita, H., Kitai, S.T., 1990. Intracellular study of rat entopeduncular nucleus neurons in an vitro slice preparations: elctrical membrane properties. Brain Res. 527, 81–88.

Nakanishi, H., Kita, H., Kitai, S.T., 1991. Intracellular study of the rat entopeduncular nucleus neurons in an in vitro slice preparation: response to subthalamic stimulation. Brain Res. 549, 285–291.

Nambu, A., Llinas, R., 1997. Morphology of globus pallidus neurons: its correlation with electrophysiology in guinea pig brain slices. J. Comp. Neurol. 377, 85–94.

Nambu, A., Takada, M., Inase, M., Tokuno, H., 1996. Dual somatotopical representations in the primate subthalamic nucleus: evidence for ordered but reversed body-map transformations from the primary motor cortex and the supplementary motor area. J. Neurosci. 16, 2671–2683.

Nambu, A., Tokuno, H., Hamada, I., Kita, H., Imanishi, M., Akazawa, T., Ikeuchi, Y., Hasegawa, N., 2000. Excitatory cortical inputs to pallidal neurons via the subthalamic nucleus in the monkey. J. Neurophysiol. 84, 289–300.

Nambu, A., Kaneda, K., Tokuno, H., Takada, M., 2002a. Organization of corticostriatal motor inputs in monkey putamen. J. Neurophysiol. 88, 1830–1842.

Nambu, A., Tokuno, H., Takada, M., 2002b. Functional significance of the cortico-subthalamo-pallidal 'hyperdirect' pathway. Neurosci. Res. 43, 111–117.

Nicola, S.M., Surmeier, J., Malenka, R.C., 2000. Dopaminergic modulation of neuronal excitability in the striatum and nucleus accumbens. Annu. Rev. Neurosci. 23, 185–215.

Nini, A., Feingold, A., Slovin, H., Bergman, H., 1995. Neurons in the globus pallidus do not show correlated activity in the normal monkey, but phase-locked oscillations appear in the MPTP model of parkinsonism. J. Neurophysiol. 74, 1800–1805.

Nisenbaum, E.S., Wilson, C.J., 1995. Potassium currents responsible for inward and outward rectification in rat neostriatal spiny projection neurons. J. Neurosci. 15, 4449–4463.

Nutt, J.G., Obeso, J.A., Stocchi, F., 2000. Continuous dopamine-receptor stimulation in advanced Parkinson's disease. Trends Neurosci. 23, 109–115.

Oertel, W.H., Mugnaini, E., 1984. Immunocytochemical studies of GABAergic neurons in rat basal ganglia and their relations to other neuronal systems. Neurosci. Lett. 47, 233–238.

Oja, E., 1982. A simplified neuron model as a principal component analyzer. J. Math. Biol. 15, 267–273.

Oja, E., 1991. Data compression, feature extraction, and autoassociation in feedforward neural networks. In: Kohonen, T., Makisara, K., Simula O., Kangas J. (Eds.), Artificial Neural Networks, pp. 737–745.

Oja, E., Karhunen, J., 1985. On stochastic approximation of the eigenvectors and eigenvalues of the expectation of a random matrix. J. Math. Anal. Appl. 106, 69–84.

Oja, E., Ogawa, H., Wangviwattana, J., 1991. Learning in nonlinear constrained Hebbian networks. In: Kohonen, T., Makisara, K., Simula O., Kangas J. (Eds.), Artificial Neural Networks, pp. 381–390 (Ref Type: Conference Proceeding).

Oja, E., Karhunen J., Wang L., Vigario R., 1995. Principle and Independent Components in Neural Networks—Recent Developments. Italian Workshop on Neural Networks.

Oorschot, D.E., 1996. Total number of neurons in the neostriatal, pallidal, subthalamic, and substantia nigral nuclei of the rat basal ganglia: a stereological study using the cavalieri and optical disector methods. J. Comp. Neurol. 366, 580–599.

Orieux, G., Francois, C., Feger, J., Yelnik, J., Vila, M., Ruberg, M., Agid, Y., Hirsch, E.C., 2000. Metabolic activity of excitatory parafascicular and pedunculopontine inputs to the subthalamic nucleus in a rat model of Parkinson's disease. Neuroscience 97, 79–88.

Parent, A., 1990. Extrinsic connections of the basal ganglia. Trends Neurosci. 13, 254–258.

Parent, A., Hazrati, L.N., 1993. Anatomical aspects of information processing in primate basal ganglia. Trends Neurosci. 16, 111–116.

Parent, A., Hazrati, L.N., 1995. Functional anatomy of the basal ganglia. I. The cortico-basal ganglia-thalamo-cortical loop. Brain Res. Rev. 20, 91–127.

Parent, A., Cicchetti, F., 1998. The current model of basal ganglia organization under scrutiny. Mov. Disord. 13, 199–202.

Parent, A., Mackey, A., De Bellefeuille, L., 1983. The subcortical afferents to caudate nucleus and putamen in primate: a fluorescence retrograde double labeling study. Neuroscience 10, 1137–1150.

Parent, A., Sato, F., Wu, Y., Gauthier, J., Levesque, M., Parent, M., 2000. Organization of the basal ganglia: the importance of axonal collateralization. Trends Neurosci. 23, S20–S27.

Parent, M., Levesque, M., Parent, A., 2001. Two types of projection neurons in the internal pallidum of primates: single-axon tracing and three-dimensional reconstruction. J. Comp. Neurol. 439, 162–175.

Park, M.R., Falls, W.M., Kitai, S.T., 1982. An intracellular HRP study of the rat globus pallidus. I. Responses and light microscopy analysis. J. Comp. Neurol. 211, 284–294.

Pennartz, C.M., McNaughton, B.L., Mulder, A.B., 2000. The glutamate hypothesis of reinforcement learning. Prog. Brain Res. 126, 231–253.

Penny, G.R., Wilson, C.J., Kitai, S.T., 1988. Relationship of the axonal and dendritic geometry of spiny projection neurons to the compartmental organization of the neostriatum. J. Comp. Neurol. 269, 275–289.

Percheron, G., Filion, M., 1991. Parallel processing in the basal ganglia: up to a point. Trends Neurosci. 14, 55–56.

Percheron, G., Yelnik, J., Francois, C., 1984. A Golgi analysis of the primate globus pallidus. III. Spatial organization of the striato-pallidal complex. J. Comp. Neurol. 227, 214–227.

Percheron, G., Francois, C., Yelnik, J., 1987. Spatial organization and information processing in the core of the basal ganglia. In: Carpenter, M.B., Jayaraman, A. (Eds.), The Basal Ganglia, vol. II. Plenum Press, New York, pp. 205–226.

Percheron, G., Francois, C., Yelnik, J., Fenelon, G., Talbi, B., 1994. The basal ganglia related system of primates: defintion, description and informational analysis. In: Percheron, G., McKenzie, J.S., Feger, J. (Eds.), The Basal Ganglia, vol. IV. Plenum Press, New York.

Perkel, D.H., Gerstein, G.L., Moore, G.P., 1967. Neuronal spike trains and stochastic point processes. II. Simultaneous spike trains. Biophys. J. 7, 419–440.

Plenz, D., 2003. When inhibition goes incognito: feedback interaction between spiny projection neurons in striatal function. Trends Neurosci. 26, 436–443.

Plenz, D., Kitai, S.T., 1999. A basal ganglia pacemaker formed by the subthalamic nucleus and external globus pallidus. Nature 400, 677–682.

Plumbley, M., 2001. Adaptive lateral inhibition for non-negative ICA. In: Proceedings of the International Conference on Independent Component Analysis and Blind Signal Separation (ICA2001).

Rafols, J.A., Fox, C.A., 1976. The neurons in the primate subthalamic nucleus: a Golgi and electron microscopic study. J. Comp. Neurol. 168, 75–111.

Ramanathan, S., Hanley, J.J., Deniau, J.M., Bolam, J.P., 2002. Synaptic convergence of motor and somatosensory cortical afferents onto GABAergic interneurons in the rat striatum. J. Neurosci. 22, 8158–8169.

Ravel, S., Sardo, P., Legallet, E., Apicella, P., 2001. Reward unpredictability inside and outside of a task context as a determinant of the responses of tonically active neurons in the monkey striatum. J. Neurosci. 21, 5730–5739.

Ravel, S., Legallet, E., Apicella, P., 2003. Responses of tonically active neurons in the monkey striatum discriminate between motivationally opposing stimuli. J. Neurosci. 23, 8489–8497.

Raz, A., Feingold, A., Zelanskaya, V., Vaadia, E., Bergman, H., 1996. Neuronal synchronization of tonically active neurons in the striatum of normal and parkinsonian primates. J. Neurophysiol. 76, 2083–2088.

Raz, A., Vaadia, E., Bergman, H., 2000. Firing patterns and correlations of spontaneous discharge of pallidal neurons in the normal and the tremulous 1-methyl-4-phenyl-1,2,3,6-tetrahydropyridine vervet model of parkinsonism. J. Neurosci. 20, 8559–8571.

Raz, A., Frechter-Mazar, V., Feingold, A., Abeles, M., Vaadia, E., Bergman, H., 2001. Activity of pallidal and striatal tonically active neurons is correlated in mptp-treated monkeys but not in normal monkeys. J. Neurosci. 21, 128.

Redgrave, P., Marrow, L., Dean, P., 1992. Topographical organization of the nigrotectal projection in rat: evidence for segregated channels. Neuroscience 50, 571–595.

Redgrave, P., Prescott, T.J., Gurney, K., 1999. Is the short-latency dopamine response too short to signal reward error? Trends Neurosci. 22, 146–151.

Reiner, A., Medina, L., Haber, S.N., 1999. The distribution of dynorphinergic terminals in striatal target regions in comparison to the distribution of substance P-containing and enkephalinergic terminals in monkeys and humans. Neuroscience 88, 775–793.

Rescorla, R.A., Wagner, A.R., 1972. A theory of Pavlovian conditioning: variations in the effectiveness of reinforcement and non-reinforcement. In: Black, A.J., Prokasy, W.F. (Eds.), Classical Conditioning. II. Current Research and Theory. Appelton-Century Crofts, New York, pp. 64–99.

Reynolds, J.N., Wickens, J.R., 2002. Dopamine-dependent plasticity of corticostriatal synapses. Neural Netw. 15, 507–521.

Reynolds, J.N., Hyland, B.I., Wickens, J.R., 2001. A cellular mechanism of reward-related learning. Nature 413, 67–70.

Robbins, T.W., Everitt, B.J., 1996. Neurobehavioural mechanisms of reward and motivation. Curr. Opin. Neurobiol. 6, 228–236.

Rubner, J., Tavan, P., 1989. A self-organizing network for principle component analysis. Europhys. Lett. 10, 693–698.

Rubner, J., Schulten, K., 1990. Development of feature detectors by self-organization: a network model. Biol. Cybern. 62, 193–199.

Rumelhart, D.E., Zipser, D., 1985. Feature discovery by competitive learning. Cogn. Sci. 9, 75–112.

Ryan, L.J., Clark, K.B., 1992. Alteration of neuronal responses in the subthalamic nucleus following globus pallidus and neostriatal lesions in rats. Brain Res. Bull. 29, 319–327.

Sanger, T.D., 1989. Optimal unsupervised learning in a single-layer network. Neural Netw. 2, 459–473.

Sato, F., Lavallee, P., Levesque, M., Parent, A., 2000. Single-axon tracing study of neurons of the external segment of the globus pallidus in primate. J. Comp. Neurol. 417, 17–31.

Sawyer, S.F., Tepper, J.M., Young, S.J., Groves, P.M., 1985. Antidromic activation of dorsal raphe neurons from neostriatum: physiological characterization and effects of terminal autoreceptor activation. Brain Res. 332, 15–28.

Scannell, J.W., Blakemore, C., Young, M.P., 1995. Analysis of connectivity in the cat cerebral cortex. J. Neurosci. 15, 1463–1483.

Schroder, K.F., Hopf, A., Lange, H., Thorner, G., 1975. Morphometrisch–statistische strukturanalysen des striatum, pallidum und nucleus subthalamicus beim Menschen [Morphometrical–statistical structure analysis of human striatum, pallidum and subthalamic nucleus]. J. Hirnforsch. 16, 333–350.

Schultz, W., 1982. Depletion of dopamine in the striatum as an experimental model of Parkinsonism: direct effects and adaptive mechanisms. Prog. Neurobiol. 18, 121–166.

Schultz, W., 1998. Predictive reward signal of dopamine neurons. J. Neurophysiol. 80, 1–27.

Schultz, W., Dayan, P., Montague, P.R., 1997. A neural substrate of prediction and reward. Science 275, 1593–1599.

Schwyn, R.C., Fox, C.A., 1974. The primate substantia nigra: a Golgi and electron microscopic study. J. Hirnforsch. 15, 95–126.

Selemon, L.D., Goldman Rakic, P.S., 1985. Longitudinal topography and interdigitation of corticostriatal projections in the rhesus monkey. J. Neurosci. 5, 776–794.

Sethi, K.D., 2002. Clinical aspects of Parkinson disease. Curr. Opin. Neurol. 15, 457–460.

Shepherd, G.M., 1998. The Synaptic Organization of the Brain. Oxford University Press.

Sherman, S.M., Guillery, R. W., 2001. Exploring the Thalamus. Academic Press, San Diego.

Shink, E., Smith, Y., 1995. Differential synaptic innervation of neurons in the internal and external segments of the globus pallidus by the GABA- and glutamate-containing terminals in the squirrel monkey. J. Comp. Neurol. 358, 119–141.

Shoulson, I., Glaubiger, G.A., Chase, T.N., 1975. On–off response: clinical and biochemical correlations during oral and intravenous levodopa administration in parkinsonian patients. Neurology 25, 1144–1148.

Smith, Y., Bennett, B.D., Bolam, J.P., Parent, A., Sadikot, A.F., 1994. Synaptic relationships between dopaminergic afferents and cortical or thalamic input in the sensorimotor territory of the striatum in monkey. J. Comp. Neurol. 344, 1–19.

Smith, Y., Bevan, M.D., Shink, E., Bolam, J.P., 1998a. Microcircuitry of the direct and indirect pathways of the basal ganglia. Neuroscience 86, 353–387.

Smith, Y., Shink, E., Sidibe, M., 1998b. Neuronal circuitry and synaptic connectivity of the basal ganglia. Neurosurg. Clin. N. Am. 9, 203–222.

Sporns, O., Tononi, G., Edelman, G.M., 2002. Theoretical neuroanatomy and the connectivity of the cerebral cortex. Behav. Brain Res. 135, 69–74.

Stanford, I.M., 2003. Independent neuronal oscillators of the rat globus pallidus. J. Neurophysiol. 89, 1713–1717.

Steiner, H., Gerfen, C.R., 1998. Role of dynorphin and enkephalin in the regulation of striatal output pathways and behavior. Exp. Brain. Res. 123, 60–76.

Stern, E.A., Kincaid, A.E., Wilson, C.J., 1997. Spontaneous subthreshold membrane potential fluctuations and action potential variability of rat corticostriatal and striatal neurons in vivo. J. Neurophysiol. 77, 1697–1715.

Stern, E.A., Jaeger, D., Wilson, C.J., 1998. Membrane potential synchrony of simultaneously recorded striatal spiny neurons in vivo. Nature 394, 475–478.

Suri, R.E., 2002. TD models of reward predictive responses in dopamine neurons. Neural Netw. 15, 523–533.
Suri, R.E., Bargas, J., Arbib, M.A., 2001. Modeling functions of striatal dopamine modulation in learning and planning. Neuroscience 103, 65–85.
Surmeier, D.J., Song, W.J., Yan, Z., 1996. Coordinated expression of dopamine receptors in neostriatal medium spiny neurons. J. Neurosci. 16, 6579–6591.
Sutton, R.S., 1988. Learning to predict by the methods of temporal difference. Mach. Learn. 3, 9–44.
Sutton, R.S., Barto, A.G., 1998. Reinforcement Learning—An Introduction. MIT Press, Cambridge, MA.
Szabo, J., 1980. Organization of the ascending striatal afferents in monkeys. J. Comp. Neurol. 189, 307–321.
Takada, M., Tokuno, H., Nambu, A., Inase, M., 1998. Corticostriatal projections from the somatic motor areas of the frontal cortex in the macaque monkey: segregation versus overlap of input zones from the primary motor cortex, the supplementary motor area, and the premotor cortex. Exp. Brain Res. 120, 114–128.
Takada, M., Tokuno, H., Hamada, I., Inase, M., Ito, Y., Imanishi, M., Hasegawa, N., Akazawa, T., Hatanaka, N., Nambu, A., 2001. Organization of inputs from cingulate motor areas to basal ganglia in macaque monkey. Eur. J. Neurosci. 14, 1633–1650.
Tepper, J.M., Trent, F., 1993. In vivo studies of the postnatal development of rat neostriatal neurons. In: Arbuthnott, G., Emson P.C. (Eds.), Chemical Signaling in the Basal Ganglia, vol. 99. Elsevier, Amsterdam, pp. 35–50.
Terman, D., Rubin, J.E., Yew, A.C., Wilson, C.J., 2002. Activity patterns in a model for the subthalamopallidal network of the basal ganglia. J. Neurosci. 22, 2963–2976.
Tesauro, G., 1994. TD-gammon: a self-teaching backgammon program, achieves master-level play. Neural Comp. 6, 215–219.
Thorndike, E.L., 1911. Animal Intelligence. Macmillan, New York.
Thorner, G., Lange, H., Hopf, A., 1975. Morphometrical-statistical structure analysis of human striatum, pallidus and subthalamic nucleus. II. Globus pallidus. J. Hirnforsch. 16, 401–413.
Tokuno, H., Inase, M., Nambu, A., Akazawa, T., Miyachi, S., Takada, M., 1999. Corticostriatal projections from distal and proximal forelimb representations of the monkey primary motor cortex. Neurosci. Lett. 269, 33–36.
Tremblay, L., Filion, M., 1989. Responses of pallidal neurons to striatal stimulation in intact waking monkeys. Brain Res. 498, 1–16.
Tunstall, M.J., Oorschot, D.E., Kean, A., Wickens, J.R., 2002. Inhibitory interactions between spiny projection neurons in the rat striatum. J. Neurophysiol. 88, 1263–1269.
Waelti, P., Dickinson, A., Schultz, W., 2001. Dopamine responses comply with basic assumptions of formal learning theory. Nature 412, 43–48.
Wagner, S., Castel, M., Gainer, H., Yarom, Y., 1997. GABA in the mammalian suprachiasmatic nucleus and its role in diurnal rhythmicity. Nature 387, 598–603.
Watkins, C.J.C.H., Dayan, P., 1992. Q learning. Mach. Learn. 8, 279–292.
Wichmann, T., DeLong, M.R., 1996. Functional and pathophysiological models of the basal ganglia. Curr. Opin. Neurobiol. 6, 751–758.
Wichmann, T., DeLong, M.R., 2003. Functional neuroanatomy of the basal ganglia in Parkinson's disease. Adv. Neurol. 91, 9–18.
Wichmann, T., Bergman, H., DeLong, M.R., 1994. The primate subthalamic nucleus. I. Functional properties in intact animals. J. Neurophysiol. 72, 494–506.
Wickens, J., 1993. A Theory of the Striatum. Pregamon Press, Oxford.
Wickens, J., 1997. Basal ganglia: structure and computation. Netw: Comput. Neural Syst. 8, R77–R109.
Wickens, J., Oorschot, D.E., 2000. Neural dynamics and surround inhibition in the neostriatum: a possible connection. In: Miller, R., Wickens, R. (Eds.), Brain Dynamics and the Striatal Complex. Harwood Academic Publishers, pp. 141–149.
Wickens, J.R., Wilson, C.J., 1998. Regulation of action-potential firing in spiny neurons of the rat neostriatum in vivo. J. Neurophysiol. 79, 2358–2364.
Wickens, J.R., Begg, A.J., Arbuthnott, G.W., 1996. Dopamine reverses the depression of rat corticostriatal synapses which normally follows high-frequency stimulation of cortex in vitro. Neuroscience 70, 1–5.
Wilms, H., Sievers, J., Deuschl, G., 1999. Animal models of tremor. Mov. Disord. 14, 557–571.
Wilson, C.J., 1993. The generation of natural firing patterns in neostriatal neurons. Prog. Brain Res. 99, 277–297.
Wilson, C.J., 1995. The contribution of cortical neurons to the firing pattern of striatal spiny neurons. In: Houk, J.C., Davis, J.L., Beiser, D.G. (Eds.), Models of Information Processing in the Basal Ganglia. MIT Press, pp. 29–50.
Wilson, C.J., 1998. Basal ganglia. In: Shepherd, G.M. (Ed.), The Synaptic Organization of the Brain. Oxford University Press, Oxford, pp. 329–375.
Wilson, C.J., 2000. Striatal circuitry: categorically selective, or selectively categorical? In: Miller, R., Wickens, J.R. (Eds.), Brain Dynamics and the Striatal Complex. Harwood, Amsterdam, pp. 289–306.
Wilson, C.J., Groves, P.M., 1980. Fine structure and synaptic connections of the common spiny neuron of the rat neostriatum: a study employing intracellular injection of horseradish peroxidase. J. Comp. Neurol. 194, 599–615.
Wilson, C.J., Kawaguchi, Y., 1996. The origins of two-state spontaneous membrane potential fluctuations of neostriatal spiny neurons. J. Neurosci. 16, 2397–2410.
Wilson, C.J., Groves, P.M., Kitai, S.T., Linder, J.C., 1983. Three-dimensional structure of dendritic spines in the rat neostriatum. J. Neurosci. 3, 383–388.
Wilson, C.J., Chang, H.T., Kitai, S.T., 1990. Firing patterns and synaptic potentials of identified giant aspiny interneurons in the rat neostriatum. J. Neurosci. 10, 508–519.
Wise, R.A., Rompre, P.P., 1989. Brain dopamine and reward. Annu. Rev. Psychol. 40, 191–225.
Yelnik, J., 2002. Functional anatomy of the basal ganglia. Mov. Disord. 17 (Suppl. 3), S15–S21.
Yelnik, J., Percheron, G., 1979. Subthalamic neurons in primates: a quantitative and comparative analysis. Neuroscience 4, 1717–1743.
Yelnik, J., Francois, C., Percheron, G., Heyner, S., 1987. Golgi study of the primate substantia nigra. I. Quantitative morphology and typology of nigral neurons. J. Comp. Neurol. 265, 455–472.
Yelnik, J., Percheron, G., Francois, C., 1984. A Golgi analysis of the primate globus pallidus. II. Quantitative morphology and spatial orientation of dendritic arborizations. J. Comp. Neurol. 227, 200–213.
Yelnik, J., Francois, C., Tand, D., 1997. Etude tridimensionnelle des collaterales initiales des neurones du pallidum interne chez le macaque. In: Proceedings of the Third Congress of European Neuroscience Society, Bordeaux, p. 104.
Yeterian, E.H., VanHoesen, G.W., 1978. Cortico-striate projections in the rhesus monkey: the organization of certain cortico-caudate connections. Brain Res. 139, 43–63.
Yoshida, S., Nambu, A., Jinnai, K., 1993. The distribution of the globus pallidus neurons with input from various cortical areas in the monkeys. Brain Res. 611, 170–174.
Zheng, T., Wilson, C.J., 2002. Corticostriatal combinatorics: the implications of corticostriatal axonal arborizations. J. Neurophysiol. 87, 1007–1017.