REINFORCEMENT-DRIVEN DIMENSIONALITY REDUCTION - A MODEL FOR INFORMATION PROCESSING IN THE BASAL GANGLIA

Izhar Bar-Gad[1], Gali Havazelet-Heimer[3], Joshua A. Goldberg[1,3], Eytan Ruppin[2] and Hagai Bergman[1,3]*

[1]*Center for Neural Computation, Hebrew University, Jerusalem,*
[2]*Departments of Computer Science and Physiology,*
Tel Aviv University, Tel Aviv
[3]*Department of Physiology,*
Hebrew University-Hadassah Medical School, Jerusalem, Israel

ABSTRACT

Although anatomical studies of the basal ganglia show the existence of extensive convergence and lateral inhibitory connections, physiological studies failed to show correlated neural activity or lateral interaction in these nuclei. These seemingly contradictory results could be explained with a model in which the basal ganglia reduce the dimensionality of cortical information using optimal extraction methods. Simulations of this model predict a transient change in the efficacy of the feed-forward and lateral synapses following changes in reinforcement signal, causing an increase in correlated firing rates. This process ultimately restores the steady-state situation with diminished efficacy of lateral inhibition and no correlation of firing. Our experimental results confirm the model's predictions: rate correlations show a drastic decrease between the input stage (cortex) and output stage (pallidum). Moreover, preliminary analysis revealed that pallidal correlations show a transient increase following discrepancies

* Address for correspondence:
Hagai Bergman
Department of Physiology
The Hebrew University – Hadassah Medical School
P.O. Box 12272
Jerusalem 91120, Israel
e-mail: hagaib@md.huji.ac.il

between the animal's predictions and reality. We therefore propose that by using a reinforcement-driven dimensionality reduction process the basal ganglia achieve efficient extraction of cortical salient information that may then be used by the frontal cortex for execution and planning of forthcoming actions.

KEY WORDS

cross-correlation, globus pallidus, Parkinson's disease, MPTP, monkey, dopamine

INTRODUCTION

A major pathway in the basal ganglia circuitry leads from most cortical areas to the input stage of the basal ganglia, the striatum. Subsequent projections link striatal neurons to the output stages of the basal ganglia, e.g., the internal segment of the globus pallidus (GPi). The GPi projects back to the frontal cortex through thalamic relay stations (Fig. 1a) /1,2/.

The cortico-striatal-pallidal pathway is characterized by a high degree of anatomical convergence. The number of cortical neurons projecting to the striatum is two orders of magnitude greater than the number of striatal neurons /3/ and an additional reduction of the same magnitude occurs from the striatum to the GPi /4,5/. The basal ganglia are further characterized by GABAergic inhibitory connections. Anatomical studies have shown that most striatal and pallidal neurons are GABAergic projection neurons /1,2/. These GABAergic neurons also form massive collateral connections within their nuclei of origin /2,6-8/, suggesting the possibility of strong collateral inhibition. This prediction, however, has been thwarted by recent physiological intracellular studies in which no evidence was found for functional synaptic interactions between striatal projection neurons /9/. In accord with these physiological results, but in apparent contrast to the anatomical data, recent studies of the striatum /10,11/ and pallidum /12,13/ have failed to reveal correlation between the spiking activity of simultaneously recorded neurons. This low level of correlated spiking activity in the basal ganglia is in sharp contrast with the significantly correlated discharge of neighboring cortical neurons /14-16/.

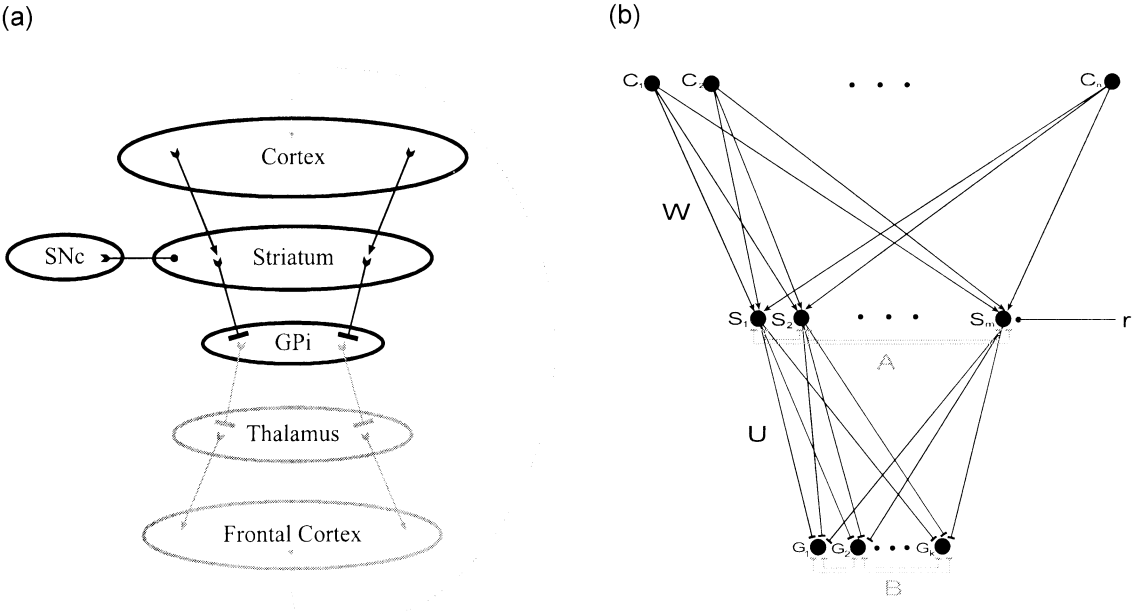

Fig. 1: Structure of the reinforcement driven dimensionality reduction network. Schematic diagram of the cortico-basal ganglia-cortical circuit (black: layers incorporated in the model; gray: layers not included in the model). The model is composed of a three-layered feed-forward network simulating the cortico-striato-pallidal circuit with lateral inhibitory connections at the intermediate (striato) and output (pallidal) layers. A reinforcement signal is provided at the intermediate layer. Arrow-head connections represent glutamatergic excitatory synapses, square-head connections represent GABAergic inhibitory synapses and round-head connections represent dopaminergic modulatory synapses.

Most previous models of basal ganglia function have been influenced by the anatomical evidence of strong lateral connectivity, and assume strong mutual inhibition between striatal neurons /17/. This leads to the predominant view that the basal ganglia function as an action selection network /18-20/. However, these models fail to incorporate the physiological data concerning the paucity of the intra-nuclear interactions, and erroneously predict strong lateral interactions and negative correlation of intra-nuclear neuronal firing. Here, we present an alternative hypothesis claiming that the basal ganglia perform efficient dimensionality reduction and decorrelation of the large and complex information space spanned by the activity of cortical neurons. Efficient reduction is achieved when all or most of the information contained within the original space is preserved. Theoretical studies demonstrate that neural networks can perform such efficient coding using competitive Hebbian learning rules for inter-layer connectivity /21/ and anti-Hebbian rules for the lateral inhibitory connectivity /22,23/.

Dimensionality reduction in a behaving animal should be affected not only by the statistical properties of the input patterns but also by their behavioral significance. Such behaviorally significant signals /24-26/ are received in the striatum from striatal cholinergic interneurons /27/ and from midbrain dopaminergic neurons /28/. Anatomical and physiological studies have shown that these reinforcement signals modulate the access of striatal neurons to cortical input /29-31/. Our working hypothesis was, therefore, extended to suggest that the basal ganglia network performs reinforcement-driven dimensionality reduction (RDDR) of their cortical inputs.

MATERIALS AND METHODS

To examine the RDDR hypothesis, we first studied a simulated neural network incorporating key aspects of basal ganglia connectivity (Fig. 1b). The three layer feed-forward network consisted of neurons with linear activation functions and lateral connectivity within the layers. The simulations employed an input layer of 64 "cortical" neurons, an intermediate layer of 16 "striatal" neurons and an output layer of four "pallidal" neurons. The feed-forward weights were initialized to random values and the lateral weights were initialized to zero.

Learning was Hebbian for the feed-forward weights and anti-Hebbian for the lateral weights. Learning in the intermediate (striatal) layer was regulated by a reinforcement signal. The cortico-striatal-dopaminergic triple synapse is modeled by a reinforcement signal controlling the feed-forward Hebbian learning. The control signal is positive for reward-related events and zero for non-reward-related events (baseline dopamine levels). Dopamine depletion, such as in Parkinson's disease, is modeled by negative reinforcement values /28/. All learning rules were local and no global information was used /22,23/. The synaptic weights were constrained according to the known physiology and anatomy of the basal ganglia. Thus, positive weights simulated glutamatergic synapses and negative values simulated GABAergic synapses.

To measure the information loss of the network due to the RDDR process, the output layer was expanded back to an input-size space using the same weights. This procedure recreates the reconstructed decompressed pattern. The reconstruction error is the mean squared difference between the normalized original and reconstructed elements over the input patterns. The high-dimensional input patterns actually lie within a lower dimensional sub-space. Input patterns were created by summing a small number of independent patterns /32/. The number of originating patterns is 16; out of these patterns 12 receive low reinforcement and four receive higher reinforcement. After reaching steady state values, the reinforcement signal changes to either reinforcing a new set of input patterns or to negative reinforcement values.

In the second stage we tested the predictions of the model by extending the studies of simultaneous activity of neurons from the frontal cortex /33/ and the pallidum /34/ of macaque and vervet monkeys engaged in a similar "Go/No-Go" behavioral task. The spike train of each neuron was converted to a continuous rate signal, and the covariance of the two rate functions was calculated for lags of up to ± 10 seconds for the entire data stream (Fig. 2a, left). The significance of the covariance at zero time lag was calculated relative to the variance of the rate cross correlation at large lags (Fig. 2a, right).

Finally, we tested the predictions of the model by studying the dynamics of firing covariance of simultaneously recorded pallidal cells. We trained a vervet monkey to perform a key-pressing task and recorded its pallidal activity during task performance. The monkey

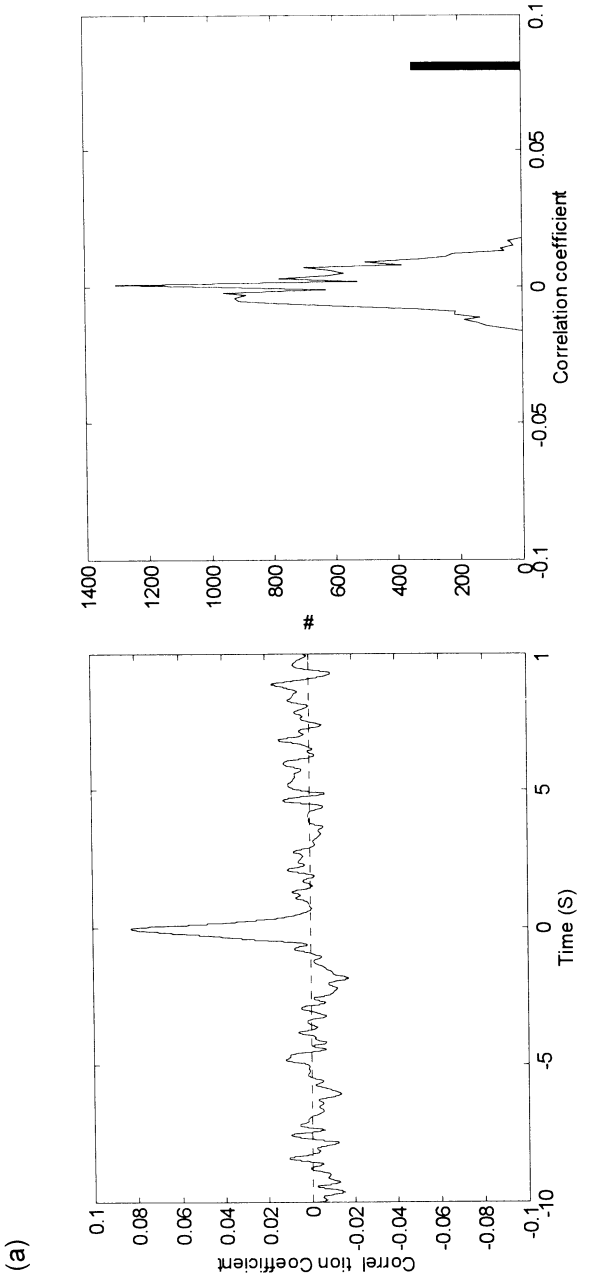
(a)
Correlation Coefficient
0.1
0.08
0.06
0.04
0.02
0
-0.02
-0.04
-0.06
-0.08
-0.1
-10
-5
0
5
1
Time (S)
#
1400
1200
1000
800
600
400
200
0
-0.1
-0.05
0
0.05
0.1
Correlation coefficient

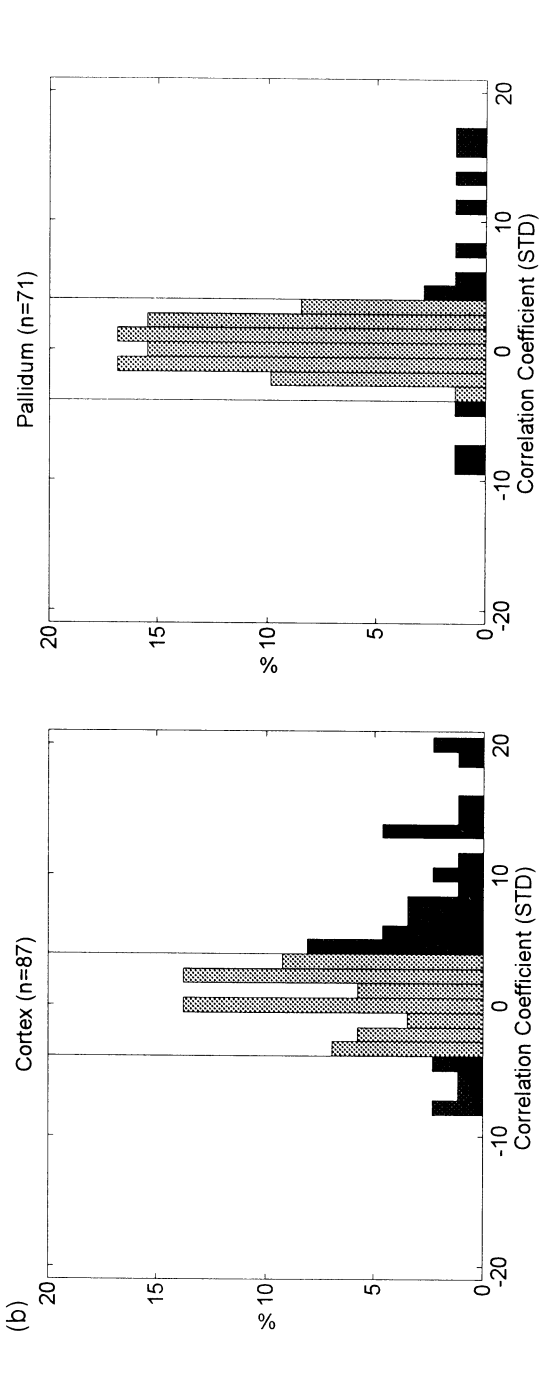

Fig. 2: **Rate cross-correlation of cortical and pallidal neuron pairs.** (**a**) An example of rate cross-correlations in the range of ±10 seconds lag (left). A bar indicating the value of the cross correlation at time shift 0 is plotted versus the distribution of correlation values at large time shifts (right). The variance of the distribution of correlation values at large time shifts equals 0.0062. The value of the correlation coefficient at time shift 0 is 0.0812, and hence equals 0.0812/0.0062=12.9 standard deviation units of the distribution. (**b**) Histogram of cross-correlation values at time shift 0 (in units of standard deviation). The lines are at the $p<10^{-4}$ significance levels (3.89 STD units). Significant positive and negative values are given by black bars, non-significant values are given by gray bars.

was facing three pushbuttons and rewarded with water for touching the central key. The spike trains of the pallidal neurons were converted to rate functions and a sliding correlation coefficient was calculated for pairs of pallidal neurons as a function of time. Behavioral events leading to transient changes in rate correlation were defined as events appearing up to 60 seconds before an occurrence of a significant value ($p<0.01$) of the correlation coefficients.

RESULTS

Following a change in the reinforcement signal, the simulated basal ganglia network (Fig. 1b) performed sub-optimal information compression and the activity of the output neurons became correlated. This correlation caused an increase in the efficacies of the inhibitory lateral synapses (Fig. 3a) and changes in the efficacies of the feed-forward connections. These changes, in turn, resulted in decorrelation of neuronal activity within the output layer (Fig. 3b) and improvement of the information compression. The reinforcement signal caused the extraction to become discriminative, performing better for reward-related inputs but not for unrelated events (Fig. 3c). Similar behavior was observed following presentation of novel input patterns (data not shown). Overall, the modification of the neural network replicates the paradoxical findings concerning the basal ganglia: uncorrelated activity of the output neurons and diminished efficacy of lateral synaptic interactions in the steady-state period.

The analysis of the simultaneous activity of neurons from the frontal cortex /33/ and the pallidum /34/ revealed that the firing rate of 41.4% of cortical pairs covaried significantly, whereas only 15.5% of the pallidal pairs showed significant rate covariance ($p<10^{-4}$) at zero time lag. The fraction of negatively correlated pairs out of all significantly correlated pairs was small in both recording areas (Fig. 2b). The mean absolute rate covariance of pallidal pairs was 2.6 standard deviations while cortical pairs had a mean absolute covariance of 4.8 standard deviations. The finding that the fraction of significantly correlated pallidal pairs is smaller than that of cortical pairs is a robust one. Consistent results were obtained when different level of confidence limits were used (Table 1a) and for different time windows of the correlation peaks (Table 1b).

TABLE 1

Significant rate correlations in the cortex and the pallidum for various significance levels and time windows

(a)

Significance level (STD)	Cortex			Pallidum		
	+	–	Total	+	–	Total
10^{-3} (3.29)	40.2	11.5	**51.7**	15.5	5.6	**21.1**
10^{-4} (3.89)	34.5	6.9	**41.4**	11.3	4.2	**15.5**
10^{-5} (4.42)	28.7	5.7	**34.4**	9.9	2.8	**12.7**

(b)

Time window - Significance level (STD)	Cortex			Pallidum		
	+	–	Total	+	–	Total
1 ms - 10^{-3} (3.29)	40.2	11.5	**51.7**	15.5	5.6	**21.1**
100 ms - 10^{-5} (4.42)	33.3	6.9	**40.2**	11.3	2.8	**14.1**
1000 ms - 10^{-6} (4.89)	36.8	8.0	**44.8**	16.9	4.2	**21.1**

a) Percentage of cortical and pallidal pairs with significant correlations at time shift 0. Percentage is calculated for various significance levels. + = significant positive correlations, – = significant negative correlations.

b) Percentage of cortical and pallidal pairs whose maximal correlation within a certain time window around 0 exceeds a significance level. Significance levels are adjusted to reflect the duration of the examined window.

Finally, correlation coefficients of pairs of pallidal neurons were calculated as a function of time and were related to behavioral events. Our preliminary analysis revealed that the correlation coefficients were low during performance of a known task leading to an expected reward and during rest periods. However, significant positive and negative transients in correlations were observed following unpredicted behavioral events. The similar time courses of these transients enabled the detection of population phenomena by averaging the absolute correlation values over all simultaneously recorded pairs. A dramatic increase in absolute correlation values occurred following

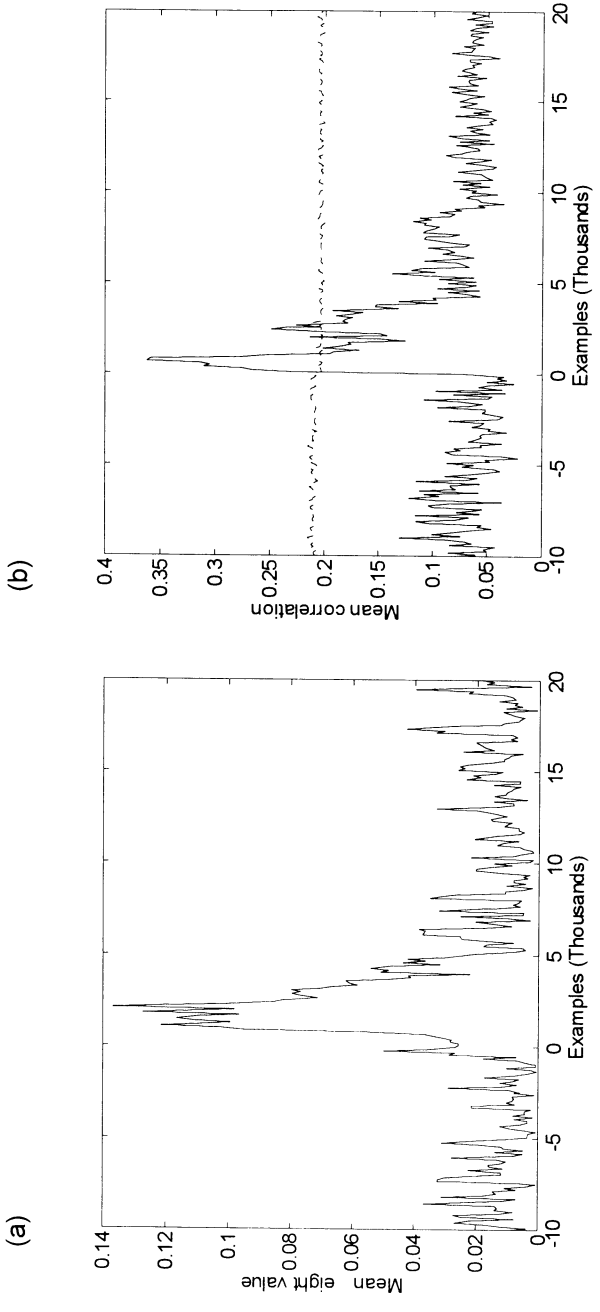
(a)
Mean eight value
0.14
0.12
0.1
0.08
0.06
0.04
0.02
0
-10
-5
0
5
10
15
20
Examples (Thousands)
(b)
Mean correlation
0.4
0.35
0.3
0.25
0.2
0.15
0.1
0.05
0
-10
-5
0
5
10
15
20
Examples (Thousands)

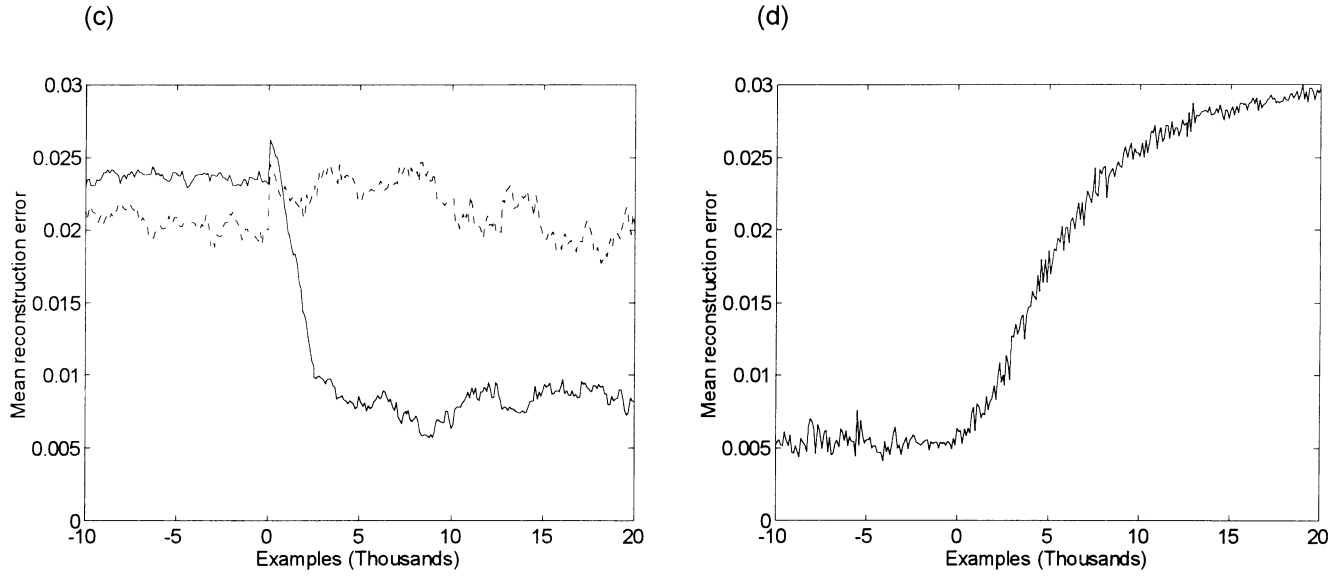

Fig. 3: Dynamic changes of the reinforcement driven dimensionality reduction network. The simulation starts at random values and reaches a steady state level. Then, at time=0, the reinforcement signal changes to either reinforcing a new set of input patterns (**a-c**) or to negative values (**d**). Presentation of examples to the model (abscissa) is analogous to the time axis and to stimuli presentation in the physiological experiments. The following variables are displayed as a function of the number of input examples: (**a**) mean absolute value of lateral weights of the output (pallidal) layer; (**b**) mean intra-layer correlation within the output (pallidal) layer (solid line) and within the input (cortical) layer (dotted line); (**c**) information loss due to the compression (reconstruction error). The error for reinforced patterns (solid line) decreases but does not change for non-reinforced patterns (dotted line). The reconstruction error is defined as the mean squared difference between the normalized original and reconstructed elements over the input patterns; (**d**) increase in reconstruction error due to constant negative reinforcement signal (simulating Parkinson's disease).

unexpected extra rewards given to the monkey. Significant changes in the correlation values were also found after withdrawal of reward for previously rewarded actions. Both types of events, i.e., unrewarded action and unpredicted reward, are characterized by a mismatch between prediction and reality. The proportion of events followed by high absolute correlation values was significantly higher following such mismatches than following predictable events.

DISCUSSION

The reinforcement-driven dimensionality reduction (RDDR) model emphasizes the role of the basal ganglia in extraction and pre-processing of information from the whole cortex. This compression process is eminently useful since it allows the transmission of large amounts of information within a limited number of axons and synapses /35,36/. In contrast to action-selection models of the basal ganglia, the RDDR model does not force a single selection but rather performs multi-dimensional encoding. For example, in a case of N binary units, an action selection process can encode only N states while a model similar to the RDDR could conceivably encode 2^N states. The RDDR network also provides a vehicle by which reinforcement learning may be carried out in the brain in a central, parsimonious location.

The model provides an explanation for the apparent lack of physiological expression of the lateral inhibitory connections and uncorrelated spiking activity observed in previous studies of the basal ganglia /9-13/. These physiological studies were carried out in adult animals which were not engaged in learning new skills and situations, whereas the RDDR model maintains that the lateral connections are functional only during the learning phase. After this phase, their efficacy decreases (Fig. 3a) as the activity of the output neurons becomes uncorrelated (Fig. 3b).

A critical prediction of the RDDR model is that rate cross-correlation will be reduced in neurons of the output stage (pallidum) relative to the input stage (cortex). This is shown in recordings of neuron pairs in the monkey: pallidal pairs have lower correlation coefficients and a smaller percentage of them are significantly correlated. Even those pallidal pairs with significant rate correlations show low spike-to-spike correlation /12/, further indicating their high information capacity. The highly correlated activity of the input

cortical neurons suggests that the lack of correlated activity in the basal ganglia cannot be explained merely by the sparse cortico-striatal connectivity /3/. These decreased correlations rather suggest an active decorrelating process.

Another important prediction of the RDDR model is that transient changes in basal ganglia intra-layer cross-correlation will occur during periods of learning. In general, periods of active learning follow detection of discrepancies between expectations and reality (e.g., unpredicted rewards or disappointing outcomes) /26/. Our preliminary finding of prolonged transient changes in the cross-correlations of pallidal pairs following such discrepancies suggests that reinforcement-driven dimensionality reduction is a major characteristic of basal ganglia physiology.

Finally, the RDDR mechanism also offers explanations for some open questions in the pathophysiology of movement disorders. For example, dopamine depletion (a negative reinforcement signal) as in Parkinson's disease substantially damages the RDDR process since no discrimination is possible between important and negligible information (Fig. 3d). Conventional dopamine replacement therapy restores the background level of dopamine. However, the pulsatile nature of the conventional treatment causes inevitable fluctuations in striatal dopamine /37/. These fluctuations are randomly timed relative to the actual reinforcement of the environment and therefore may result in the generation of random encoding and the development of dyskinesia.

ACKNOWLEDGEMENTS

This study was supported in part by the Israeli Academy of Science, AFIRST, and the US-Israel Binational Science Foundation. We thank Opher Donchin, Thomas Wichmann, Steve Wise, Eilon Vaadia and Moshe Abeles for their critical reading and helpful suggestions. We thank Hamutal Slovin, Aeyal Raz, Joshua A. Goldberg, Sharon Maraton and Genella Morris for helping with data acquisition.

REFERENCES

1. Parent A, Hazrati LN. Functional anatomy of the basal ganglia. I. The cortico-basal ganglia-thalamo-cortical loop. Brain Res Rev 1995; 20: 91-127.

2. Gerfen CR, Wilson CJ. The basal ganglia. In: Swanson LW, Björklund A, Hökfelt T, eds. Handbook of Chemical Neuroanatomy, Vol. 12: Integrated Systems of the CNS, Part III. Amsterdam: Elsevier Science, 1996; 371-468.
3. Kincaid AE, Zheng T, Wilson CJ. Connectivity and convergence of single corticostriatal axons. J Neurosci 1998; 18: 4722-4731.
4. Percheron G, Francois C, Yelnik J, Fenelon G, Talbi B. The basal ganglia related system of primates: Defintion, description and informational analysis. In: Percheron G, McKenzie JS, Feger J, eds. The Basal Ganglia IV. New York: Plenum Press, 1994; 3-20.
5. Oorshcot DE. 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 1996; 366: 580-599.
6. Kita H. Two pathways between the cortex and the basal ganglia output nuclei and the globus pallidus. In: Ohye C, Kimura M, McKenzie JS, eds. The Basal Ganglia V. New York: Plenum Press, 1996; 77-94.
7. Yelnik J, Francois C, Tand D. Etude tridimensionnelle des collaterales initiales des neurones du pallidum interne chez le macaque. 3rd Congress of European Neuroscience Society, Bordeaux, 1997; 104.
8. Sato F, Lavallee P, Levesque M, Parent A. Single-axon tracing study of neurons of the external segment of the globus pallidus in primate. J Comp Neurol 2000; 417: 17-31.
9. Jaeger D, Kita H, Wilson CJ. Surround inhibition among projection neurons is weak or nonexistent in the rat neostriatum. J Neurophysiol 1994; 72: 2555-2558.
10. Jaeger D, Gilman S, Aldridge JW. Neuronal activity in the striatum and pallidum of primates related to the execution of externally cued reaching movements. Brain Res 1995; 694: 111-127.
11. Stern EA, Jaeger D, Wilson CJ. Membrane potential synchrony of simultaneously recorded striatal spiny neurons in vivo. Nature 1998; 394: 475-478.
12. Nini A, Feingold A, Slovin H, Bergman H. 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 1995; 74: 1800-1805.
13. Bergman H, Feingold A, Nini A, Raz A, Slovin H, Abeles M, Vaadia E. Physiological aspects of information processing in the basal ganglia of normal and parkinsonian primates. Trends Neurosci 1998; 21: 32-38.
14. Eggermont JJ. The Correlative Brain. Theory and Experiment in Neuronal Interaction. Berlin: Springer-Verlag, 1990.
15. Nelson JI, Salin PA, Munk MH, Arzi M, Bullier J. Spatial and temporal coherence in cortico-cortical connections: a cross-correlation study in areas 17 and 18 in the cat. Vis Neurosci 1992; 9: 21-37.
16. Vaadia E, Haalman I, Abeles M, Bergman H, Prut Y, Slovin H, Aertsen A. Dynamics of neuronal interactions in monkey cortex in relation to behavioral events. Nature 1995; 373: 515-518.

17. Beiser DG, Hua SE, Houk JC. Network models of the basal ganglia. Curr Opin Neurobiol 1997; 7: 185-190.
18. Wickens J. A Theory of the Striatum. Oxford: Pergamon Press, 1993.
19. Mink JW. The basal ganglia: focused selection and inhibition of competing motor programs. Prog Neurobiol 1996; 50: 381-425.
20. Berns GS, Sejnowski TJ. A computational model of how the basal ganglia produce sequences. J Cogn Neurosci 1998; 10: 108-121.
21. Oja E. A simplified neuron model as a principal component analyzer. J Math Biol 1982; 15: 267 -273.
22. Kung SY, Diamantaras KI. A neural network learning algorithm for adaptive principal component extraction (APEX). Proceedings IEEE International Conference on Acoustics, Speech and Signal Processing 1990; 2: 861-864.
23. Foldiak P. Adaptive network for optimal feature extraction. Proceedings International Joint Conference on Neural Networks 1989; 1: 401-405.
24. Redgrave P, Prescott TJ, Gurney K. Is the short-latency dopamine response too short to signal reward error? Trends Neurosci 1999; 22: 146-151.
25. Berridge KC, Robinson TE. What is the role of dopamine in reward: hedonic impact, reward learning, or incentive salience? Brain Res Brain Res Rev 1998; 28: 309-369.
26. Robbins TW, Everitt BJ. Neurobehavioural mechanisms of reward and motivation. Curr Opin Neurobiol 1996; 6: 228-236.
27. Graybiel AM, Aosaki T, Flaherty AW, Kimura M. The basal ganglia and adaptive motor control. Science 1994; 265: 1826-1831.
28. Schultz W. Predictive reward signal of dopamine neurons. J Neurophysiol 1998; 80: 1-27.
29. Freund TF, Powell JF, Smith AD. Tyrosine hydroxylase-immunoreactive boutons in synaptic contact with identified striatonigral neurons, with particular reference to dendritic spines. Neuroscience 1984; 13: 1189-1215.
30. Centonze D, Gubellini P, Picconi B, Calabresi P, Giacomini P, Bernardi G. Unilateral dopamine denervation blocks corticostriatal LTP. J Neurophysiol 1999; 82: 3575-3579.
31. Calabresi P, Centonze D, Gubellini P, Pisani A, Bernardi G. Acetylcholine-mediated modulation of striatal function. Trends Neurosci 2000; 23: 120-126.
32. Rumelhart DE, Zipser D. Feature discovery by competitive learning. Cogn Science 1985; 9: 75-112.
33. Slovin H, Abeles M, Vaadia E, Haalman I, Prut Y, Bergman H. Frontal cognitive impairments and saccadic deficits in low-dose MPTP-treated monkeys. J Neurophysiol 1999; 81: 858-874.
34. Raz A, Feingold A, Zelanskaya V, Vaadia E, Bergman H. Neuronal synchronization of tonically active neurons in the striatum of normal and parkinsonian primates. J Neurophysiol 1996; 76: 2083-2088.
35. Linsker R. Perceptual neural organization: some approaches based on network models and information theory. Ann Rev Neurosci 1990; 13: 257-281.
36. Dan Y, Atick JJ, Reid RC. Efficient coding of natural scenes in the lateral geniculate nucleus: experimental test of a computational theory. J Neurosci 1996; 16: 3351-3362.

37. Sage JI, Mark MH. The rationale for continuous dopaminergic stimulation in patients with Parkinson's disease. Neurology 1992; 42: 23-28.