

UNIVERSITY OF MINNESOTA

This is to certify that I have examined this copy of a doctoral Dissertation by

Neil Chadwick Schmitzer-Torbert

and have found that it is complete and satisfactory in all respects,
and that any and all revisions required by the final
examining committee have been made.

A. David Redish

Name of Faculty Adviser

Signature of Faculty Adviser

Date

GRADUATE SCHOOL

THE INVOLVEMENT OF THE RODENT STRIATUM IN
NAVIGATION

A DISSERTATION
SUBMITTED TO THE FACULTY OF THE GRADUATE SCHOOL
OF THE UNIVERSITY OF MINNESOTA
BY

Neil Chadwick Schmitzer-Torbert

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

A. David Redish, Adviser

December 2004

© Copyright by Neil Chadwick Schmitzer-Torbert 2005

Acknowledgements

The experiments that I have done over the past four years in the Redish lab would not have been possible without help from many people. In the lab, I have had the benefit of working with many great technicians, including Mallika Arudi, Deborah Bang, Dan Bernal, Chris Boldt, Giuseppe Cortese, Sarah Jutila, Monica Kumar, Susan Nwoke, Liz Ortiz, Kelsey Seeland, Saumya Rao, and Tim Singewald, who all assisted with the experiments described in this thesis. I have also been fortunate to work with a number of graduate students in the lab, including Pratibha Aia, Jadin Jackson, Adam Johnson, Beth Masimore, Saumya Rao, and Jon Waataja (during his lab rotation). I am particularly indebted to Deborah Bang and Sarah Jutila for technical work on the histology, and Maureen Reidl of the Elde lab for her assistance in this area. Kelsey Seeland, Chris Boldt and Giuseppe Cortese have each given heroic efforts over the years to assemble the hyperdrives used for chronic recordings. Mallika Arudi and Giuseppe Cortese have assisted in most of the surgeries, and I have had no shortage of willing volunteers to lend a hand during hyperdrive implantations. Chris Boldt was also kind enough to run some of my experiments while I was out of town, for which I am very thankful.

Jadin Jackson, Adam Johnson and Beth Masimore deserve special thanks for helpful discussions about experiments and in the development of interesting analyses. They have also been good friends, which is, I think, more valuable.

During my graduate career, I have been fortunate to receive funding from several sources, including an NSF Graduate Research Fellowship, over a year of support from NSF-IGERT Training Grant #9870633 and a year of funding by a University of Minnesota Graduate Fellowship.

Abstract

The basal ganglia are a collection of nuclei which have important roles in motor control and cognitive processing, and are the target of several human neurodegenerative disorders. The experiments described here examine the normal functioning of the basal ganglia by examining neural activity in the striatum (also termed caudoputamen, the major input structure of the basal ganglia). The thesis begins with a review of literature related to our current understanding of striatal function, with a special emphasis on the anatomy of the basal ganglia and striatum, the behavioral correlates of striatal neurons, and the function of the striatum in learning and memory. Experiments are then presented which address current issues in striatal function, including the identification of projection neurons and interneurons in extracellular recordings from awake, behaving rats, the behavioral correlates of striatal neurons in navigation tasks, and the representation of task parameters by ensembles of striatal neurons. On the basis of extracellularly recorded spike trains, neurons in the rodent striatum could be differentiated into phasic and tonic subtypes, which are believed to correspond respectively to projection neurons and interneurons of the striatum. Tonic neurons could be further differentiated into 3 subtypes on the basis of firing properties and extracellular action potential parameters, and may each correspond to distinct striatal neural types. Phasic neurons which were responsive during navigation tasks were active either during navigation or during reward-receipt. Tonic neuron subtypes were also behaviorally modulated: two subtypes showed spatial oscillations as rats were running in a sequential navigation task, while the third subtype was only modulated following the presentation of a cue which signaled food delivery. Ensembles of striatal neurons provided high-quality representations of task parameters such as spatial location and reward-delivery. However, a strong representation of space was only obtained in a sequential navigation task, and not in a navigation task in which spatial location was ambiguously associated with reward-delivery. Also, in the sequential navigation the striatal spatial representation developed with a time-course that preceded the development of a stable route through the environment, suggesting that the striatum may participate in developing a stable stimulus-response strategy in navigation.

Contents

Acknowledgements	i
Abstract	iii
1 Introduction and overview of the thesis	1
2 The Basal Ganglia	5
2.1 Organization of the basal ganglia	5
2.1.1 Anatomy	5
2.1.2 Striatal neuron types	7
2.2 Behavioral correlates of striatal neurons	17
2.2.1 Tonic versus phasic neurons	19
2.2.2 Phasically active neurons (PANs)	20
2.2.3 Tonicly active neurons (TANs)	25
2.2.4 Changes in striatal activity	26
2.3 Striatal function	27
2.3.1 Learning and memory	27
2.3.2 Sequence learning	31
2.3.3 Reinforcement Learning and Basal Ganglia Function	32
3 Methods	37
3.1 Animals	37
3.2 Tasks	37
3.3 Surgery	39
3.4 Recording	40
3.5 Histology	41
3.6 Neural Data Analysis	41
3.7 Behavioral measures.	42
4 Classification of striatal neurons	47
4.1 Results	48
4.1.1 Data sets	48
4.1.2 Phasic/Tonic separation	48
4.1.3 Multiple subtypes of tonic neurons	49
4.1.4 Cell classification is stable between behavioral conditions	51
4.1.5 Spike timing relationships between cell types	52
4.1.6 Relationship of cell type classification to extracellular waveforms	53

4.2	Discussion	56
4.2.1	Cell type correspondence	57
5	The Multiple T Task	59
5.1	Introduction	59
5.2	Results	59
5.2.1	Behavior	59
5.2.2	Neurophysiology	61
5.2.3	Phasic-firing neurons	62
5.2.4	Tonic-firing neurons	69
5.2.5	Striatal representation of task parameters	73
5.2.6	Neural learning correlates	75
5.3	Discussion	77
5.3.1	Phasic firing neurons	78
5.3.2	Tonic firing neurons	80
5.3.3	Striatal representation	80
6	The Take 5 task	81
6.1	Introduction	81
6.2	Results	81
6.2.1	Behavior	81
6.2.2	Neurophysiology	82
6.2.3	Phasic-firing neurons	83
6.2.4	Tonic-firing neurons	83
6.2.5	Striatal representation of task parameters	85
6.3	Discussion	86
7	Conclusions	95
	Glossary	101

Chapter 1

Introduction and overview of the thesis

The basal ganglia constitute a large volume of subcortical tissue and include several distinct nuclei. Over the past two centuries, the basal ganglia have been the subject of a significant number of investigations, but today our understanding of basal ganglia function remains murky. In 1664, based on observations of paralysis in patients with basal ganglia degeneration, Thomas Willis proposed that the corpus striatum (including the caudate nucleus, putamen and globus pallidus) represented the origin of motor behavior (Finger, 1994). Although Willis' theory was not accepted without dispute, the corpus striatum was widely believed to be the origin of the motor tract controlling the movement of skeletal muscles until the discovery of the motor cortex in 1870. Experimental approaches in animals and clinical observations in humans led to the development by 1940 of a view of basal ganglia function in which the corpus striatum was seen to exert inhibitory control over the production of movements (Finger, 1994), a view which persists today in current models of basal ganglia anatomy and function.

In addition to motor function, the basal ganglia have been implicated in cognitive abilities. In particular, the striatum (also termed the caudoputamen, the major input structure to the basal ganglia) is critical for diverse functions including some forms of procedural memory or skill learning (Packard and Knowlton, 2002), behavioral flexibility (Ragozzino et al., 2002), sensory neglect (VanVleet et al., 2002) and sequence learning (Miyachi et al., 1997; Matsumoto et al., 1999). In spite of this functional diversity, the cellular architecture of the striatum is remarkably homogenous, composed primarily of projection neurons and several types of interneuron. On the basis of the the uniformity of striatal architecture, it seems probable that the striatum performs a uniform computational function, and that the range of behaviors that have been attributed to the basal ganglia reflect differences in the connectivity of striatal regions to other brain areas. An important, current view of the basal ganglia is that striatum receives input from several regions of the cortex, and in turn impacts one of the cortical areas which provided the original inputs (Alexander et al., 1986). Other contemporary views of basal ganglia function focus on basal ganglia outputs to brainstem and spinal motor regions (Swanson, 2000).

Key insights into the function of the basal ganglia have come from recordings of neural activity in awake, behaving animals. Many of these studies have focused on neural activity in the striatum, and have found that striatal neurons respond to a multitude of task parameters such as movements, spatial location, head direction, rewards, and stimuli which provide task-relevant cues. The diversity of reports from the primate and rodent literature suggest that for almost any task parameter that can be identified, some population of striatal neurons can be found which is responsive to that parameter. This observation leaves open the question of how the striatal representation is organized into useful behavior, and if there are any limits to the responsiveness of striatal neurons to task parameters.

Also, while neurons in the striatum are believed to change their responses during some types of learning (Carelli et al., 1997; Jog et al., 1999), and single striatal neurons have been shown to change their responsiveness during classical conditioning tasks (Kimura et al., 1984; Aosaki et al., 1994; Tremblay et al., 1998; Schultz et al., 2003), how changes in the behavioral correlates of striatal neurons relate to changes in performance of behavioral tasks is still unclear.

This thesis addresses several current issues in basal ganglia function. 1) Can the known types of striatal neurons (projection neurons, different types of interneurons) be differentiated on the basis of extracellular recordings? 2) How do striatal neurons contribute to the learning and performance of navigation tasks? 3) What type of information is represented by the striatum, and how does the striatal representation change as a function of experience? To answer these questions, experiments were conducted in which ensembles of neurons in the rodent striatum were recorded as animals performed navigation tasks. These experiments provided several important findings related to basal ganglia function. First, rats were able to learn to perform navigation tasks in which they mastered novel experimental conditions in a single behavioral session, indicating that the activity of single neurons in the rodent can be studied during learning of some navigation tasks. Second, on the basis of firing patterns and extracellular waveforms, striatal neurons were differentiated into several types which likely correspond to different striatal neuron categories (e.g. projection neurons, and various types of interneurons). Third, each striatal neuron type was modulated by some aspect of task performance (such as spatial location, reward delivery, etc.), and these behavioral correlates were strikingly different for each neuron type. Fourth, ensembles of striatal neurons represented task parameters such as position and temporal events, the striatal representation spatial location was task-dependent. Fifth, in a sequential navigation task the striatal representation of space developed as a function of the animal's experience on the task.

The thesis is organized into 7 chapters. Chapter 2 begins with a review of basal ganglia anatomy, with a special emphasis on the organization of the striatum. The types of neurons present in the striatum and the function of those neurons in the striatal network is a critical question for our developing understanding of basal ganglia function, and that issue is the primary focus of first section of the chapter. Firing patterns, connectivity, protein expression, and vulnerability to disease of striatal neurons are reviewed, with the goal of developing an understanding not only of the functioning of the striatal network, but how striatal neurons could be identified in awake, behaving animals. Section 2.2 continues with an examination of the behavioral correlates of striatal neurons in awake, behaving animals. Behavioral correlates of striatal neurons from both primates and rodents are reviewed, with a special emphasis on the function of putative striatal projection neurons and cholinergic interneurons. Section 2.3 reviews the role of the striatum in learning and memory. Lesions, inactivations and striatal dysfunction due to pathological states are considered with respect to the deficits that are produced on tasks of implicit and explicit learning and memory. Section 2.3 closes with a consideration of a current computational approach to striatal function based on the theory that the striatum is part of a larger brain network implementing a reinforcement learning algorithm.

Chapter 3 presents the collected set of methods used in the experiments described subsequently in Chapters 4 through 6.

Chapter 4 presents the a classification scheme for striatal neurons recorded extracellularly in the awake, behaving rodent. On the basis of firing patterns and extracellular action potential shapes, striatal neurons were differentiated into phasic and tonic categories, which are presumed to correspond respectively to projection neurons and interneurons. Interneurons were then subdivided into three types on the basis of firing patterns. One type of interneuron had properties consistent with

those reported for tonically active neurons (TANs) in the primate, and may correspond to cholinergic striatal neurons.

Chapters 5 & 6 use the proposed classification of striatal spike trains to identify the behavioral correlates of each type of striatal neuron during the performance of navigation tasks. Chapter 5 examines striatal neural activity in rats performing a sequential navigation task (the Multiple T task). In the Multiple T task, rats navigate through a set of sequentially organized T maze choices in order to receive food rewards. In the course of each experimental session, rats demonstrated two behavioral learning rates, a fast elimination of errors (incorrect turns) and the slow refinement of the path taken through the maze. Phasic neurons which were responsive during the Multiple T task could be separated into mutually-exclusive categories, those that were active as rats ran on the maze, and those that were active during the receipt of food rewards. Some tonic neurons demonstrated rapid, narrow spatial oscillations in firing rate as rats ran on the Multiple T maze, while other tonic neurons demonstrated slow, broad spatial oscillations in firing rate on the task. Presumed cholinergic interneurons were unresponsive in the Multiple T task. Ensembles of striatal neurons provided high-quality representations of task parameters, such as the spatial location of rats on the maze and the interval in which food rewards were obtained. The striatal representation of space also developed as a function of experience, with a rate that was faster than the development of a stable route (which was the slower behavioral learning rate that was observed in rats on the Multiple T maze).

Chapter 6 further examines the striatal representation of spatial and sequence parameters during navigation by running rats on a task (the Take 5) in which the spatial location of the rat is dissociated from the set of actions rats perform to receive food. In the Take 5 task, rats ran on a rectangular track for food rewards that could be delivered at four locations. On each trial, the location of the next food reward rotated with respect to the environment, such that the rewarded location was entirely predictable, but not stable in space. In the Take 5 task, rats were able to learn the rule predicting food delivery. Similar to the Multiple T task, phasic striatal neurons responded either during navigation or reward delivery, but not both. No type of tonic neuron demonstrated robust spatial tuning on the Take 5 task, although many tonic neurons were modulated during reward receipt. Also, presumed cholinergic interneurons demonstrated short-latency phasic responses to cues predicting food delivery. As was the case for the Multiple T maze, striatal ensembles recorded as rats ran the Take 5 task provided a high-quality representation of the interval in which rats obtained their food rewards. However, in the Take 5 task there was an absence of a high-quality representation of the location of the rat on the maze.

Finally, Chapter 7 draws together the results of these experiments and relates the findings described in Chapters 5 & 6 to issues of striatal function outlined in Chapter 2.

Chapter 2

The Basal Ganglia

2.1 Organization of the basal ganglia

The term basal ganglia refers to a collection of subcortical nuclei which classically include the caudoputamen (also termed the dorsal striatum) and the pallidum (i.e. globus pallidus and entopeduncular nucleus). Other basal ganglia systems are recognized, such as the ventral striatum (comprised of the nucleus accumbens, the olfactory tubercle and ventral region of the caudoputamen) and ventral pallidum (the substantia innominata, Heimer et al., 1995). Recently, medial and caudorostral extensions of the striatum and pallidum have also been proposed (Swanson, 2000).

The primary focus of this thesis is on the neural activity of the dorsal striatum, which receive prominent inputs from the cortex and projects to pallidal structures (the globus pallidus) as well as to the substantia nigra (Swanson, 2000).

2.1.1 Anatomy

This section will describe the known connectivity of the dorsal striatum, and models of dorsal striatal connectivity, which will guide the subsequent discussion of striatal function. The focus in this thesis is on the rat as an experimental subject, and all anatomy terms will refer to those commonly used in the rat. Where it is deemed relevant, data from primates will be introduced, and in these cases, the anatomical terms used in primates will be specified and related to proposed homologous structures in the rodent.

The connectivity of the dorsal striatum and closely related structures (globus pallidus, entopeduncular nucleus, substantia nigra and the subthalamic nucleus) are summarized in Figure 2.1. This schematic omits projections to brainstem regions such as the GABAergic projections originating in the substantia nigra pars reticulata (SNpr) and the entopeduncular nucleus which project to the superior colliculus and the periaqueductal gray. For a more complete description of these anatomical connections, the reader is referred to Swanson (2000). The dorsal striatum is the recipient of strong glutamatergic inputs from nearly the entire neocortex, as well as a dopaminergic input from the substantia nigra pars compacta (SNpc), and GABAergic inputs from the globus pallidus and the entopeduncular nucleus. The dorsal striatum in turn gives rise to a GABAergic projection to the globus pallidus and the entopeduncular nucleus, SNpr, SNpc as well as a weak projection back to the cortex (McGeorge and Faull, 1989; Swanson, 2000). The connections of these striatal targets (globus pallidus, entopeduncular nucleus, and the substantia nigra) are quite extensive (reviewed in Swanson, 2000). Among other targets, the globus pallidus projects to SNpc, SNpr, the subthalamic

nucleus, thalamus, entopeduncular nucleus, dorsal striatum and cortex. Among other targets, the entopeduncular nucleus projects to the globus pallidus, dorsal striatum, thalamus (including projections to the parafascicular nucleus of the intralaminar nuclei, and the ventral anterior-lateral complex) and cortex. The SNpc gives rise to the dopaminergic projection innervating the dorsal striatum, as well as the subthalamic nucleus, globus pallidus, entopeduncular nucleus, and other targets. The SNpr has many projections, which include thalamic targets, the SNpc, the entopeduncular nucleus and the dorsal striatum. Both SNpr and the entopeduncular nucleus project to motor regions of the thalamus (VA/VL), which in turn project to motor cortical areas.

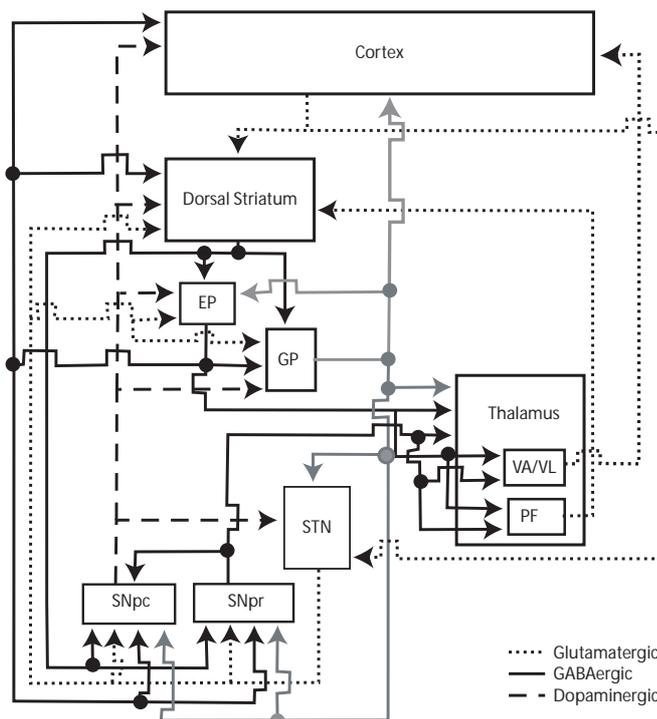


FIGURE 2.1: Connectivity of the dorsal striatum. Shown are the major inputs and outputs of the dorsal striatum, and the connections of the principal targets of striatal projections. Omitted are the extensive brainstem targets of the pallidum (EP and GP). Figure based on Swanson (2000). *Abbreviations:* EP, entopeduncular nucleus, GP: globus pallidus, PF: parafascicular nucleus of the intralaminar thalamic nuclei, SNpc: substantia nigra pars compacta, SNpr: substantia nigra pars reticulata, STN: subthalamic nucleus, VA/VL: ventral anterior, lateral nuclei of the thalamus. Solid lines indicate GABAergic projections, dotted lines indicate glutamatergic projections, and dashed lines indicate dopaminergic projections. Based on data reviewed by Smith et al. (1998); Swanson (2000)

Based on these anatomical connections, and a review of the literature delineating the motor functions and neural correlates of the basal ganglia, DeLong and Georgopolus (1981) proposed a model of basal ganglia function which emphasized cortico-basal ganglia-thalamocortical loops. In this view, the dorsal striatum was the recipient of convergent cortical inputs, and made projections to the pallidum, which projected to motor areas of the thalamus which in turn projected back to the cortex. The original model, as well as more recent treatments, have delineated the existence of multiple, segregated loops, which each receive input from multiple, overlapping cortical regions, and project back to distinct cortical regions (Alexander et al., 1986; Alexander and Crutcher, 1990a; Middleton and Strick, 2000). At least five such loops have been proposed, but the degree of segregation and convergence of information process in the basal ganglia remains a matter of debate (Parent and Hazrati, 1995).

Subsequent refinements of the model (for instance, Albin et al., 1989; Smith et al., 1998) have added additional complexity while retaining the fundamental assumption that the major functions of the basal ganglia are described by these cortico-basal ganglia-thalamocortical loops. Albin et al. (1989) incorporated the concept of indirect and direct pathways from the striatum to the thalamus (see Figure 2.2A). In this view, inhibitory outputs of the striatum project to the globus pallidus (or the external segment of the globus pallidus in the primate, GPe in Figure 2.2A) and the entopeduncular nucleus (or the internal segment of the globus pallidus in the primate, GPi in Figure 2.2A) and the SNpr. The entopeduncular pallidus and SNpr in turn project to motor areas of the thalamus (VA/VL), allowing striatal activity to accomplish a disinhibition of motor thalamic areas which in turn excites motor cortical areas. In contrast, the globus pallidus (GPe in primates) projects to the subthalamic nucleus, which in turn projects to the entopeduncular nucleus and the SNpr, allowing striatal activity to produce a net inhibition of motor areas of the thalamus and motor cortices (by stimulating the entopeduncular nucleus). The striatum-entopeduncular nucleus-thalamus-cortex pathway constitutes the direct pathway (with direct projections from the dorsal striatum to the entopeduncular nucleus), while striatum-globus pallidus-subthalamic nucleus-entopeduncular nucleus-thalamus-cortex constitutes the indirect pathway. Albin et al. (1989) further integrated the dopaminergic innervation of the dorsal striatum into this model to account for the dopaminergic involvement in hypokinetic movement disorders (such as Parkinson's Disease) and hyperkinetic disorders associated with excess DA levels (which can occur with treatments affecting DA levels, such as L-DOPA).

Another recent refinement to this anatomical model is the existence of a third, "hyperdirect" pathway from the cortex directly to the subthalamic nucleus (Nambu et al., 2002). In this model, glutamatergic inputs from the cortex excite subthalamic nucleus neurons, which in turn increase inhibition of the thalamus via excitation of the entopeduncular nucleus. The hyperdirect pathway is activated before the direct or indirect pathways, and allows for an immediate inhibition of motor activity at the level of the thalamus, before activity in the direct pathway is used to select an action to perform (Nambu et al., 2002).

A second perspective on basal ganglia anatomy emphasizes feedforward connections from the cortex and basal ganglia to brainstem motor regions, rather than cortico-basal ganglia-thalamocortical loops (see Figure 2.2B, Swanson, 2000). In this anatomical scheme, the cortex and basal ganglia impact behavior through projections to "behavioral control columns" which encompass motor control structures ranging from motor neurons in the spinal cord, to local circuits in the spinal cord which act as central pattern generators, to brainstem areas such as the hypothalamus which coordinate larger patterns of movements, including behaviors such as reproduction, food intake, and exploration. The basic organization of forebrain input to the behavioral control columns is an excitatory (glutamatergic) input from the cortex, which also projects to the striatum, an inhibitory (GABAergic) input from the striatum (which also projects to the pallidum), and a disinhibitory input from the pallidum (a GABAergic projection which is under striatal inhibition).

2.1.2 Striatal neuron types

As can be appreciated from the diagram presented in Figure 2.1, and in the reduced descriptions shown in Figure 2.2, the dorsal striatum is located at a critical node in information processing in the basal ganglia. A better understanding of basal ganglia function, then, depends on our understanding of striatal function: how do neurons in the striatum integrate glutamatergic, GABAergic, and dopaminergic inputs? How are neurons in the striatum organized with respect to these inputs, and what is the internal architecture of the striatal network? The present section addresses these

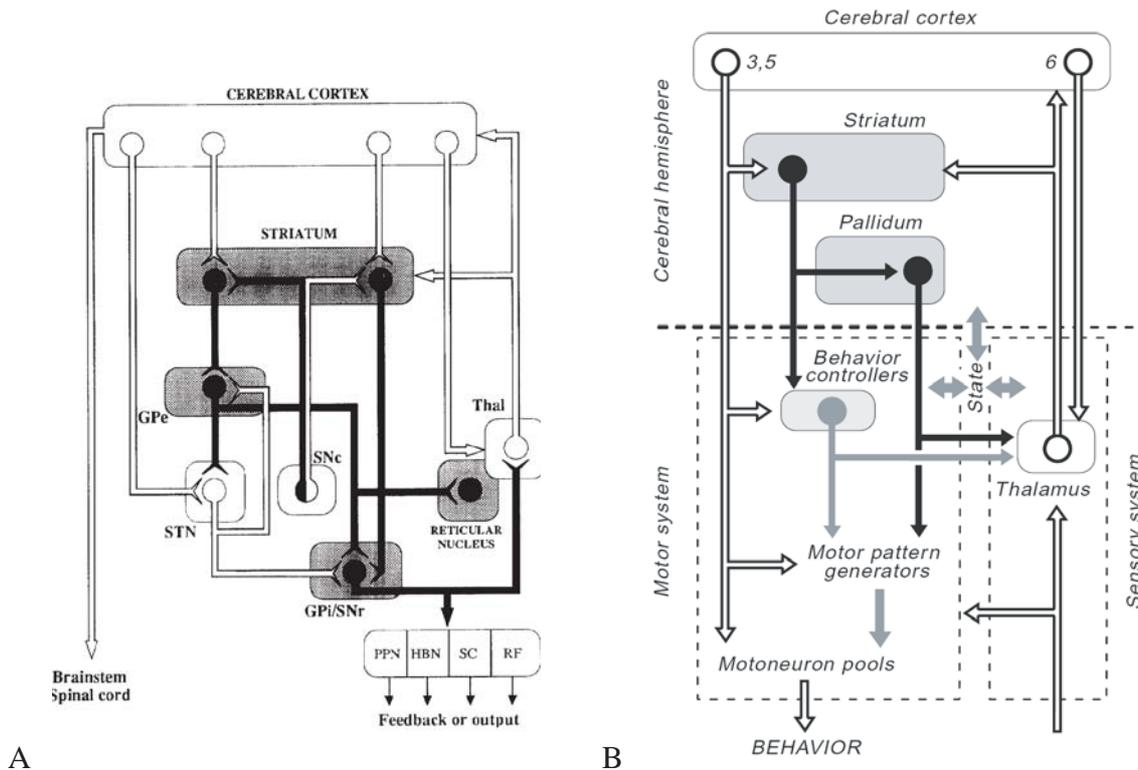


FIGURE 2.2: Models of basal ganglia information flow. *A:* An expanded model of cortico-basal ganglia-thalamocortical loops (diagram reproduced from Smith et al., 1998) originally proposed by Alexander and Crutcher (1990a) *B:* A recently proposed model of how basal ganglia and cortical projections converge on motor centers in the brainstem and spinal cord. Diagram reproduced from Swanson (2000).

questions by examining the fine structure of the dorsal striatum: the types of neurons of which the striatum is composed, the organization of those neurons with respect to one another and extrinsic inputs to the striatum.

Within the striatum, at least five distinct neuron types are known to exist, including projection neurons and four types of striatal interneurons (Kawaguchi et al., 1995). These distinct striatal neuron types can be identified on the basis of morphology, protein expression patterns, and physiology. The majority of striatal neurons are projection neurons (>95% of all striatal neurons in the rodent, Graveland and DiFiglia, 1985), which use GABA as their classical neurotransmitter. Projection neurons are medium-sized, and are further recognized by the high density of synaptic spines covering their dendrites (see Figure 2.3A & B). In this thesis, the striatal medium-sized spiny projection neurons will be referred to as MSPs. MSPs can be subdivided into those which express enkephalin and those which co-express substance P and dynorphin (Reiner and Anderson, 1990).

Of the four known types of striatal interneurons, three use GABA as their classical neurotransmitter (Kubota et al., 1993). These GABAergic striatal interneurons are also medium-sized, but do not have a high density of dendritic spines (see Figure 2.3 C – F). One of the GABAergic interneurons is immunopositive for the calcium-binding protein parvalbumin (PV+, Cote et al., 1991; Kita et al., 1990) and neurotensin-related hexapeptide LANT6 (Reiner and Anderson, 1993). A second GABAergic interneuron is immunopositive for the calcium binding protein calretinin (CR+, Figuerdo-Cardenas et al., 1996; Parent et al., 1995). The third GABAergic interneuron is immunoreactive for somatostatin (SS+), nitric-oxide synthase (NOS+), neuropeptide Y (NPY+), and has nicotinamide adenine dinucleotide phosphate diaphorase activity (NADPH-d+, (Kowall et al.,

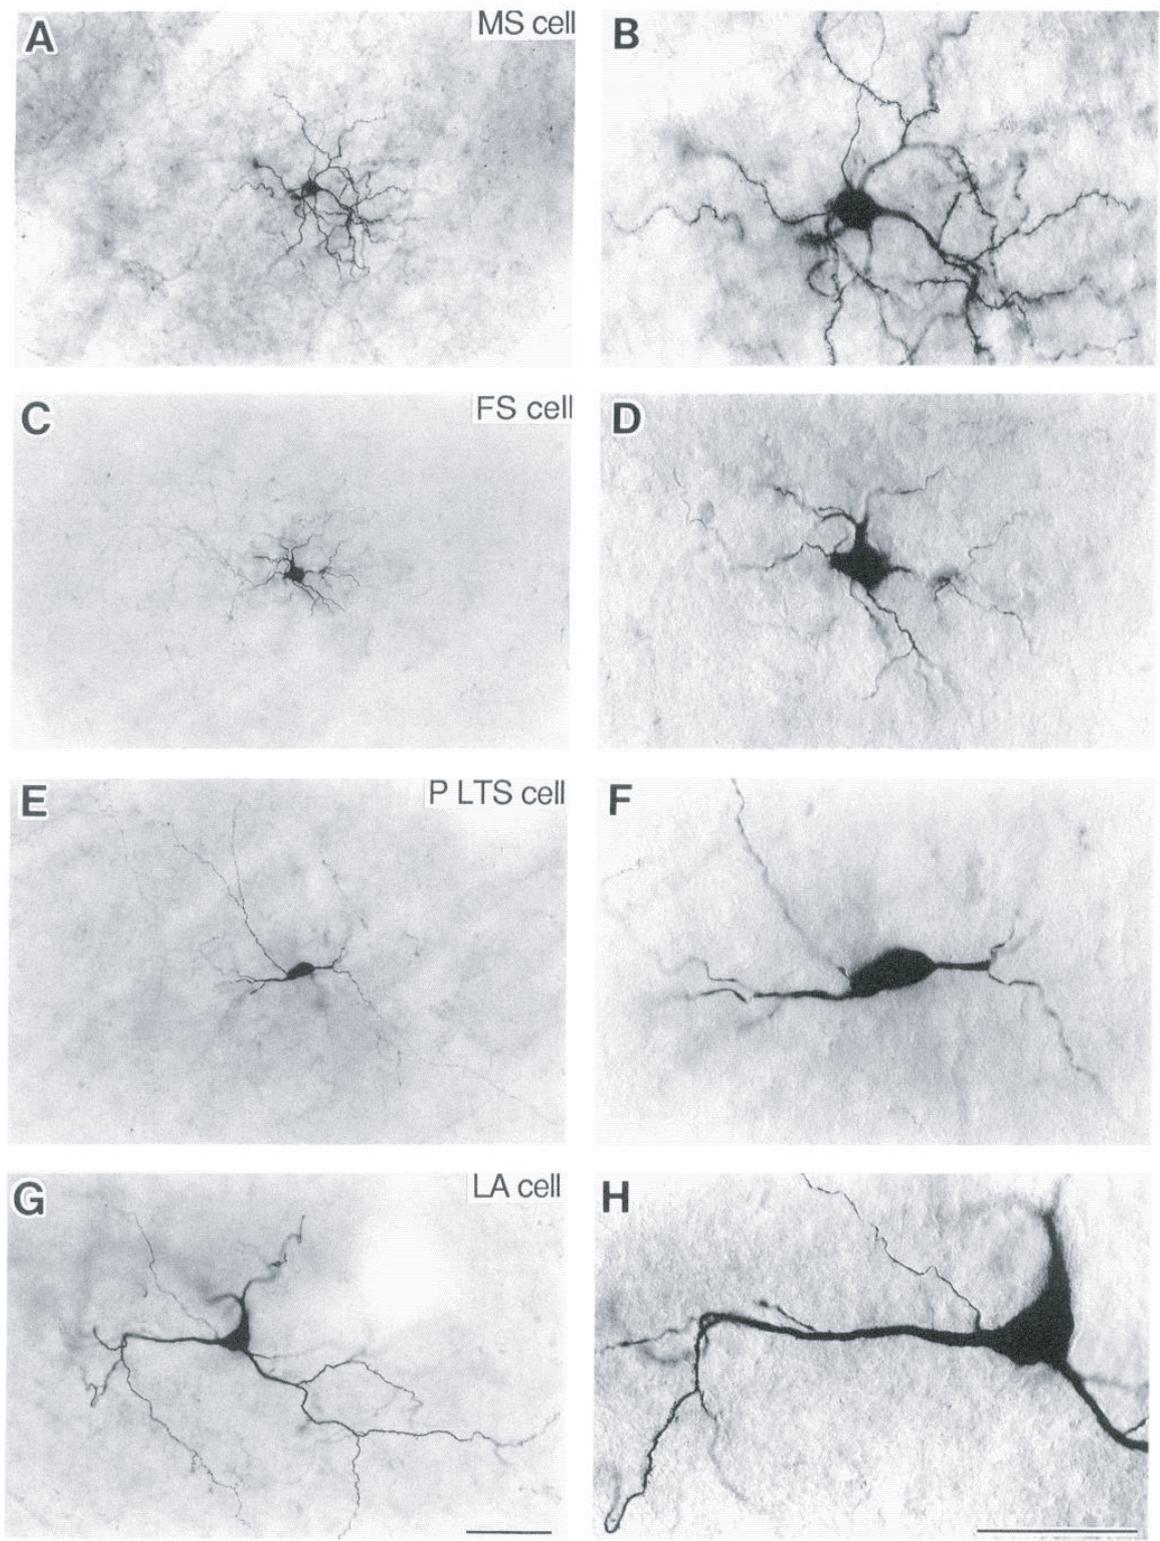


FIGURE 2.3: Photomicrographs of four classes of neostriatal cells intracellularly stained by biocytin. B, D, F and H are higher magnifications of A, C, E and G, respectively. A,B: An MS (medium spiny) cell. Note the dendrites with many spines. C,D: An FS (fast-spiking) cell. E,F: A PLTS (persistent and low-threshold spike) cell. Note the dendrites with fewer spines in FS and PLTS cell than MS cell. G,H: An LA (longer-duration afterhyperpolarization) cell. Note a larger soma than other cells. Scale bars: G, 100 μ m for A, C, E, and G; H, 50 μ m for B, D, F, and H. Figure and legend reproduced from Kawaguchi (1993).

1987; Selden et al., 1994; Smith and Parent, 1986; Figueredo-Cardenas et al., 1996)). The fourth striatal interneuron type uses acetylcholine as its classical neurotransmitter, and can be identified by expression of choline acetyltransferase (ChAT+ Selden et al., 1994). In contrast to other striatal neurons, ChAT+ neurons are large, and like other interneurons, lack prominent dendritic spines (see Figure 2.3A & B, Kawaguchi, 1993). In this thesis, these striatal interneurons will be commonly referred to as PV+, CR+, NOS+ or ChAT+ neurons. NOS+ neurons will also be referred to as NADPH-d+ neurons.

Striosome/matrix subcompartmentalization

On the basis of protein expression patterns, the dorsal striatum can be differentiated into two separate subcompartments, termed striosomes and matrix. Using a number of chemical and protein markers, striosomes can be visualized 2-dimensionally as patches, or 3-dimensionally as a tube-like network situated within a surrounding matrix (Groves et al., 1988). Striosomes are enriched in immunoreactivity for enkephalin and substance P (Graybiel et al., 1981), as well as μ -opiate receptors (Canales and Graybiel, 2000; Prensa and Parent, 2001). Matrix is enriched in staining for acetylcholinesterase (AChE, Graybiel and Moratalla, 1989), somatostatin-immunoreactive fibers (Gerfen, 1984), calbindin immunoreactivity (Gerfen et al., 1985; Cote et al., 1991) and is μ -opiate receptor-poor. (Canales and Graybiel, 2000; Prensa and Parent, 2001).

The distribution of striatal neuron types (their cell bodies and processes) are also organized with respect to striosome/matrix compartmental boundaries. The axons and dendrites of MSPs which are located in either striosomes or matrix respect striosome/matrix boundaries: MSPs located in matrix typically have axons and dendrites which do not cross into nearby striosomes and vice versa (Kawaguchi et al., 1989). Also, when a single injection of biocytin stains multiple neurons, all neurons are located in the same compartment, indicating that if MSPs are connected by gap junctions (as has been suggested by Onn and Grace, 1994), such electrical connections may also be restricted to other MSPs on the basis of striatal compartmentalization (Kawaguchi et al., 1989, and a similar argument is made by Onn and Grace, 1994, who have also suggested that administration of the DA agonist apomorphine can reveal dye coupling between striosome and matrix MSPs). A population of MSP with calbindin immunoreactivity is also localized to the striatal matrix (Cote et al., 1991; Gerfen et al., 1985).

Also, in humans NADPH-d+ neurons have been reported to be primarily located in the AChE rich matrix, and possessed dendrites which crossed into both subcompartments (Ferrante et al., 1987; Kowall et al., 1987). However, other researchers have reported that somatostatin-labelled (SS+) neurons (which are presumed to be NADPH-d+/NOS+ neurons) are present in both striosomes and matrix. In this study, the dendrites of SS+ neurons crossed into both compartments, but the axons of SS+ neurons were presumed to ramify specifically in the matrix (based on the distribution of somatostatin immunoreactivity, which is enriched in the matrix, Gerfen, 1984), arguing that SS+ neurons could mediate information transfer between compartments (Gerfen, 1984). While a patchy distribution of neuropeptide Y immunoreactive fibers (which co-localized to a large extent with somatostatin immunoreactivity) has also been shown in the cat, NPY immunoreactivity is homogeneously distributed in the basal ganglia of squirrel monkeys (Smith and Parent, 1986). Studies have indicated that large aspiny striatal neurons (ChAT+ interneurons) obey striatal striosome/matrix boundaries: The cell bodies and axons of large aspiny neurons lie outside of striosomes, as defined by calbindin staining, while their dendrites cross over into both matrix and striosomes (Kawaguchi, 1992; Aosaki et al., 1995). Axons which do cross into striosomes do not show arborizations within the striosome,

but do give off fine processes and boutons after passing back into the matrix (Kawaguchi, 1992). In contrast, PV+ neurons do not show a preference for striatal compartments (Cowan et al., 1990).

Striosomes and matrix have different connectivity patterns, receiving different cortical inputs and projecting to different nuclei. In rodents, Gerfen (1989) used the anterograde tracer, *Phaseolus vulgaris*-leucoagglutinin (PHA-L) to examine the topography of corticostriatal projections. A general finding was that deep cortical injections (deep layer V and later VI) tended to label striosomes, while superficial layer V and layers III and II tended to label matrix. Also, though there are cortical areas that tend to label striosome or matrix specifically, striosomes that are innervated by one cortical area are surrounded by matrix which also is innervated by the same cortical area. Selectivity for striosome or matrix may be determined by the emphasis of specific cortical layers in areas like the agranular cortices. The densest striosome input came from IL and PL injections, moderate striosome input came from anterior cingulate cortex, and sparse inputs from medial and lateral agranular cortices. Also, there was a rough topographic relation such that dorsal and ventral cortical areas projected to dorsal and ventral striatum.

Striosomes and matrix also differ in their connections to dopaminergic neurons (Gerfen, 1986). Striatal neurons in the matrix project to non-dopamine parts of the SNpr, while neurons in striosomes project to the locations of the cell bodies of dopaminergic neurons in the SNpc and the proximal dendrites of dopaminergic neurons of the SNpc which extend into the SNpr.

Physiological profiles

Intracellular recordings have been published for four of the five striatal neuron types (Kawaguchi, 1993; Kawaguchi et al., 1995). Each of these four types of striatal neuron has a distinct physiological profile, produced largely by intrinsic currents and synaptic inputs. MSPs are typically quiescent in culture (Plenz and Kitai, 1998), slice, and anesthetized preparations (Wilson and Groves, 1981). In relatively intact preparations, MSPs have bistable membrane potentials, and spend most of their time in a hyperpolarized “Down” state which is relatively far from the neuron’s action potential threshold. Periodically, MSPs transition to a depolarized “Up” state in which they are able to fire action potentials. These Up state transitions depend on coordinated synaptic inputs and active potassium conductances (Wilson and Kawaguchi, 1996). Corticostriatal neurons also have Up and Down states, and most have subthreshold oscillations (gamma frequency) in their Up states, and these frequencies are not observed in striatal cells (Stern et al., 1997), suggesting that cortical frequencies are not directly transmitted to striatal neurons.

Of the four striatal interneurons, physiological profiles generated from intracellular recordings have been characterized for three interneurons types (the PV+, NOS+, and ChAT+ interneurons), while the firing properties of the fourth type (CR+ interneurons) are unknown (Kawaguchi et al., 1995). Using whole cell recordings and intracellular staining for biotin and double staining or histochemistry for interneuron specific markers (ChAT, parvalbumin and NADPH diaphorase activity) Kawaguchi (1993) describes unique characteristics of each interneuron type. Physiologically, interneurons were divided into three types: fast-spiking cells (FS), persistent and low-threshold spike cells (PLTS) and long-lasting afterhyperpolarization cells (LA). FS cells were positive for parvalbumin (PV+), and all PV+ cells were GABAergic. PLTS cells had NADPH-diaphorase activity (NADPH-d+), and all NADPH-d+ cells were positive for NOS. LA cells were cholinergic (choline acetyltransferase immunoreactive, ChAT+).

Striatal neurons could be further differentiated on the basis of properties of the intracellularly recorded action potential. Using the duration of the action potential measured at half amplitude,

MSPs and PLTS/NADPH-d+ neurons had similar spike durations (~ 1 ms), while LA/ChAT+ neurons had shorter duration spikes (~ 0.8 ms) and FS/PV+ neurons has the shortest duration spikes (0.29 ms). LA/ChAT+ neurons had a very long time to peak of afterhyperpolarization (53 ms), which cleanly differentiated these neurons from PLTS/NADPH-d+ neurons and FS/PV+ neurons which had shorter times to peak AHP (14 ms and 1.3 ms respectively Kawaguchi, 1993).

Fast-spiking, parvalbumin immunoreactive interneurons (FS/PV+) At resting potential, synaptic stimulation of FS/PV+ neurons induced single spikes, at depolarized potentials a synaptic input induced trains of spikes. Depolarizing steps induced a slow depolarization followed by single spikes and long and short bursts of spikes at a constant frequency (20 - 100 Hz) (Kawaguchi, 1993). Bracci et al. (2003) have shown that the pattern of action potential firing in FS/PV+ bursts are highly variable in slices even when the same (step) currents are injected. Bursts of spikes in FS neurons did not show spike frequency adaptation, and between bursts of spikes the membrane potential of FS cells undergo high frequency oscillations. FS/PV+ interneurons are also modulated by dopamine: FS neurons are depolarized directly by dopamine in a D1 receptor dependent manner, and GABAergic inputs to FS neurons are inhibited by dopamine in a D2-dependent manner (Bracci et al., 2002). Endogenous catecholamines can have similar effects, as the application of cocaine (which will prevent the reuptake of catecholamines) will substitute for exogenous dopamine.

Koós and Tepper (2002) have also explored the cholinergic influence on FS/PV+ activity. In slices from rats aged 13-30 days, ACh had two effects on FS/PV+ cells: a direct depolarization via nondesensitizing soma-dendritic nicotinic receptors, and an attenuation of FS GABAergic inhibition of MSPs via presynaptic muscarinic receptors. The first mechanism may be relevant to the pauses in firing rate demonstrated by cholinergic striatal interneurons in primates in response to behaviorally relevant stimuli (see Section 2.2.3), as pauses in the firing of cholinergic striatal interneurons may impact FS/PV+ interneuron firing rates (Koós and Tepper, 2002).

Persistent and low-threshold spike cells, NADPH-d+ interneurons (PLTS/NADPH-d+) In PLTS/NADPH-d+ cells, depolarizing pulses could elicit fast spikes (Kawaguchi, 1993). At hyperpolarized potentials, depolarizations, cessation of hyperpolarization or synaptic inputs could elicit low-threshold single spikes in PLTS cells, and spikes with long-lasting depolarizations (Kawaguchi, 1993). In thalamic cells, low-threshold spikes are produced by transient Ca^{2+} currents and allow thalamic cells to enter a bursting mode (Gutierrez et al., 2001), and may indicate that PLTS/NADPH-d+ interneurons are capable of high firing-rate bursts. Striatal interneurons have also been identified which can fire low-threshold spikes, but do not demonstrate persistent, depolarizing plateau potentials (Koós and Tepper, 1999) which characterize PLTS cells (Kawaguchi, 1993). The LTS may be variants of the PLTS neurons (Koós and Tepper, 1999), and will be treated here as likely to be PLTS/NADPH-d+ striatal interneurons. LTS neurons are depolarized by the application of dopamine, and this DA-dependent depolarization is blocked by D₁-like, but not D₂-like antagonists (Centonze et al., 2002).

Long afterhyperpolarization, cholinergic interneurons (LA/ChAT+) LA/ChAT+ interneurons fire action potential which are characterized by long afterhyperpolarizations, are able to fire repetitively in response to depolarizing current steps, and show spike frequency adaptation (Kawaguchi, 1992, 1993). In response to increasing current injections, LA/ChAT+ firing rates (measured as the first interspike interval of the sequence of spikes) increase less than is seen in MSPs. Also, LA/ChAT+ firing rates show strong spike frequency adaptation, while little adaptation is seen in MSP firing rates

(firing rate measured using the second interspike interval, Kawaguchi, 1992). These data suggest that compared to other striatal cells, such as medium-spiny projection neurons, the striatal cholinergic interneurons are limited from firing at high firing rates.

In vitro, cholinergic interneurons recorded from rat slices (2 to 4 weeks postnatal) are spontaneously active, firing between 0-9 Hz (Bennett and Wilson, 1999). In the slice, the tonic firing of LA/ChAT+ neurons is not due to tonic excitatory inputs, and blockade of GABAergic, AMPA, NMDA, dopamine D1 and D2 type receptor as well as muscarinic cholinergic receptors all failed to significantly modify the spontaneously active cholinergic interneurons.

Synaptic organization of the striatum

Extrinsic inputs to the dorsal striatum As shown in Figure 2.2, the major inputs to the striatum are derived from several structures, including the cortex, the parafascicular and centromedian intralaminar nuclei of the thalamus, the pallidum, the subthalamic nucleus and the SNpc. The largest source of striatal afferents comes from the cortex, which projects in a roughly topographic fashion onto the striatum. Corticostriatal axons arise from two populations of cortical projections: from collaterals of brainstem projecting neurons which course through the striatum in the bundles of fibers of the internal capsule, and from cortical neurons projecting bilaterally to the striatum (Cowan and Wilson, 1994; Zheng and Wilson, 2002). The arborizations of corticostriatal axons within the striatum can be of a focal (1-4 0.5 mm arborizations separated by > 0.5 mm) or extended (> 1 mm) type. Both types of arborizations were found in brainstem projecting cortical neurons, and extended types found predominantly in bilaterally projecting corticostriatal axons (Cowan and Wilson, 1994; Zheng and Wilson, 2002). The density of synaptic contacts of the neurons with extended arborizations have been described as a matrix-like pattern combined with a baseline frequency of termination (Zheng and Wilson, 2002). As reported by Kincaid et al. (1998), corticostriatal projections to striosomes tended to be focal, while projections to matrix had both focal and extended terminations.

Medium-spiny projection neurons receive many of the corticostriatal inputs, which terminate densely on the synaptic spines decorating MSP dendrites (Kemp and Powell, 1971), but interneurons are also the targets of corticostriatal afferents (Bennett and Bolam, 1994; Lapper and Bolam, 1992; Lapper et al., 1992; Ramanathan et al., 2002; Vuillet et al., 1989). While synapses between corticostriatal axon terminals and cholinergic interneurons are absent in some studies using ChAT immunoreactivity (Lapper and Bolam, 1992), a small number of cortical axons terminate on synapses which are immunoreactive for m2 muscarinic receptors, which are specifically located in the striatum in cholinergic and NADPH-d+ interneurons (Thomas et al., 2000), and cortical axons are observed to terminate on the distal dendrites of ChAT+ neurons (Dimova et al., 1993). While sparse and positioned distal to the cell bodies of cholinergic neurons, electrical stimulation of cortical inputs to ChAT+ neurons produce post-synaptic depolarizations, and at shorter latencies than are seen in MSPs in response to contralateral corticostriatal stimulation (which preferentially stimulates corticostriatal axons derived from the bilaterally projecting cortical neurons, Reynolds and Wickens, 2004). PV+ interneurons also receive inputs from corticostriatal axons (Ramanathan et al., 2002).

Examining the innervation patterns of MSPs by corticostriatal axons, Kincaid et al. (1998) found little convergence of corticostriatal afferents on to single, or adjacent MSPs. The distribution of corticostriatal axons on MSPs was distributed exponentially, with a minimum separation of $\sim 4\mu\text{m}$, and an average distance of $10\mu\text{m}$. The authors estimate that a single axon made no more than 40 synapses within the dendritic tree of a single MSP, and that because 2840 MSPs were estimated to lie in the volume of one MSP dendritic tree, each corticostriatal axon contacting an MSP is likely to make synaptic contacts with no more than 1.4% of the MSPs sharing the volume of the dendritic tree.

Also, Kincaid et al. (1998) estimate that there are 30.5 million asymmetric synapses in the volume of the MSP dendritic tree, and if half of these synapses are cortical there are 380,000 different corticostriatal axons innervating the volume of one MSP dendritic tree. Even striatal neurons whose dendritic trees are restricted to the same volume will thus be unlikely to share a cortical input.

Corticostriatal axons also make synaptic contacts with PV+ interneurons, and single PV+ neurons receive convergent inputs from different cortical areas (primary motor and somatosensory cortices, Ramanathan et al., 2002). The terminals of cortical axons often formed multiple varicosities on a single PV+ neuron, and sometimes continued on to synapse on another, unlabelled cell. Thus, while single corticostriatal axons make a negligible contribution to MSPs, single axons from different cortical areas are likely to have a large impact on single PV+ interneurons.

An important thalamostriatal projection arises from the centromedian/parafascicular intralaminar thalamic nuclei (Kemp and Powell, 1971; Smith et al., 2004). Similar to corticostriatal afferents, thalamostriatal afferents innervating MSPs contact dendritic spines (Kemp and Powell, 1971). In primates, the centromedian nucleus projects to the putamen posterior to the anterior commissure (Sidibé and Smith, 1999), while the parafascicular nucleus projects to the caudate and the rostral and ventral putamen (Smith et al., 2004). Thalamic input to the striatum terminates on projection neurons and all interneuron types with the exception of calretinin interneurons (Lapper and Bolam, 1992; Sidibé and Smith, 1999; Smith et al., 2004; Rudkin and Sadikot, 1999). However, some reports suggest that the thalamostriatal input may only weakly innervate some interneuron types, such as the PV+ interneurons. In rats, when the anterograde tracer BDA (biotinylated dextranamine) placed in the parafascicular nucleus, only 4% of PV+ dendritic spines were contacted by labelled synaptic boutons (Rudkin and Sadikot, 1999), and only 1.3% of labelled parafascicular synaptic boutons contacted PV+ dendrites in this study. In contrast, PV+ striatal interneurons in primates receive a large input from the centromedian intralaminar nucleus (42% of labelled terminals, Sidibé and Smith, 1999). These results may indicate that there are species differences in the degree of thalamostriatal innervation of PV+ interneurons (Smith et al., 2004), or may reflect differences in the selection of tissue for examination in each study.

Pallidostriatal afferents preferentially terminate on PV+ and NOS+ interneurons in rats (Bevan et al., 1998). After labelling neurons in the globus pallidus, between 19-66% of the striatal synapses of made by labelled pallidostriatal axons were made with PV+ neurons, and a smaller, more variable number of synapses were made with NOS+ striatal neurons (3-32%). Individual PV+ neurons received multiple contacts from single pallidostriatal axons. About 25% of the pallidal cells labelled projected to the striatum, and pallidostriatal axons were observed to make collateral projections to other structures such as the globus pallidus, subthalamic nucleus, SNpc, SNpr, and entopeduncular nucleus). Other evidence suggests that pallidostriatal inputs to MSPs may be minimal: synaptic events in intracellularly recorded MSPs are increased following glutamate application in the striatum, but not in the globus pallidus (Guzmán et al., 2003). However, this does not rule out a specific projection from the globus pallidus to a subpopulation of MSPs.

Striatal inputs from the SNpc provide the major dopaminergic input to the striatum, and nigrostriatal axons from the SNpc make synaptic contacts with MSPs, PV+ interneurons (Bennett and Bolam, 1994; Kita et al., 1990), neuropeptide Y-immunoreactive (NADPH-d+ interneurons, Vuillet et al., 1989), and ChAT+ interneurons (Dimova et al., 1993). Striatal DA release is inhibited by nitric oxide (NO) synthesis, indicating that the local NOS+ striatal interneurons play a role in regulating striatal dopamine levels (Silva et al., 2003).

Intrastriatal connectivity The axons of spiny projection neurons have axon collaterals that are either largely restricted to the area of the dendritic tree or extended over long distances, with the latter pattern only observed in projection neurons located in the matrix subcompartment (Kawaguchi et al., 1989). While the existence of extensive GABAergic MSP axon collaterals suggests that MSPs may experience pronounced lateral inhibition, experimental studies have had difficulty in demonstrating the presence of inhibitory interactions. Jaeger et al. (1994) found no evidence for lateral inhibition of MSPs by nearby MSPs, either by antidromic stimulation of MSPs using electrodes placed in the substantia nigra, or using dual intracellular recordings of nearby MSPs. The authors speculated that MSP-MSP synapses may not use GABA, but rather peptide receptors (substance P, enkephalin). Fujiyama et al. (2000) also provides evidence that a significant number of GABAergic synapses on MSPs and interneurons contain low densities of GABA receptors. However, recent studies have demonstrated GABAergic MSP-MSP synaptic interactions. For instance, Guzmán et al. (2003) report that GABAergic IPSCs can be observed using antidromic stimulation of MSP collaterals projecting to the globus pallidus. Dopamine modulates the GABAergic transmission between MSPs mediated by local axon collaterals via both D1 and D2 mechanisms, with D1 agonists facilitating and D2 agonists decreasing GABAergic IPSCs (Guzmán et al., 2003)

MSPs are under profound GABAergic inhibition, and the application of GABA_A receptor antagonists increase MSP firing rates increase by more than 300% (Nisenbaum and Berger, 1992). Some GABAergic inputs to MSPs are derived from other MSP axon collaterals (Guzmán et al., 2003), but MSPs also receive GABAergic synapses from PV+ interneurons (Bennett and Bolam, 1994; Kita et al., 1990), which are interconnected via gap junctions. Koós and Tepper (1999) describe how PV+ neurons could act as a syncytial network to provide most of the inhibition of MSPs. In slice recordings, they found strong inhibition of MSPs by FS (PV+) neurons and a new cell they describe as an LTS cell (which may be a variant of the PLTS/NADPH-d+ neurons described by Kawaguchi, 1993). Koós and Tepper (1999) saw evidence for electrotonic coupling of FS cells, and a lack of inhibitory feedback from MSPs to FS cells. From these experiments, it was estimated that each FS cell might contact 135-541 MSPs, and each MSP might receive contacts from 4-27 FS cells. Single spikes from FS and LTS cells could delay MSP action potentials, and bursts of FS spikes could profoundly delay MSP spikes.

Electrotonic coupling via gap junctions is able to produce short time-scale synchronization of neural firing (Bennett, 1999). Interneurons in other structures such as the cerebellum (Mann-Metzer and Yarom, 1999) and cortical fast-spiking and LTS interneurons (Galarreta and Hestrin, 1999; Gibson et al., 1999) are connected by gap junctions and are capable of firing synchronized action potentials. These properties suggest that networks of FS/PV+ striatal interneurons, with their strong cortical and thalamic inputs could function as a synchronized, feed-forward inhibitory input to MSPs.

Plenz and Kitai (1998) describes results from MSPs and FS cells recorded from cortex-striatum-substantia nigra organotypic cultures, where simultaneous intracellular recordings were made of MSP and FS neuron pairs. Up state transitions in MSPs were blocked by CNQX (glutamate antagonist), and the occurrence of MSP Up states was correlated with depolarized periods in FS neurons. Blackwell et al. (2003), also using the triple co-culture, estimate that MSPs and FS cells receive synaptic inputs at 10-40Hz in down states and ~800 Hz in up states. Up states lasted for 300 ms for MSPs and 246 ms for FS cells. GABA_A currents were present in both up and down states with similar proportions of total synaptic inputs.

Chang and Kita (1992) explored the connectivity of cholinergic and parvalbumin immunoreactive (PV+) striatal neurons in rats. Cholinergic and PV+ neurons were observed to synapse on

PV+ neurons and putative projection neurons, but PV+ were not observed to synapse on cholinergic neurons.

Unilateral elimination of cholinergic interneurons using an immunotoxin induces an acute phase of contralateral rotations which dissipates, and a chronic contralateral turning in response to systemic administration of apomorphine (Kaneko et al., 2000). Coincident with the acute phase, there is an increase in substance P mRNA and decrease in enkephalin mRNA in the striatum, and in the chronic phase, there is a decrease in both D1 and D2 receptor binding. The results support opponent interactions between the cholinergic and dopaminergic system, which had been previously proposed. Elimination of both the NADPH-d+ and ChAT+ interneuron populations (which both express the neurokinin-1, NK-1, receptor) by administration of an excitotoxin conjugated with substance P (SP-PE35) induces no basal rotations, but ipsilateral rotations in response to DA agonists (Saka et al., 2002). After approximately 10 days, the response changes to a contralateral rotation. Striosome/matrix interactions are also affected, and in response to coactivation of D1 and D2 receptors acutely or chronically, the specific activation of striosome neurons was blurred, with more matrix neurons activated in lesioned animals. Normal animals have Fos expression in these conditions almost exclusively in dynorphin containing neurons, and in lesioned animals there is a recruitment of PV+ neurons (Saka et al., 2002).

Vulnerability to disease

Striatal projection neurons and interneurons are differentially vulnerable to neurodegenerative diseases such as Parkinson's, Huntington's and Alzheimer's Disease, and to injury. Projection neurons are specifically vulnerable in Huntington's Disease (HD), but can be differentiated into more or less vulnerable populations on the basis of their projections. In early and middle stages of HD, MSPs containing enkephalin and projecting to the external segment of the globus pallidus were more strongly affected than substance P containing projection neurons projecting internal segment of the globus pallidus (Reiner et al., 1988). In the substantia nigra, substance P projection neurons projecting to SNpr were lost to a greater degree than those projecting to SNpc. In later stages of the disease, all projections were lost except some substance P fibers remaining in the substantia nigra pars compacta (Reiner et al., 1988). Also in HD, NADPH+ neurons, medium-sized CR+ neurons and PV+ neurons are increased in density, but PV+ neurons are missing in advanced stages of the disease (Ferrante et al., 1987; Harrington and Kowall, 1991; Cicchetti and Parent, 1996). The increase in density of CR+ neurons is likely due to general striatal atrophy rather than the genesis of new CR+ neurons (Cicchetti and Parent, 1996). In HD, large CR+ neurons, which are predominantly cholinergic, are absent, but this result may not indicate the loss of cholinergic cells in HD, but rather a loss of the calretinin marker for the large, cholinergic cells (Cicchetti and Parent, 1996). For example, AChE+ neurons are preserved in HD, while ChAT levels are decreased, indicating that cholinergic neurons are likely still present in HD, but are not functioning normally (Ferrante et al., 1987). In Alzheimer's disease, ChAT+ neurons are lost in the ventral striatum (nucleus accumbens Selden et al., 1994), while NADPH-d+ interneurons are spared in both AD and Parkinson's disease (PD, Mufson and Brandabur, 1994; Selden et al., 1994), as well as Huntington's disease (Dawbarn et al., 1985; Ferrante et al., 1985, 1987). In models of ischemic injury in gerbils, projection neurons are severely reduced while cholinergic and SS+ interneurons are preserved (Chesselet et al., 1990).

Some neurons which are preserved in disease also show changes in the regulation of protein expression and morphology. In Parkinson's disease, NADPH-d neurons are preserved but show increased expression of neuropeptide Y mRNA (Cannizzaro et al., 2003), which may represent a

loss of dopaminergic neurotransmission or the effects of pharmacological treatments for PD (such as the administration of L-DOPA).

In Huntington's disease, there is a marked reduction of the size of AchE rich matrix, while AchE poor striosomes are unaffected (Ferrante et al., 1987)

2.2 Behavioral correlates of striatal neurons

While anatomical studies have shown the striatum to be a homogenous structure in terms of cell type constitution (e.g. section 2.1.1), disruption of striatal function has diverse effects, depending on the subregions targeted and the tasks used in assessment (e.g. section 2.3). In this situation, where a network of uniform structure has a critical involvement in different behavioral capabilities can suggest two possibilities. Either the deficits that have been observed to date on a variety of behavioral tasks can be organized under a single deficit (e.g. attention, perception of time, etc.), or regional specificity can be produced by region specific striatal inputs (from the cortex, thalamus, etc.), upon which the striatum performs a universal operation. In either case, a more clear picture of striatal function will only be possible with a better understanding of striatal information processing is developed. An invaluable tool in this respect comes from extracellular recording of neural activity in the awake, behaving animal. While such data are correlative in nature, when combined with the constraints imposed by our understanding of the striatal network connectivity, and the effects of striatal damage or inactivation, neural recordings in awake, behaving animals are of critical importance for our understanding of striatal information processing.

Given the intimate connections of the basal ganglia with cortical motor areas, a large amount of research in primates, and to some extent in rodents, has examined the relationship between striatal neural activity and movements. Other studies, motivated by the involvement of the striatum in more cognitive tasks (learning and memory, behavioral sequencing, behavior flexibility) have examined the plasticity of striatal responses, or the dependence of stimulus or motor neural correlates on higher-order variables such as sequence or behavioral context. As the striatum is the recipient of widespread inputs from the cortex and thalamic intralaminar nuclei, among other sources, striatal neurons have access to a diverse set of information. Accordingly, extracellular recordings of striatal neurons have revealed that populations of neurons can be identified whose firing rates are well correlated with a wide variety of task parameters, including motor activity, sensory cues, sequence information (described below in Sections 2.2.2 and 2.2.3). Similar results have been obtained in rats, where experiments in rodent navigation have demonstrated that dorsal striatal neurons respond to a variety of navigation-related parameters, such as the location, head direction or motor activity of rats performing navigation tasks (Wiener, 1993; Ragozzino et al., 2001; Yeshenko et al., 2004). Also, in instrumental tasks rodent striatal neurons are observed which respond to the presentation of sensory cues, the motor response animals perform, or both (White and Rebec, 1993). Rodent striatal neurons which respond to movements are sensitive to the context in which these movements occur. Striatal neurons in the rodent which are active during grooming movements occurring during stereotyped, sequenced grooming are often not active during the same grooming movement occurring outside of the context of sequenced grooming (Aldridge and Berridge, 1998).

However, in general past experiments using extracellular recordings in the dorsal striatum of awake, behaving rats have not made a distinction between projection neurons and interneurons, or between different types of interneurons. For the goal of understanding the information processing functions of the basal ganglia, such information is vital in order to understand what information is conveyed by the striatum to other brain structures as well as how information is processed locally

within the striatum. In this area, a good deal of progress has been made over the past two decades in primate experiments. In extracellular recordings made in the dorsal striatum of awake, behaving primates, a distinction has been made between the activity of phasic and tonic striatal neurons, which are believed to correspond to projection neurons and cholinergic interneurons, respectively (Kimura et al., 1990, 1996; Aosaki et al., 1995). In rodents, such distinctions have not been as common, but recent studies have begun to separate phasic units from presumed interneurons (see for instance, Redish et al., 2002; Berke et al., 2003; Pennartz et al., 2004; Daw, 2003; Schmitzer-Torbert and Redish, 2004). While cholinergic interneurons have been identified in primates, and some tonic neurons in the rodent have been hypothesized to correspond to striatal fast-spiking interneurons (Berke et al., 2003), at the present the full complement of striatal neurons have not been identified in any awake, behaving animals and studied in detail using extracellular recordings.

When different classes of striatal neuron have been identified on the basis of extracellular recordings, the behavioral correlates of striatal neurons are found to be highly dependent on cell type. In the primate, tonically active neurons (TANs, cholinergic interneurons) exhibit short-latency, short-duration pauses to unexpected rewards (see the example in Figure 2.4A), and to stimuli which are highly predictive of reward (Kimura et al., 1984; Ravel et al., 2001). In contrast, phasically active neurons (PANs, striatal projection neurons) are primarily quiescent, but respond at high rates (in some cases in excess of 50 Hz) under the appropriate behavioral conditions (see the example in Figure 2.4B). While TANs respond somewhat indiscriminately (they may respond to multiple stimuli, such as the delivery of reward, or to visual or auditory stimuli which predict reward), PANs are highly specific, sometimes responding to a unique combinations of stimuli and/or movements. The responses of TANs in some cases also depend on the motivational significance of the stimuli, as Ravel et al. (2003) have demonstrated that the responses of individual primate TANs discriminate appetitive stimuli (juice rewards) from aversive stimuli (aversive air puffs delivered to the face and loud auditory stimuli).

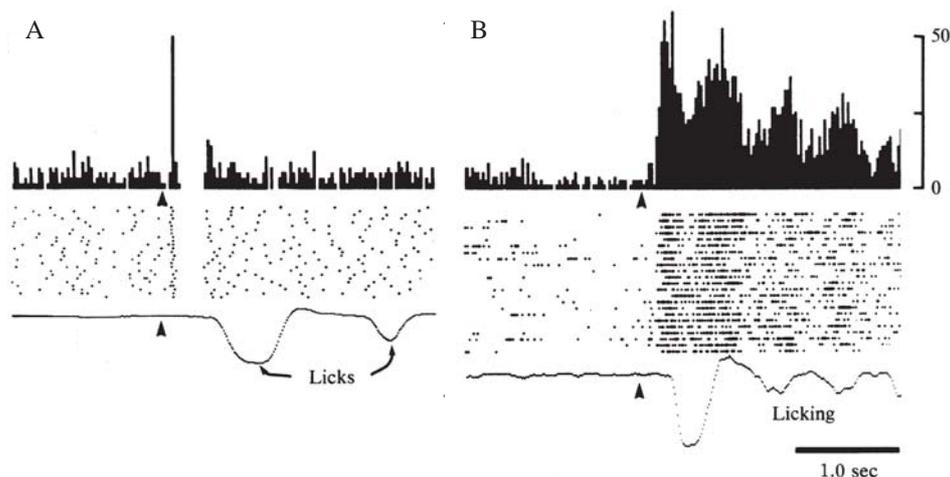


FIGURE 2.4: Example of a primate TAN (A) and PAN (B). Classically, TANs lack strong correlates to motor behavior, and typically pause in response to presentation of salient stimuli. PANs are often responsive to motor or sensory stimuli. Figure reproduced from Kimura et al. (1984).

The rest of the present section provides a review of extracellular recordings in the striatum of awake, behaving animals, focusing on the classification of striatal neurons on the basis of extracellular recordings, and the known correlates of striatal neurons, including the literature on PANs and TANs in primates, and the available neural data in rodents.

2.2.1 Tonic versus phasic neurons

As described in section 2.1.2, different types of striatal neurons can be identified on the basis of anatomical connectivity, physiology, striosome/matrix compartmentalization and vulnerability to disease and injury. In an examination of the behavioral correlates of striatal neurons with a goal of understanding striatal information processing, a critical issue extracellular recording is the identity of the neurons that are observed. Currently, no general classification scheme has been proposed in either primates or rodents which can identify the five known types of striatal neurons, but methods exist for the identification of MSPs and ChAT⁺ neurons in the primates, and some classifications have also been proposed for rat striatal neurons.

The recognition that striatal neurons in the primate could be separated into two types, phasic and tonic, on the basis of their baseline firing rates has been known for at least the last twenty years. Also, it has been clear that phasic and tonic striatal neurons have different behavioral correlates and connectivity to basal ganglia output nuclei. Alexander and DeLong (1985) described that striatal neurons recorded extracellularly in the primate putamen and caudate could be divided into Type I and Type II neurons, with Type I firing phasically (also referred to as phasically active neurons, or PANs) and making up 92% of their sample, while Type II units were active tonically (also referred to as tonically active neurons, or TANs) and had slightly different extracellular waveforms (see Figure 2.5). These proportions are comparable to the distribution of projection neurons and interneurons that would be expected on the basis of anatomical data, and Kimura et al. (1990) have shown that phasic neurons, but not tonic neurons, project to the globus pallidus, supporting the identification of phasic neurons as striatal projection neurons. Also, Aosaki et al. (1995) have provided evidence that TANs correspond to cholinergic striatal interneurons.

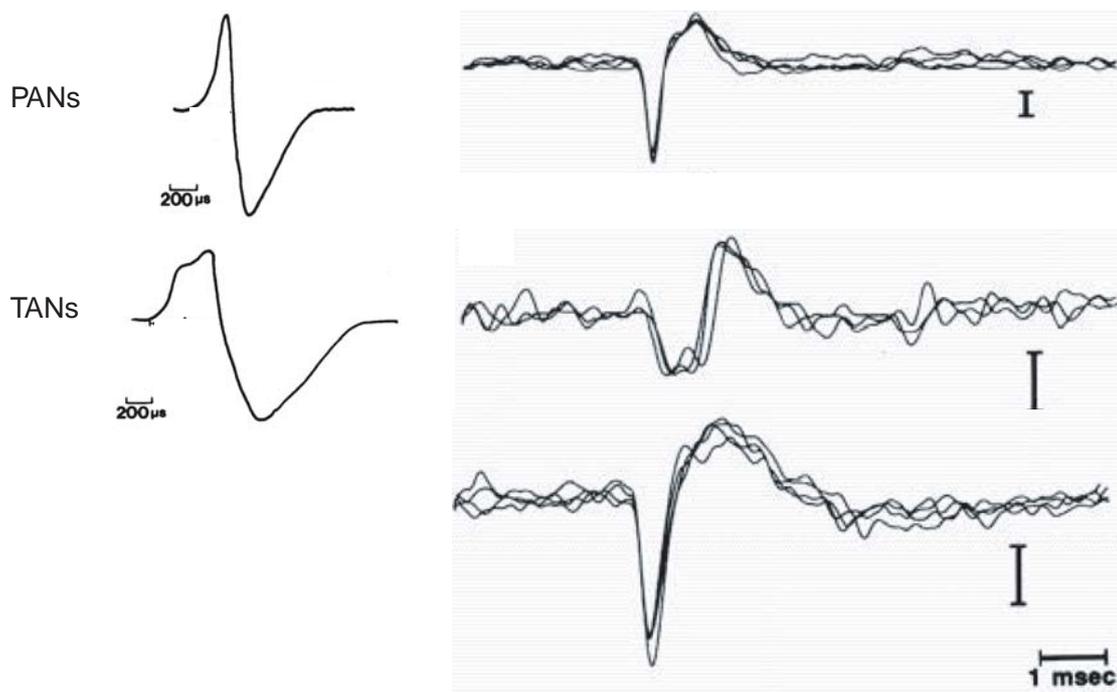


FIGURE 2.5: Examples of PANs and TANs from primates. Figure reproduced from Alexander and DeLong (1985); Hikosaka et al. (1989a)

In rats, projection neurons and interneurons are less frequently separated on the basis of firing patterns (but see Redish et al., 2002; Daw, 2003; Berke et al., 2004; Pennartz et al., 2004; Schmitzer-Torbert and Redish, 2004). On the basis of extracellular waveforms, two types of striatal neurons have been identified in the rodent which differ on the basis of their latency of response to corticostriatal stimulation (Nisenbaum et al., 1988), and response to the administration of acute or chronic haloperidol (Skirboll and Bunney, 1979). However, the correspondence of these neuron types to projection neuron versus interneuron types is not known. Type I striatal neurons have an extracellular waveform with an initial negativity, while Type II striatal neurons have an extracellular waveform with an initial positivity (Skirboll and Bunney, 1979). Extracellular waveform shapes are also stable over a wide range of bandpass filtering ranges and recording distances from the cell, but differences between these types were reduced with large high-pass filters (1 kHz), large distances between electrode tips and recorded neurons, and low impedance electrodes (1-2 M Ω , Nisenbaum et al., 1988). Type II units are most often “bursty”, while some Type I units fire both tonically or in bursty sequences (Skirboll and Bunney, 1979). Type I and Type II striatal neurons also have complementary responses to pair-pulse stimulation of the corticostriatal pathway. At short inter-pulse intervals (10-30 ms), Type I units have a facilitated response to the second pulse, while Type II units are inhibited. At longer intervals (50-240 ms), the relationship reverses, and Type I units are inhibited, while Type II units are facilitated (Nisenbaum et al., 1988). The relative prevalence of each type depends strongly on experimental conditions: Under halothane anesthesia, the majority of spontaneously active units in the striatum are Type I (86%, Skirboll and Bunney, 1979), while under chloral hydrate anesthesia, the majority of units which respond to corticostriatal stimulation are Type II (86%, Nisenbaum et al., 1988).

Type II units have been proposed to be medium-spiny projection neurons, and the short latency inhibition of Type II units has been proposed to be evidence of the inhibitory actions of local axon collaterals of MSPs (Nisenbaum et al., 1988). However, using intracellular recordings Onn et al. (1994) have shown that striatal neurons with both Type I and Type II pair-pulse relationships are predominantly medium-spiny neurons. The authors were not able to compare extracellular waveforms to the Type I and Type II shapes, so it is not clear if their Type I/II classification is entirely in register with those defined by extracellular waveform shapes. In Onn et al. (1994), when non-MSPs were encountered, they always demonstrated Type I paired-pulse profiles, and a total of 20% of Type I neurons recorded were not MSPs in this study. Thus, it is likely that in extracellular recordings, MSPs and non-MSPs may be differentiable on the basis of extracellular waveforms in the rodent, as is the case in primate extracellular recordings.

2.2.2 Phasically active neurons (PANs)

Putamen activity

Recordings in the putamen have revealed a large region in which the responses of phasic neurons are correlated with movements of the limbs and body, and an electrically excitable zone in which bodily movements can be produced by microstimulation. In the primate, Alexander and DeLong (1985) describe a detailed mapping of the responses of putamen and caudate neurons to passive and active sensorimotor stimulation and to microstimulation. There was dorsolateral to ventromedial somatotopic mapping, with the lower extremities represented in the dorsolateral putamen and the face represented in the ventralmedial putamen. Phasic neurons in the putamen were more responsive to sensorimotor stimulation than those in the caudate (50 versus 8 percent respectively were responsive to at least one type of sensorimotor stimulation, such as passive manipulation, active

movement, cutaneous stimulation, etc.), and while there was a wide region of the putamen in which microstimulation could elicit movements, none of the caudate stimulations produced movements. In regions of the putamen in which microstimulation elicited movements (which the authors termed “striatal microexcitable zones”) the body areas that were activated were generally in register with the response properties of neurons in the same locations to passive and active stimulation. Stimulation always produced movements of the contralateral body when the extremities were involved, and usually produced movement of the contralateral body when the face and axial muscles were involved. While there was a coarse somatotopic mapping in the putamen, there was no evidence of an orderly, fine-grained map (there was no consistent progression from arm to wrist to fingers in the striatal representation of the upper extremities). Phasic neurons also tended to be observed in spatially organized clusters, and neurons found in a cluster tended to respond to similar types of movements.

In rodents, the lateral aspect of the dorsal striatum is considered to be the homolog of the primate putamen, and neurons in the dorsolateral striatum exhibit neural correlates to movements (Carelli and West, 1991; Carelli et al., 1997; West et al., 1990) and to somatosensory stimulation (Carelli and West, 1991). Movement-related neurons in the dorsolateral striatum discharge during “whole-body movements,” or are related to movements of specific body parts, such as the whiskers, limbs, head, etc. (Carelli and West, 1991; West et al., 1990).

In behavioral tasks, primate putamen neurons are related not only to movement production (Crutcher and Alexander, 1990), but also to “cognitive” processes such as movement preparation (Alexander and Crutcher, 1990c) and goal locations (Alexander and Crutcher, 1990b). In these experiments (Crutcher and Alexander, 1990; Alexander and Crutcher, 1990c,b), monkeys were presented with visual targets (located either to the left or the right of the monkey) and were required to make an arm flexion or extension to acquire the targets. Movements were either made on the basis of cues, or on the basis of a memory of the goal location. Movements could also be either spatially congruent or incongruent (movements were made either toward or away from the stimuli), and movements were made under various levels of resistance loading to the manipulandum. Also, recordings were made in the arm areas of the supplementary motor area, motor cortex and putamen to examine the nature of neural coding at several levels of the corticobasal gangliathalamic loop related to motor activity.

Comparing neural activity under varying levels of resistance loading of the manipulandum, a higher percentage (52% vs 24%) of putamen neurons that were movement-related (i.e. cells that discharged significantly during some movement condition) were better related to the *direction* of movement (termed direction-specific neurons) than to *muscle activations* (termed muscle-like neurons, those neurons whose activity varied according to the amount of force required). In SMA and MC, approximately equal percentages (SMA: 38% vs 41%, MC: 41% vs 36%) of neurons were direction-specific versus muscle-like (Crutcher and Alexander, 1990). In each area, approximately 20% of the movement-related neurons were classified as “other”, indicating that their pattern of activity under varying loads was not simply explained by the amount of resistance applied to the manipulandum, or that the neuron did not differentiate between movement directions (flexion versus extension). The distribution of lead times of neural activity relative to movement initiation was largest for SMA, followed by MC than the putamen, suggesting that the movement related activity could be initiated in SMA, then propagate to MC and finally the putamen.

Comparing movements in which the direction of movement is known or unknown in advance, Alexander and Crutcher (1990c) describe that many neurons in SMA, MC and the putamen are specifically activated during the preparation period before an instructed, memory-guided movement.

Of the cells with significant activations during either preparation, movement or both, SMA had a higher percentage of preparation and preparation/movement neurons (55%) than MC (37%) or the putamen (33%). In each structure, >80% of the preparation related neurons did not vary their activity on the basis of the resistance applied to the manipulandum, and the majority of neurons in each structure (>77%) were direction-specific, in that their preparatory activity differentiated flexion and extension movements. As with movement-related neurons, preparatory activity in SMA tended to lead such activity in MC, which led activity in the putamen.

In order to dissociate movement-specific responses (selective for flexion or extension) from goal-specific responses (selective for targets appearing on the left versus the right of the monkey), Alexander and Crutcher (1990b) performed another condition of the task in which monkeys performed blocks of trials where the normal stimulus-response relationships were inverted. In the Inverted condition, a target to the right of the animal was reached by movements of the manipulandum to the left, and vice versa. By comparing Inversion trials with Normal trials, the dependence of neural coding on action versus target location was determined. Across both Normal and Inverted trials for preparatory and movement related neurons, *target-dependent* neurons were those whose activity depended on which location the target was presented at, while *limb-dependent* neurons were those whose activity depended on the movement required (flexion or extension) to capture the target. There were more limb-dependent than target-dependent preparatory related units in the SMA (40% vs 36%), while the opposite was true of the MC (15% vs 40%) and putamen (3% vs 12%). For movement related units, there were more limb-related units than target-related units in each area (60-70%). Importantly, a large proportion of putamen units (53% of the preparatory related units and 31% of the movement related units) were classified as “indeterminate” because although they discriminated between flexion and extension in Normal trials, they did not show discriminate firing in Inversion trials. Thus, a very common response for putamen neurons was that their responses were not well described as either target-dependent or limb-dependent.

Kimura and colleagues (Kimura, 1986, 1990; Kimura et al., 1990; Kimura, 1992; Kimura et al., 1996) have used a similar flexion/extension experiment in which monkeys made a series of three flexion-extension movements following the presentation of a visual target. Two types of movement-related phasic neuron were found: those that responded in the interval between the presentation of a sensory cue and the initiation of the first movement (sensorimotor PANs) and those that responded during each movement of the series (motor PANs, Kimura, 1986). Some PANs also responded to the presentation of the sensory stimulus (sensory PANs), and usually differentiated stimuli which were followed by movements from stimuli which were not followed by movements.

Of the movement-related PANs that were active following the sensory stimulus that initiated the first movement, the responses were better related to the time of the presentation of the sensory stimulus than the subsequent initiation of movement, but the responses of these neurons depended on the type of movement that was subsequently performed (flexion or extension), suggesting a role for these neurons in sensorimotor integration. The responses of the neurons depended on both the type of visual stimulus presented, and on whether or not the instructed movement was performed (Kimura, 1990). The response of these sensorimotor PANs was bimodal, with an early mode that appeared to be dependent on the delivery of the sensory stimulus, and a second mode that preceded the first movement of subsequent flexion or extension. When an additional visual stimulus was added that was presented before the trigger for movement initiation, the response of sensorimotor PANs became unimodal, losing the sensory response to the trigger, but keeping the activity preceding the first movement of the sequence.

The motor PANs demonstrated less specificity for visual cues and the type of movement sequence that was being performed. In the repetitive sequence of three flexion-extension movements, these neurons fired during each of the three movements, with the neural response well correlated with particular phases of the response (flexion or extension, Kimura, 1990).

Sensorimotor PANs had an earlier onset of activation than did motor PANs, which in turn were activated before the EMG of the major muscles involved in the flexion/extension response (biceps, triceps, etc., Kimura, 1990), but even the sensorimotor PANs did lead EMG by less than 100 ms.

Other experiments, using more complex tasks have supported the idea that PANs in the putamen have highly specific responses to the context within which stimuli are presented and movements are elicited. A recent study by Ueda and Kimura (2003) trained monkeys to make a sequence of 2 movements to either the monkey's left or right, and consisting of either pressing a button or turning a lever. Monkeys were presented with one sequence at random, and moved through it in a cued fashion, then repeated the sequence from memory. In the remembered condition, phasic neurons had a stronger relationship to the direction of movement, and also represented combinations of movements. Neurons which preferred the remembered condition responded preferentially to combinations of movements. Neurons that preferred the instructed condition or did not differentiate the two had low combination selectivity, and those that did not differentiate also had strong directional selectivity.

Complex stimulus-movement mapping have also been described by Boussaoud and Kermadi (1997), indicating that striatal neurons in the putamen were active after attention-directing cues and movement-instructing cues. In general, the responses of striatal neurons differentiated movements and cues according to the particular sequence of actions required (i.e. the same sensory stimulus or movement produced different responses in the majority of neurons based on what other stimuli and actions had just occurred). The authors also show that some of these stimulus-movement mappings were orderly: they show a unit which is responsive after an attention-directing cue (cues were presented in six locations arranged radially around a central fixation point), and the unit shows directional tuning. Another neuron, responsive after a movement-instruction cue, was only responsive for right movements cued by stimuli in the lower visual field. They also report a few cells with interesting properties, one of which fired in anticipation of cue offsets, and other neurons that anticipated go-signals.

Caudate activity

A second area of striatal neural responses that has been studied in detail comes from tasks in primates performing saccadic eye movements. Hikosaka and colleagues have shown that in simple, visually cued saccades, caudate neurons are responsive during saccades, to the presentation of visual and auditory stimuli, and in apparent preparation or expectation of task cues (Hikosaka et al., 1989a,b,c). In a sample of 2,559 caudate neurons, 867 (34%) were related to some task parameter (visual cues, auditory cues, saccades, memory, etc.), while an additional 502 (20%) caudate neurons were not responsive to any task parameter tested, but did respond outside of the task to some variable (visual, auditory, reward, motor activity, etc.). Neurons were tested in various versions of a saccade task in which monkeys made saccades to visual targets, or to remembered locations where previous targets had been presented.

Of the PANs which were saccade-related (306/867, 35% of task-related neurons), neural activity was specific to saccades directed to specific regions of the visual space (Hikosaka et al., 1989a). The "movement fields" of saccade related neurons were predominantly contralateral. This spatially restricted firing field was also modulated by whether the saccade was initiated in response to visual information (21%), memory for a previously presented target (33%), or both (26%). Some neurons

were also active as monkeys prepared to initiate a saccade (8%), demonstrating an increasing ramp in activity as the anticipated trigger stimulus for the saccade approached. Also, a sizable proportion of saccade-related neurons (13%) did not fall into any of these four categories (Hikosaka et al., 1989a). Saccade related neurons were highly specific for saccades within the context of task-performance: only 4 saccade related PANs were also active during the performance of spontaneous saccades made outside of the behavioral task.

Using a memory-guided saccade task, Hikosaka and colleagues have also probed the dependence of caudate activity on reward expectations. Monkeys were required to fixate while a visual target was presented in 1 of four locations (upper right, lower right, lower left, upper left, for example) and after a delay, was allowed to make a saccade to the correct location. Two conditions were tested: ADR, in which all four directions were rewarded, and 1DR, in which only one of the four directions was rewarded (or was assigned a higher reward relative to the other three locations: relative 1DR condition Kawagoe et al., 1998). In the 1DR condition, monkeys were still required to make a correct saccade to unrewarded targets in order to advance through the block of trials. Importantly, to keep the overall amount of reward constant over blocks of ADR and 1DR trials, the amount of reward delivered in 1DR trials was 4 times larger than that delivered in ADR trials. Kawagoe et al. (1998) used the ADR and 1DR tasks to study the responses of visual and memory related caudate neurons, and found that the responses of these neurons were strongly dependent on the reward condition. Some neurons maintained a similar spatial preference in the 1DR condition as the ADR condition, but decreased their activity when the preferred location was not rewarded. In other cases, neurons which had a direction specific response in the ADR condition remapped in each 1DR phase, following the rewarded direction and responding very little to unrewarded saccades. Somewhat more bizarrely, a few caudate neurons were also observed with anti-reward responses, only responding for saccades to unrewarded locations in 1DR trials. Greater than 80% of the visual and memory related caudate neurons demonstrated a reward dependence, and of these, 84% had enhanced activity for rewarded versus unrewarded saccades (i.e. reward-facilitated) while the remaining 16% had reduced activity for rewarded versus unrewarded saccades (i.e. reward-suppressed).

In the ADR and 1DR memory guided saccades, approximately half of the task-related caudate neurons are of the preparatory/anticipatory type, with activity preceding the presentation of the visual target. Remarkably, >90% of these anticipatory PANs are modulated by the location of the rewarded target: in each block of trials the activity of the anticipatory PANs depends on which locations are rewarded (Takikawa et al., 2002). Using variants of the task, Takikawa et al. (2002) demonstrated that the anticipatory activity was spatially dependent, and the preferred directions in the 1DR task were predominantly contralateral. In a visually guided saccade task (using only left and right targets), similar results are obtained (Lauwereyns et al., 2002). Changes in anticipatory activity in these neurons after switching the rewarded direction (from right to left, or vice versa) were tightly correlated with changes in the reaction times of the monkey (which were faster for rewarded than unrewarded saccades). Hikosaka and colleagues have proposed that the caudate activity may represent the motivational bias of the monkey to make saccades in one direction over another in tasks in which only one direction is rewarded. They envision the caudate working in feed-forward manner, taking inputs from the cortex, and influencing activity in the superior colliculus via projections to the substantia nigra pars reticulata, and may be the neural basis for the reward effects on saccade latency and velocity (Lauwereyns et al., 2002).

Other recordings made from the caudate nucleus of the primate have indicated that primate striatal neurons are modulated by information about the sequence of actions the animal is performing. Kermadi and colleagues (Kermadi et al., 1993; Kermadi and Joseph, 1995) studied the activity of

caudate neurons as monkeys performed visuomotor sequences (three different button presses) and found that neurons could be related to the performance of one or more arm movements, movements occurring in specific sequence positions (by responding when a button press was first in a sequence, but not second or third for instance), and could be sequence specific (by responding only to one movement in one sequence of movements).

2.2.3 Tonicly active neurons (TANs)

Tonicly active neurons (TANs) are presumed cholinergic interneurons possessing behavioral correlates very different from those of PANs (Aosaki et al., 1995; Kimura et al., 1984). TANs have baseline firing rates between 2-10Hz and have broad action potentials. Following the delivery of rewards or stimuli which are predictive of rewards, TANs respond with a brief pause in activity, which is less reliably flanked by brief periods of excitation. Kimura et al. (1984) gave the first description of the behavioral responses of these neurons. When monkeys were presented with juice rewards following a click, the monkey developed an association between the click and reward delivery. At the same time, TANs recorded in the monkeys developed a phasic inhibition in tonic firing that started 60 msec after the click, and lasted slightly longer than the ISI (see Figure 2.4A). The structure of the phasic inhibition of TAN firing was a single spike occurring 60 msec after the click, then another single spike after an ISI slightly longer than the TAN's average ISI. This microstructure produces the appearance of a triphasic activation pattern of successive activation-inhibition-activation. In other published reports, (Aosaki et al., 1994, 1995; Raz et al., 1996), TAN responses are not as reliable as those presented in Kimura et al. (1984), but possess a similar pattern of activation-inhibition-activation.

Aosaki et al. (1995) have demonstrated that tonicly active neurons are the large, aspiny, cholinergic neurons of the striatum, and that half of TANs were at the border between striosomes and matrix, suggesting that TANs may coordinate some information transfer between striatal subcompartments. TAN pause responses (268.3 ms) were longer than the mean ISI of TANs, suggesting that the TAN pause response is a suppression of TAN activity, rather than a change in spike timing per se. In a sample of 553 TANs recorded after monkeys had learned a CS-US relationship, 62% TANs responded to the presentation of the conditioned stimulus. Comparing visual and auditory stimuli, 2/3 of TANs were responsive to a visual CS, and 2/3 were responsive to an auditory CS (thus, 1/3 responded to both stimuli). For visual stimuli, caudate TANs had a faster onset than putamen TANs, but in general the responses were similar in both nuclei.

TAN responses are also dependent on the predictability of reward delivery. Ravel et al. (2001) found that TANs responded strongly to unpredictable rewards, responded weakly to rewards that immediately followed a trigger, and also responded weakly to rewards that followed a trigger presentation with a fixed delay. Shimo and Hikosaka (2001) examined Japanese monkey TANs and observed that 66.2% were responsive to rewards delivered unpredictably, while 63.1% responded to a visual stimulus presented 500 msec before a reward, and had no response to reward delivery itself. They also tested TANs on the ADR (All Directions Rewarded) and 1DR (1 Direction Rewarded) memory-guided saccade task (described above in Section 2.2.2). Many TANs showed responses to fixation point and cue presentation, and little response to reward. Only the cue presentation response was examined in depth. Less response was seen to cue presentation by TANs in ADR (25.2%) than 1DR (47.1%). In ADR, though, once you see fixation, the rest of the task, up to reward, is determined. The lack of certainty about the delivery of reward may reflect differences in cue response. Population responses show that TANs preferred to respond to contralateral directions in 1DR (60.3% showed

stronger responding to contralateral cues over ipsilateral cues, 5.5% preferred ipsilateral cues, and the remainder had no preference). In contrast, only 5.5% of TANs showed a preference for rewarded directions in 1DR.

In normal monkeys, TANs fire synchronously in a broad window of ± 50 -100ms, and are uncorrelated with basal ganglia output neurons in the globus pallidus (GP, Raz et al., 1996, 2001). In animals depleted of dopamine (by injection of MPTP, 1-methyl-4-phenyl-1,2,3,6-tetrahydropyridine), simultaneously recorded TANs show highly synchronized 15 Hz oscillations (Raz et al., 1996), and TAN-GP cell pairs also develop synchronized oscillations at 10 or 15 Hz (Raz et al., 2001).

TAN responses to reward-predictive stimuli depends critically on excitatory inputs from the thalamus (Matsumoto et al., 2001), and dopaminergic neurotransmission (Aosaki et al., 1994; Watanabe and Kimura, 1998). Inactivation of thalamic inputs from the centromedian and parafascicular intralaminar nuclei or reduction of dopaminergic transmission eliminates the conditioned pause and rebound excitation of TANs, but does not reduce the short-latency TAN excitations. Reduction of TAN responses parallels disruption of the conditioned behavior of animals, indicating a correlation of TAN responses with behavior.

In the rodent, no comparable demonstrations of TAN responses to behavioral stimuli have been presented. In most studies, neurons are not differentiated into phasic and tonic types (for recent exceptions, see Redish et al., 2002; Daw, 2003; Berke et al., 2004; Schmitzer-Torbert and Redish, 2004). The closest demonstration of a TAN-like pause in response to behaviorally relevant cues is presented in Gardiner and Kitai (1992), in which a TAN-like pause is shown in response to an auditory cue (see Figure 4B, Gardiner and Kitai, 1992). No systematic study of rodent TANs exists, but it is not clear if the lack of data represents a species difference in the behavioral correlates of cholinergic interneurons, or in the lack of appropriate experimental studies. Evidence that cholinergic neurons in the behaving rodent undergo learning related plasticity has been demonstrated indirectly in cases such as Bonsi et al. (2003), who demonstrated an increase in spontaneous GABAergic synaptic currents in striatal cholinergic interneurons in slices collected from rats trained in a instrumental task. As rodents are the primary subject of *in vitro* studies of cholinergic interneurons of the striatum, it is very important for our understanding of the cellular mechanisms of TAN responses and their contribution to information processing that the rodent TAN behavioral correlates be studied in detail.

2.2.4 Changes in striatal activity

Neural activity in the striatum also undergoes modification as animals learn new tasks and associations (Aosaki et al., 1995; Carelli et al., 1997; Jog et al., 1999; Schultz et al., 2003; Tremblay et al., 1998)

Carelli et al. (1997) have shown that in the lateral dorsal striatum of rats, neurons which respond to forelimb movement initially are responsive to a lever press for food, but gradually show a decrease in responsiveness over several days of training. This decrease in responsiveness was task-specific: these neurons remained responsive to passive and active limb movement after losing their responses to the lever press. Simultaneous EMG recordings of the biceps and deltoid muscles during lever pressing suggested that the decrease in task-related firing occurred without changes in muscle activation, making it unlikely that a change in responsiveness was due to a change in the manner in which rats were performing the task.

Jog et al. (1999) have shown that rats trained on a cued single T maze task, in which the frequency of a tone indicates whether to turn left or right, neurons in the lateral dorsal striatum respond to task

events (start of the trial, tone presentation, turning, and end of the trial). However, over the course of training, more neurons are found that respond to the first and last task events (the beginning and end of the maze) and fewer neurons are found that respond to the turn, while little change is seen in the number of neurons that respond to the tone presentation.

These two studies (Carelli et al., 1997; Jog et al., 1999) suggest that neuronal activity in the rodent striatum undergoes remodelling during the acquisition of a behavior. However, because the comparisons in neural responses are made across days, it is not clear from these data the extent to which the responses of single rodent striatal cells undergo remodelling during task acquisition. Other studies, examining cells in the primate striatum, have addressed this type of question by examining learning which can occur within in a single recording session, which allows the responses of single neurons to be examined across task acquisition.

Tonically active neurons of the caudate and putamen, presumed to be striatal cholinergic interneurons, show the development of a response to conditioned cues which predicts the delivery of reward (Aosaki et al., 1995). As monkeys learn that an auditory (solenoid click) or visual (LED flash) conditioned stimulus predicts a juice reward, TANS show a short duration pause in their firing (~ 260 ms) (Aosaki et al., 1995). These pauses depend on tonic levels of striatal dopamine, and TANS fail to respond in the MPTP treated monkey, but will resume responding after administration of apomorphine (a dopamine receptor agonist) (Aosaki et al., 1994).

Using a visual go-nogo task with monkeys, Tremblay et al. (1998) have shown that a number of slow firing neurons in the anterior striatum (caudate, putamen and ventral striatum) are responsive to task events (reward delivery, instruction cue, trigger) and undergo changes in responsiveness when new stimuli are being learned. The pattern of changes in responsiveness was diverse, including transient increases or decreases in responsiveness to familiar items, and changes in responsiveness to the new items which were being learned. Neurons were observed which responded to all three task events (instruction, trigger, reward) throughout the anterior striatum (dorsal and ventral). Neurons were also observed which differentiated between rewarded and unrewarded movements, with responses to rewarded movements being more common.

2.3 Striatal function

The intimate relationship of the basal ganglia and movement was recognized by Thomas Willis in the 1600s (Finger, 1994), but for many years the basal ganglia were incorrectly believed to be the origin of motor activity. Today, current anatomical models of basal ganglia function (section 2.1.1, see Figure 2.2) and the movement symptoms characteristic of human diseases such as Parkinson's Disease and Huntington's Disease have supported a modulatory role for the basal ganglia in voluntary movement (Albin et al., 1989). In parallel, a recognition of the role of the basal ganglia in higher cognitive functions has been also been developing. One area in particular in which a great deal of progress has been made is in our understanding of the involvement of the dorsal striatum in learning and memory, and in the performance of highly trained sequences of behavior.

2.3.1 Learning and memory

Data from humans, nonhuman primates and rodents, among other species, have indicated that the dorsal striatum is important for the learning and performance of habitual, automatic, procedural, stimulus-response behaviors. The striatum is not required for the performance of flexible, voluntary, declarative, goal-directed behaviors, which depend on the integrity of the the hippocampus among

other areas. This section will focus on the involvement of the basal ganglia in the development of reflexive, stimulus–response behaviors, with an emphasis on the rodent. Where relevant, examples from the human and non-human primate literature will also be introduced.

Over the past 25 years, evidence has accumulated for the existence of separate neural systems supporting the learning of “declarative” (or explicit) and “habitual” (or implicit) information. For instance, human amnesiacs are impaired in learning new explicit or declarative information, but are nonetheless able to acquire new implicit or procedural skills, such as a mirror-reading skill (Cohen and Squire, 1980). Also, amnesiacs demonstrate intact procedural learning of a weather-prediction task, but simultaneously report that they do not remember being trained on the task (Knowlton et al., 1996). Patients with Parkinson’s disease show the opposite pattern of results: They are impaired in learning the weather-prediction task, but can remember being trained to perform the task. Patients with Parkinson’s disease or Huntington’s disease are also impaired in other tests of implicit memory, such as rotor-pursuit learning (Heindel et al., 1989).

Analogous dissociations between memory systems have been obtained in rodents in the areas of navigation and instrumental learning. In rodent navigation, the task of learning how to reach a goal located somewhere in the environment can in principle be solved in one of two ways. If the goal is always located in a stable position in the environment, rats could learn a *place*-strategy in which rats plan a route from their current location to the goal on the basis of their knowledge of the environment (i.e. on the basis of a mental map of the environment). On the other hand, if the goal is located always in the same location relative to the rat’s current location, rats could learn a *response*-strategy in which rats learn to perform a fixed set of movements or a chain of stimulus–response relationships that will bring them to the goal. Place-strategies are thus analogous to explicit, or declarative learning, in that they require an explicit representation of the environment in order to plan the appropriate trajectory to reach the animal’s goals. Response-strategies are an instance of reflexive, or habitual learning, in that they require only stimulus–response associations and are inflexible.

Under the appropriate conditions, rats can learn to navigate to a goal using either a place- or a response-strategy, as was demonstrated by Tolman et al. (1946). In this experiment, rats were trained to run to a goal on a *plus-maze* (a maze formed by two linear paths intersecting at right angles, see the schematic shown in Figure 2.6). On each trial, rats were placed in either the north or south arm of the plus maze, and were rewarded for entering either the west or east arm of the maze. Rats were placed in either a place-learning condition or a response-learning condition. In the place-learning condition, rats were rewarded for always choosing the same arm (west or east), irrespective of which arm they were started in. In the response-learning condition, rats were rewarded for always making the same response (a left or a right turn), irrespective of which arm they were started in. Tolman et al. (1946) found that both groups of rats were able to learn their tasks, but the place-learning group acquired the task at a faster rate than the response-learning group. The authors concluded that rats were capable of implementing both place- and response-strategies in navigation, but that place-learning was a more “primitive” or basic ability.

Within the last decade, Packard and McGaugh (1996) have demonstrated that individual rats use both place- and response-strategies, with the former used early in training, and the latter used after extensive training. By converting a plus-maze into a T by blocking one arm of the maze (for instance, the north arm of the maze, creating a T maze with east-west arms forming the choices, and the south arm forming the stem of the maze, see Figure 2.6) rats were trained to run from the south arm of the maze into one of the choice arms to obtain food rewards. Probe trials were also used in which the south arm of the maze was blocked off, and rats were placed in the north arm of the maze (see Figure 2.6). On probe trials, rats demonstrated the use of a place-strategy if they chose to enter the

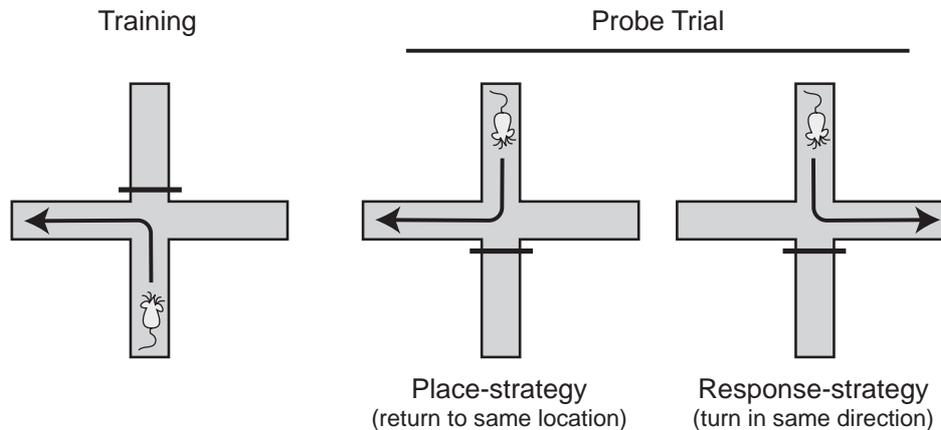


FIGURE 2.6: Schematic representation of the task used Packard and McGaugh (1996). A plus-maze was converted to a T by the use of a barrier blocking access to one arm. Rats were trained to run from the stem of the T to one of the arms for a food reward. After a period of training, rats were presented with a probe trial in which the arm that was blocked off and the arm rats were started in were interchanged. On probe trials, if rats entered the arm that was reinforced during training, they were classified as using a place-strategy. If rats entered the opposite arm (thus making the same *response* that was reinforced during training) they were classified as using a response-strategy.

same arm of the maze that was rewarded during training, whereas rats demonstrated the use of a response-strategy if they made the same response (turning in the same direction) that was required during training. In contrast to the experiment of Tolman et al. (1946), rats could use either place- or response-strategies during training in order to obtain food rewards, but the use of probe trials by (Packard and McGaugh, 1996) allowed for a direct test of which strategy rats were employing after learning the task.

In a probe trial conducted after one week of training, rats predominantly chose to enter the same arm that had been rewarded in training and thus were using a place-strategy. In a second probe trial conducted after a further week of training, rats predominantly entered the opposite arm by making a turn in the same direction that had been rewarded during training and thus were using a response-strategy. These behavioral data are consistent with the work of Tolman et al. (1946), and further indicate that when place-learning and response-learning are both consistent with task demands, place-strategies are typically acquired before response-strategies.

Packard and McGaugh (1996) explored the neural substrates of the place- and response-learning systems by inactivating the hippocampus or dorsolateral striatum with lidocaine before each of the probe trials. When the hippocampus was inactivated on the first probe trial (after one week of training, when rats typically used a place-strategy), rats chose randomly between the two arms. Dorsal striatal inactivations on the first probe trial had no effect on performance. When the hippocampus was inactivated on the second probe trial (after two weeks of training, when rats typically used a response-strategy), rats showed no impairments. However, dorsal striatal inactivations on the second probe trial caused rats to revert back to the earlier learned, place-strategy. These results indicated that the hippocampus was critical for the expression of the place-strategy, while the dorsolateral striatum was critical for the expression of the response-strategy. Further, after the striatal response system had taken over control of the behavior, the earlier learned hippocampal strategy was still functional: When the striatum was inactivated, rats immediately began to use the place-strategy on the second probe trial.

A number of subsequent studies have supported the hippocampal/place-strategy, dorsal striatal/response-strategy conceptualization of rodent navigation (reviewed by Packard and Knowlton, 2002; Poldrack

and Packard, 2003). Packard (1999) have found that infusions of glutamate made into the hippocampus after each training session biases rats to use place-strategies, while infusions of glutamate into the striatum biases rats to use response-strategies (using the same experimental paradigm as Packard and McGaugh, 1996). Using a water version of the plus-maze and training similar to that of Tolman et al. (1946) in which rats are trained either on a place-learning or response-learning task, delivery of the anesthetic bupivacaine into the hippocampus immediately following training sessions attenuates place-learning and enhances response-learning (Schroeder et al., 2002). Similarly, injections of lidocaine into the hippocampus before each training session delays place-learning and enhances response-learning (Chang and Gold, 2003a). Chang and Gold (2003b) have also used microdialysis techniques to demonstrate that in plus-maze training like that of Packard and McGaugh (1996), the trial at which rats switch from a place strategy to a response strategy is highly correlated with the ratio of baseline acetylcholine (ACh) release in the hippocampus to baseline ACh release in the dorsolateral striatum measured before training is initiated. During training, each structure (the hippocampus and dorsolateral striatum) demonstrates a different pattern of ACh release, with hippocampal ACh reaching plateau levels on the first day of training, and striatal ACh levels reaching plateau levels over the course of more than a week of training.

Within the striatum, place- and response-strategies may be represented respectively by the medial and lateral aspects of the dorsal striatum. The dorsomedial striatum is the target of inputs from structures in the hippocampal formation, including the entorhinal cortex (McGeorge and Faull, 1989), and would thus appear to be a likely component of the place-learning system. Consistent with this prediction, NMDA lesions of the posterior dorsomedial striatum (PDMS) significantly increase the use of a response-strategy by rats on the first probe trial in plus-maze training (Yin and Knowlton, 2004, using the paradigm of Packard and McGaugh, 1996). Dorsolateral striatal lesions increased the use of a place-strategy in both early and late probe trials, but the result was not significant. In this study, visual extramaze cues were minimized which may have led to a large number of rats (more than half) using response-strategies even on the first probe trial (Yin and Knowlton, 2004). In another experiment, when rats are trained in the Morris water maze concurrently to find either a visual or hidden platform and presented with a competition test, rats with dorsolateral striatal lesions navigate to the location of the hidden platform (a place-strategy) while rats with dorsomedial striatal lesions navigate to the visual platform (a response-strategy, Devan and White, 1999). Finally, when training rats to learn a reference memory version of the eight arm radial maze, NMDA antagonism in both the dorsomedial striatum and a dorsal posterior, more lateral, striatal location impairs acquisition of the task (Smith-Roe et al., 1999). After task acquisition, NMDA antagonism in the posterior striatum, but not the dorsomedial striatum, impairs performance (Smith-Roe et al., 1999).

The involvement of the dorsolateral striatum in response-learning and the hippocampus and the dorsomedial striatum in place-learning can also provide a general framework with which earlier research in rodent navigation is accounted for. For instance, lesions of the hippocampus or the fornix (which disrupts subcortical inputs and outputs of the hippocampus) disrupt performance on a win-shift version of the radial maze (in which each of 8 arms are baited once) but striatal lesions do not impair win-shift behavior (Packard et al., 1989; McDonald and White, 1993). Conversely, lesions of the striatum impair performance on a win-stay version of the 8-arm radial maze (in which 4 of the 8 arms were baited twice, and the baited arms are cued with lights). Also, rats with dorsolateral striatal lesions are impaired in learning to enter arms in the radial maze on the basis of associations with visual or tactile cues (McDonald and Hong, 2004), but not based on location (McDonald and White, 1995). These data are consistent with the place-strategy and response-strategy functions described for the hippocampus and dorsal striatum, in that the win-shift task requires a representation of the

environment and locations which have recently been visited (which is dependent on the integrity of the hippocampus, O'Keefe and Nadel, 1978; Redish, 1999). Stimulus-response associations, which are thought to underly striatally dependent response-strategies provide useful responding on the win-stay version of the radial maze, in which light-food associations are used to guide behavior. Similarly, Packard and McGaugh (1992) found that rats with dorsal striatal lesions were impaired in a visual discrimination version of the hidden platform Morris watermaze (in which one of two cues indicates where the platform is located), while fornix lesions impaired a spatial version of the task (where the same two visual cues were presented, but did not accurately predict the location of the hidden platform). Working memory impairments following hippocampal and striatal damage are also support these conclusions. In a delayed-matching-to-sample task testing spatial information, rats with hippocampal lesions are impaired, but rats with dorsal striatal lesions are not (Kesner et al., 1993). In a delayed-matching-to-sample task for egocentric responses, rats with dorsal striatal lesions are impaired but rats with hippocampal lesions are not (Kesner et al., 1993).

Within the animal learning literature, a related topic to place- and response-learning is found in the area of instrumental learning. When the delivery of a reward, such as food or water, is contingent on the behavior of an animal, such as the classical lever-press, rats quickly increase their rate of responding, thus obtaining the rewards. Such responding could have a basis in at least two representational schemes: goal-directed or stimulus-response. The goal-directed account would hold that rats have learned a relationship between lever-pressing and food delivery, and subsequent responding is driven by motivation to obtain the reward. Such responding is goal-directed, voluntary, and depends on an encoding of action-outcome relationships (A-O behavior). The stimulus-response account would hold that rats have associated the stimulus of the lever or the lever in the context of the testing apparatus, and by virtue of the rewards that are obtained after lever pressing, have developed a bias to press the lever when it is available, without any explicit encoding of the subsequent outcome. Such responding is reflexive, and depends on encoding the stimulus-response relationships (S-R behavior).

One of the key differences between A-O and S-R responding is the predicted sensitivity to the motivation of the animal to obtain the reward. Specifically, because A-O responding is guided by the value of the outcome, and S-R responding is guided by associations between the stimulus and response, the former but not the latter should be sensitive to changes in the value of the outcome (Dickinson, 1985). By devaluing the outcome (by allowing the animal to consume the outcome before testing, or by pairing the outcome with an aversive stimulus such as lithium chloride), A-O and S-R responding can be differentiated. Analogous to the results of Packard and McGaugh (1996), rats performing a lever pressing task for food show A-O responding after a moderate amount of training, and S-R responding after extensive training (Dickinson, 1985). S-R responding also appears to involve the striatal regions associated with place-strategies in navigation, as Yin et al. (2004) have demonstrated that lesions of the dorsolateral, but not dorsomedial, striatum impair the development of S-R habits.

2.3.2 Sequence learning

Beyond the tasks described above, in which animals were trained to perform a single response (make a single turn on a T maze, make a single barpress), the dorsal striatum is also involved in the production of sequences of actions. In rodents, sequenced grooming depends on an ~ 1 mm area in the anterior, dorsolateral striatum (Cromwell and Berridge, 1996). Lesions to this area do not block the ability of rats to perform individual grooming movements, but do disrupt the highly stereotyped

sequences that these grooming movements are normally observed to occur in. In contrast, lesions of the entire cortex and cerebellum do not abolish the ability of rats to perform stereotyped grooming sequences (Berridge and Whishaw, 1992).

In primates, Hikosaka and colleagues (Hikosaka et al., 2002; Miyachi et al., 1997; Rand et al., 1998, 2000) have shown that humans and monkeys can learn and retain visuomotor sequences and that in monkeys, learning and performance of these sequences depends on the dorsal striatum. In their task, subjects are presented with pairs of visual stimuli, and must select each item of the stimulus pair in the correct order. A single sequence is formed by 5 (for monkeys) or 10 (for humans) such stimulus pairs, presented sequentially. Humans and monkeys show a decrease in errors and time required to complete the sequence during acquisition of novel sequences (Hikosaka et al., 2002; Rand et al., 1998). Once learned, sequences are well retained for at least 18 months without intervening rehearsal (Hikosaka et al., 2002). Rather than memorizing the appropriate response for each stimuli pair (a simple stimulus-response strategy), monkeys appear to represent sequences: When stimulus pairs from a learned sequence were presented in the opposite order (preserving the stimulus-response relationships but perturbing the order within the sequence) the monkeys' performance was comparable to that for a novel sequence (Rand et al., 1998). Injections of muscimol, a GABA antagonist, into the anterior caudate or putamen inhibits learning of a new sequence, while injections into the middle and posterior putamen inhibited performance of learned sequences (Miyachi et al., 1997).

Using a different visuomotor sequencing task, Matsumoto et al. (1999) demonstrated that unilateral DA depletion by injection of MPTP (1-methyl-4-phenyl-1,2,3,6-tetrahydropyridine) directly into the caudate-putamen impairs learning and retrieval of previously learned sequences with the contralateral arm, but not the ipsilateral arm. If the sequence was learned prior to DA depletion, the monkey was able to relearn the sequence with the arm contralateral to the depleted caudate-putamen.

The dorsal striatum of primates and rodents is also involved in the production of repetitive, stereotyped movements following the administration of drugs of abuse, or co-administration of dopamine D1/D2 agonists (Canales and Graybiel, 2000; Saka et al., 2004). Increases in stereotypical behaviors are also produced by the chronic administration of drugs of abuse, and the production of stereotypical behaviors is well correlated with the specific activation of striatal projection neurons in striosomes relative to matrix neurons, and activation of the dorsal striatum, specifically the putamen (Saka et al., 2004). General activation in the caudate is also well-correlated with the stereotypies produced by drug free animals (Saka et al., 2004).

2.3.3 Reinforcement Learning and Basal Ganglia Function

While the data presented above indicate that the striatum is involved in the learning and production of habitual, sequential behavior, they do not address the question of *how* the basal ganglia in general, or the dorsal striatum in particular, accomplish such learning. Recently, however, the question of striatal function has increasingly been addressed from a computational framework in which the learning and behavioral functions of the striatum are well-described.

One theory of striatal function which has generated a considerable amount of research is that the striatum may be implementing a *reinforcement learning* algorithm (Daw, 2003; Schultz et al., 1995; Sutton and Barto, 1998). Reinforcement learning is a computational algorithm that focuses on an agent interacting with its environment, and addresses the question of how that agent should learn to interact with the environment in order to maximize the the amount of reward received in the future (Sutton and Barto, 1998). Agents implementing reinforcement learning algorithms possess a minimum of three elements: a *reward function* that specifies the amount of numerical reward that

is associated with each state of the environment, a *value function* which describes the amount of future reward that can be obtained from any state, and a *policy* that specifies how actions should be selected in each state of the world (Sutton and Barto, 1998). The value of a given state (the amount of reward that the agent can expect to receive in the future from that state) is often dependent on the policy that the agent is using, and so the agent is faced with two problems: First, for any given policy (any mapping of world states to actions), what is the value of each world state? and second, what is the policy that produces the maximal value function? The model considered below from the reinforcement learning literature tackles these questions simultaneously: it considers the situation in which agents are simultaneously estimating the value of a given policy and searching for the optimal policy in an online, trial-and-error fashion. It should also be mentioned at this point that some reinforcement learning systems incorporate a *model* of the environment which allows the agent to predict the outcomes of its actions, and to engage in planning rather than simple trial-and-error search (Sutton and Barto, 1998).

The situation faced by agents in reinforcement algorithms is a general description of the one presented to animals in many operant learning paradigms (Barto, 1995). The hypothesis that the learning strategies employed by animals may parallel the algorithms developed by reinforcement learning is most strongly supported by experiments in primates in which the activity of dopaminergic neurons has been recorded as animals perform various tasks. The work of Wolfram Schultz and colleagues (Fiorillo et al., 2003; Hollerman and Schultz, 1998; Ljungberg et al., 1992; Mirenowicz and Schultz, 1994; Schultz et al., 1995, 1997; Schultz, 1998; Tobler et al., 2003; Waelti et al., 2001) has demonstrated that the firing of primate dopaminergic neurons is highly consistent with the predictions of a specific type of reinforcement learning model: temporal difference reinforcement learning (TDRL). One specific TDRL model which has been proposed to account for the activity of dopamine neurons are *actor-critic* TDRL methods, in which separate components of the model track the current estimates of the policy (represented by the actor) and the value function (maintained by the critic Sutton and Barto, 1998). In any given time step, the amount of reward received (which can be referred to as *primary reinforcement* and is determined by the reward function) is compared to an estimate of the amount of reward that was expected for the state that the agent is currently at, and the state that immediately preceded the current state (which can be referred to as *secondary reinforcement* and is determined by the current estimate of the value function for each state (Barto, 1995). The difference between the the amount of primary and secondary reinforcement can be used to define a temporal-difference error term (TD error, also termed the *effective reinforcement*, Barto, 1995), denoted in most models by δ , which is used to update the value function and the policy (see Figure 2.7). Because the error term depends not only on the current state, but the immediately preceding one, δ allows the agent to associate high value to states which are not themselves rewarded but which lead to rewarded states. In an incremental fashion, this model (presented schematically in Figure 2.7) can learn long sequences of actions in which reward is delayed until sequence completion.

The interest in actor-critic TDRL models in explaining striatal function comes from recordings in primates which demonstrated that DA neurons fired in response to the delivery of rewards that were unexpected, and transferred their responses to cues which predict reward delivery (see Figure 2.8), and were inhibited at the time of omission of an expected reward. All of these properties are consistent with the properties of the δ signal proposed by actor-critic models, and suggest that the targets of DA neurons may be the synapses involved in the learning of value functions and policy estimation. The striatum receives a dense dopaminergic input, with DA neurons making synapses on the necks of dendritic spines of medium spiny GABAergic projection neurons, which also receive glutamateric synapses from corticostriatal axons (Sesack and Pickel, 1990). DA input to the striatum

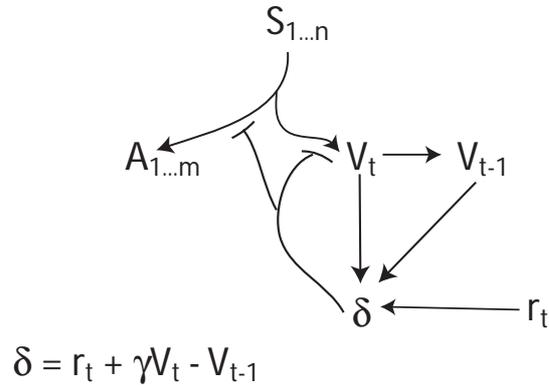


FIGURE 2.7: Schematic representing the actor-critic temporal difference reinforcement learning model structure. The model experiences a world defined by n world states, each of which is mapped to a single value unit (the critic), and m action units (together forming the actor). State transitions are accomplished by selecting one of the available actions, and updating the state of the world. The model learns by modification of the weights between world states and actions and worlds states to value using a comparison between experienced reward (r_t) and predicted reward (V_t) in each time step. The value of the state experienced at time t is compared with a memory of the value of the state experienced at time $t - 1$ and the reward experienced at time t to generate δ , the reward prediction error signal, which modifies the $S_t - A_{1...m}$ and $S_t - V$ model weights. An equation for δ is also given in the schematic.

is then well positioned to influence the mapping of cortical input, which could provide input representing the state of the world, to striatal responses, which in turn have access to the motor system through brainstem projections and to motor centers in the cortex via thalamic connections.

On the basis of these observations, recent proposals have suggested that the striatum may implement a temporal-difference reinforcement-learning algorithm (Barto, 1995; Sutton and Barto, 1998), in which striatal neurons select an action to perform based on a policy which is modified in order to maximize the receipt of reward over time (Brown and Sharp, 1995; Daw and Touretzky, 2000; Daw, 2003; Doya, 1999, 2000; Houk et al., 1995; Foster et al., 2000; Montague et al., 1996; Schultz et al., 1997). In reinforcement learning models of the striatum, the nigrostriatal dopaminergic system provides a reward-prediction error signal, and the striatum implements an actor-critic architecture. The actor is responsible for selecting which action would be appropriate given the current sensory input, while the critic uses the reward-prediction error signal to change the value of sensory inputs so that the most advantageous action will be chosen. The roles of actor and critic may be separated implemented by the matrix and striosome striatal subcompartments, respectively. Matrix receives inputs from sensorimotor cortex and projects to the substantia nigra pars reticulata and pallidal output nuclei (Gerfen, 1984, 1989; Kawaguchi et al., 1990; Ragsdale and Graybiel, 1984). Striosomes receive input from “limbic” cortex (including infralimbic, prelimbic, and anterior cingulate cortex) and project to dopaminergic cells in the substantia nigra pars compacta (Gerfen, 1984, 1989; Ragsdale and Graybiel, 1984). With its inputs to the substantia nigra pars compacta, striatal patches are well-suited to be involved in reward-related processing, while matrix is well-suited to be involved in action (Houk et al., 1995; Graybiel, 1998; Kimura, 1995; White, 1989). White and Hiroi (1998) have shown that electrodes placed in striosomes, but not matrix, will support self-stimulation in rats, supporting a relationship between striosomes and reward. Trytek et al. (1996) have shown that motor related neurons tended to be located in the matrix, supporting a relationship between matrix and action.

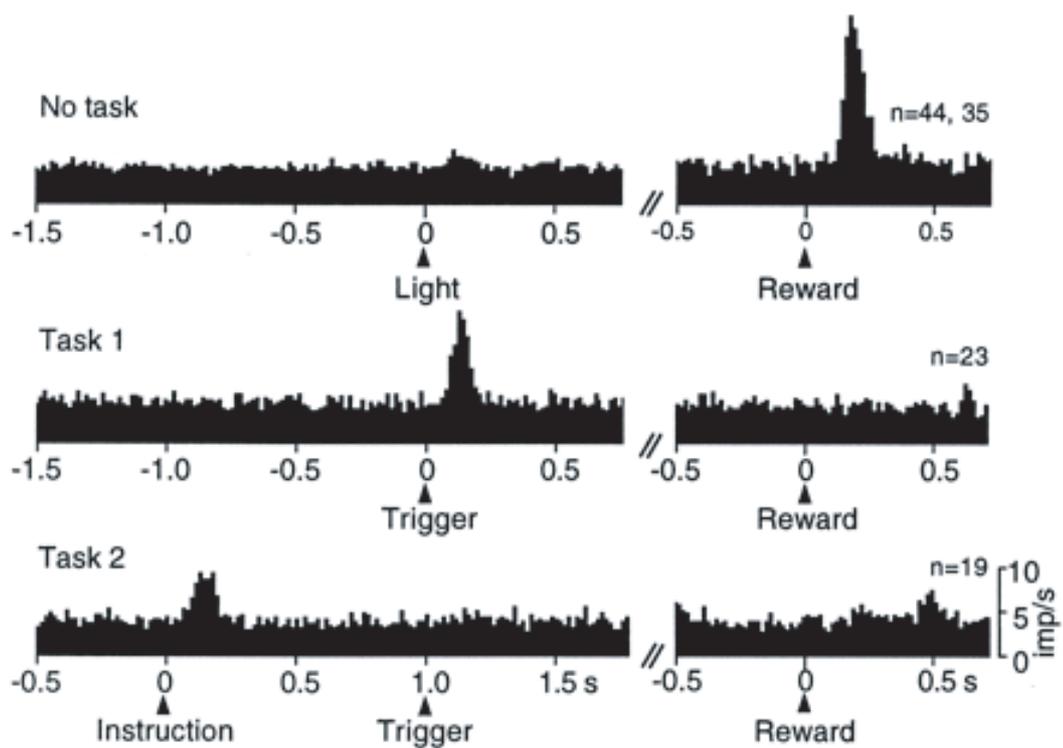


FIGURE 2.8: Population activity of dopaminergic neurons in the primate. Average population responses of dopaminergic neurons in primates in response to an unpredictable reward (A), a reward preceded by a sensory cue (B), and a reward preceded by two cues (C). In each case, the population response is a phasic burst of activity following the presentation of the earliest predictor of reward. When reward delivery is unpredictable, DA neurons fire phasically, when reward is predicted by earlier cues, this phasic activity moves to the earlier cues, and no response is seen to the delivery of reward itself. Figure reproduced from Schultz (1998).

Chapter 3

Methods

3.1 Animals

Data was collected from 21 rats, of which 10 were implanted chronically with hyperdrives located over the striatum. Rats were 13-15 months old at the time of experiments, and were food-restricted during behavioral training and testing. All experiments were conducted in accordance with NIH guidelines for animal care and approved by the IACUC at the University of Minnesota. Following each day's experiments, rats were handled for 15 minutes and fed additional food if required to keep the animal's weight at greater than 80% of its baseline weight.

Neural data from rats running the Nosepoke task was kindly provided by Jadin Jackson, and a number of sessions from rats running the Multiple T task were kindly provided by Pratibha Aia, Mallika Arudi, and Dan Bernal.

3.2 Tasks

Rest For 5 minutes immediately before and after each recording session, neural activity was recorded while rats rested in a terra cotta pot lined with towels.

Multiple T Rats were trained to run an elevated linear multiple T task which consisted of 3-5 T choices arranged sequentially to form a *turn sequence* (see Figure 3.1). On either side of the turn sequence, *return rails* led from the end of the maze back to the beginning, so that rats ran the maze as a continuous loop. On each return rail, two automatic food dispensers (Med-Associates, St. Albans VT) delivered pellets to locations on the track separated by ~ 45 cm. On completion of each trial the rat received two 45 mg pellets (Research Diets, New Brunswick, NJ) at each food delivery location, for a total of 4 pellets per trial. If the rat made an incorrect turn on the final T and ran back along the wrong return rail (thus passing the incorrect pair of pellet dispensers), no pellets were delivered, and the rat had to repeat the turn sequence in order to finish the trial and receive food. Throughout the task, rats were blocked from moving backwards on the maze, but were allowed to make incorrect choices. In practice, rats tended not to turn around and were rarely blocked.

The maze was constructed of plywood boards measuring 10 cm wide and covered with carpet. Each T consisted of a stem (40 cm long), and two choice arms (each 22 cm long) oriented at 90° to the stem. The return rails were 215 cm long, and were separated by two rails (140 cm long), located at either end of the return rails, which led from the turn sequence to the return rails (see Figure 3.1).

Take 5 Rats were trained to run on an elevated, rectangular track for food. The track measured 61 cm by 92 cm, and was created out of plywood boards measuring 15 cm wide, covered with carpet. On each side of the track, food pellets were delivered from automatic food dispensers (Med-Associates, St. Albans VT). On each trial, rats were rewarded for running 1.25 times around the track. Thus, to complete one trial, after receiving food rewards at one location (say the south side of the track as shown in Figure 3.2), rats ran clockwise around the track to return to the south side of the track, then advanced to the west side of the track to receive a food reward.

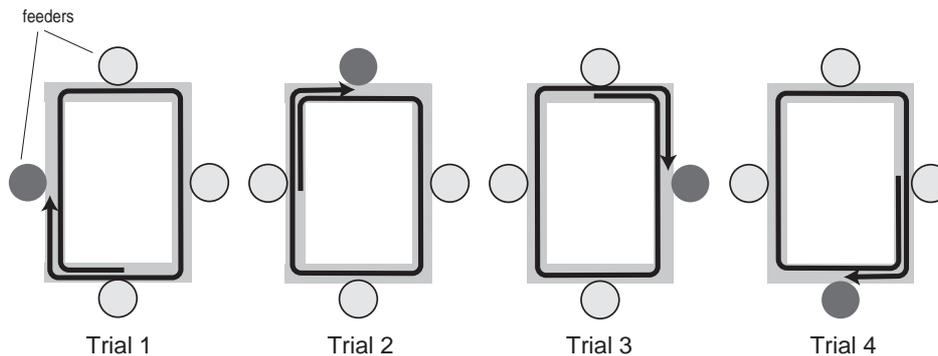


FIGURE 3.2: The Take 5 task schematic. Rats ran on a rectangular track for food which could be delivered to any side of the track. In each trial, rats were required to run 1.25 times around the track. Illustrated are four successive trials, in which the rat first received food reward on the west side of the track, then subsequently received food on the north side of the track. Thus, on each trial food was delivered in a new location, which rotated around the track. The path of the rat in each trial is shown by the black arrow, the rewarded pellet dispenser in each trial is shaded dark.

Four trials on the Take 5 task are shown schematically in Figure 3.2. On each trial, the next food delivery site was predictable based on a rule using the last rewarded food site, and how far the rat had run around the track. However, rewarded locations rotated on each trial clockwise with respect to the room. Such a procedure allowed the differentiation of behavior and neural activity that was related to *physical location* from activity related to the *sequence location* of the rat in the set of actions leading up to reward. The completion of a trial was signalled by a short (~ 100 ms) tone which immediately preceded the delivery of food. Occasional probe trials were also delivered pseudorandomly in which either the tone, food delivery or both were omitted. There was 1 probe trial in each block of 12 trials, no probe trials delivered in the first 4 trials in any block.

Nosepoke Rats were trained to make an operant response (nose-poke) in order to receive food rewards. A nosepoke port containing an infrared beam and a pellet dispenser were placed at opposite ends of an elevated linear track (see Figure 3.3). The track was constructed from a plywood board measuring 137 cm by 15 cm and covered with carpet. Rats were required to interrupt the infrared beam for ~ 100 ms in order to receive a short tone signalling the availability of food at the pellet dispenser. When rats travelled to the opposite side of the track, food pellets were delivered.

3.3 Surgery

After pretraining on one of the behavioral tasks (Multiple T, Nosepoke, or Take 5), rats were implanted with 14-tetrode hyperdrives (David Kopf Instruments, Tujunga, CA) targeting the striatum. Twelve tetrodes were used to record neural activity, and two electrodes were used as references for

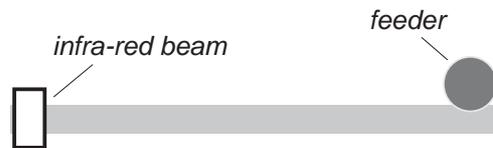


FIGURE 3.3: The Nosepoke task schematic. A nosepoke port and a pellet dispenser were placed at opposite ends of a linear track. Rats were required to interrupt the infrared beam in the nosepoke port for ~ 100 ms in order to activate the pellet dispenser.

common noise rejection. Tetrodes were constructed from four lengths of 0.0127 mm wire insulated with polyamide (Kanthal Precision Wire, Palm Coast, FL). Rats were anesthetized with Nembutal (sodium pentobarbital, 40-50 mg/kg, Abbott Laboratories, North Chicago, IL) and the area of the implantation was shaved. Rats were then placed on a stereotaxic apparatus (Kopf) and 0.1 cc Dual-cillin (Phoenix Pharmaceutical Inc., St. Joseph, MI) was injected intramuscularly into each hindlimb. During surgery, anesthesia was maintained using isoflurane (0.5-2% isoflurane vaporized in medical grade O_2). The scalp was then disinfected with alcohol and swabbed with Betadine (Purdue Frederick, Norwalk, CT). The skin overlying the skull was incised and retracted and the underlying fascia was cleared from the surface of the skull. Excess bleeding was stopped by application of hydrogen peroxide followed by cautery of the retracted fascia. Anchor screws and one ground screw were placed in the skull, and a 1.8 mm diameter craniotomy was opened using a surgical trephine (Fine Science Tools, Foster City, CA). The hyperdrive was positioned over the striatum (Bregma +0.5 mm AP, ± 3.0 mm ML Paxinos and Watson, 1998), and lowered to 1 mm below the surface of the skull. The craniotomy was protected using silastic (Dow Corning 3140) and the hyperdrive was secured in place with dental acrylic (Perm Reline and Repair Resin, The Hygenic Corp., Akron, OH). Following surgery, 10 cc sterile saline (0.9%) was administered subcutaneously, and all tetrodes were advanced ~ 1 mm. Animals were allowed to recover in an incubator until they were ambulatory, which was usually 1–2 hours following surgery. Once animals were ambulatory, 0.8 cc Children's Tylenol was administered orally. For two days following surgery, rats received water containing Children's Tylenol (25 mL in 275 mL of water). To prevent postsurgical infections, treatment regimens of topically applied Neosporin and subcutaneous 0.1 cc Baytril (2.27% enrofloxacin, Bayer Corp., Shawnee Mission, KS) were for some rats. Rats were allowed two days to recover from surgery before resuming experiments. Three animals received right side implants, and two animals received left side implants.

3.4 Recording

Over a period of a week following surgery, tetrodes were advanced 340-680 μm per day until reaching the striatum. The striatum was differentiated from the cortex by the observation of corpus callosum, which is quiet relative to the overlying cortex and underlying striatum. The striatum was further identified by the observation of extremely slow firing cells (< 1 action potential per minute). Neural activity was recorded using a 64 channel Cheetah recording system (Neuralynx, Tucson, AZ). A 72 channel motorized commutator (AirFlyte, Bayonne, NJ; Dragonfly, Ridgeley, WV; Neuralynx, Tucson, AZ) allowed the rats to run the task without twisting the tether cables which connected the hyperdrive to the recording system. Tetrode channels were sampled at 32 kHz. Signals were filtered

between 0.6-6 kHz (Multiple T and Nosepoke, 5 animals) or between 0.3-9kHz (Take 5 task, 5 animals). When the voltage on any of the four channels of a single tetrode exceeded a threshold set by the experimenter, the spike waveform on each of the four channels on the tetrode was recorded and timestamped with microsecond resolution. In 5 animals running the Multiple T and Nosepoke tasks, 1 ms (32 samples) spike waveforms were used, and in 5 animals running the Take 5 task, 2 ms (64 sample) waveforms were used. In some recordings (Take 5 task), the filtered electrical potentials were written directly to disk, and spikes were identified in these recordings offline.

Spikes were clustered offline into putative cells on the basis of their waveform properties using MClust 3.0 (Redish and Schmitzer-Torbert, 2002) or MClust 3.4, with automatic pre-clustering using KlustaKwik 1.0 (Harris, 2002) or KlustaKwik 1.5. Cluster quality was assessed by using L_{ratio} (Schmitzer-Torbert et al., in press).

Unique spike trains After reaching the striatum, tetrodes were often not advanced if cells were observed. In many cases, spike trains recorded in successive sessions represented multiple observations of the same cells. This repeated sampling of the same neurons in multiple sessions could create a bias in analyses of the proportion of cells which fired phasically or were task responsive. To correct for repeated sampling, a set of unique spike trains was defined, and analyses were either restricted to the set of unique spike trains, or results derived from the entire set of spike trains was checked against the set of unique spike trains.

Spike trains obtained from each tetrode were matched across successive sessions on the basis of the correlation of their extracellular waveforms on all four tetrode channels in both sessions. Spike trains with very similar waveforms in successive sessions were considered likely to represent multiple observations of the same cell (for an example, see Figure 5.3). After matching cells across sessions, a set of unique cells was created by selecting the spike train from each set of matched spike trains with the smallest L_{ratio} (i.e. the spike train with the best cluster quality value). This set of matched spike trains was also used in order to examine how the responses of some neurons changed between sessions.

3.5 Histology

Following the completion of all experiments, the final locations of each tetrode were marked with small lesions by passing a small amount of anodal current (5-10 μ A for 5-10 seconds) through each tetrode. At least two days later, rats were deeply anesthetized with sodium pentobarbital and perfused transcardially with saline followed by 10% formalin. Brains were stored in formalin followed by sucrose formalin until slicing. Slices were made either coronally or horizontally through the area of the implantation and stained with either ethidium bromide or cresyl violet to visualize electrode tracks.

3.6 Neural Data Analysis

Descriptive statistics For each spike train, mean firing rate was defined as the number of spikes observed divided by the length of the recording (generally 40 minutes for behavioral tasks, 5-10 minutes for rest periods). Maximum and minimum firing rates were determined by calculating the number of spikes observed in a 1 second sliding window, dividing by the window length, and taking the maximum and minimum values of the measure.

Cross-correlogram The cross-correlogram was created by counting the number of spikes observed from a given spike train (X) in temporal bins of equal duration relative to the time of occurrence of a spike in a second spike train (Y). The count in each bin was then normalized by the number of spikes in spike train Y (the number of alignment events) to generate the firing rate of spike train X relative to the occurrence of action potentials in spike train Y .

Auto-correlogram The auto-correlogram was defined in the same manner as the cross-correlogram, with spike train $X =$ spike train Y . Also, spikes observed at zero-lag were excluded.

Prop_{ISI>X} To separate phasic and tonic neurons, the proportion of time spent in long interspike-intervals was calculated by finding all ISIs which exceeded a criterion (X), summing those ISIs, and dividing by the total session time (Schmitzer-Torbert and Redish, 2004). The measure, Prop_{ISI>X}, gives a measure of what proportion of the session was spent in ISIs equal to or longer than X . In general, X was set to 2 seconds (for example, see Chapter 4), but in Chapter 5, $X = 5$ seconds. Criteria of 2 or 5 seconds yielded similar separations of phasic and tonic neurons.

3.7 Behavioral measures.

Errors Errors occurred when rats explored regions of the Multiple T maze that were not normally sampled, such as incorrect turns on the T maze choices, and regions of the maze that were accessible, but did not lead to the locations of the pellet dispensers. As errors occurred infrequently, it was possible to estimate their occurrence by finding the local density of position samples for every observed position sample, and identifying errors as those position samples found in regions of low density. Density was estimated using Gaussian kernel density estimation, using a Gaussian kernel of width 100 pixels. Potential error trials were those in which position samples were observed in regions of very low density. Each trial was then scored as an error trial if at least 100 low-density position samples were obtained, and was scored as error-free if fewer than 100 such position samples were obtained. The resultant measure was binary, having values of 0 or 1 for each trial, and performed well on a randomly selected subset of data that was examined by a human observer and in which error trials were verified. This method gave similar results to a previously published error-detection method (Schmitzer-Torbert and Redish, 2004).

Path correlation The method by which path correlation was calculated for these analyses was similar to a procedure described elsewhere (Schmitzer-Torbert and Redish, 2002). In these analyses, the video data was divided into $\sim 3.6 \text{ cm} \times 3.6 \text{ cm}$ bins, with each position sample assigned to one bin. Then, the positions sampled on one lap were correlated with the positions sampled on the subsequent lap to measure the similarity in the path taken in both laps. For reconstructed path data (see Chapter 5.2.6), the method was identical, except that the reconstructed position, again expressed in terms of sampling of the $\sim 3.6 \text{ cm} \times 3.6 \text{ cm}$ bins, was used instead of the actual position.

Idealized path For the purpose of constructing linearized spatial rastergrams and histograms (described below), an *idealized path* was created for each Multiple T or Take 5 session by selecting a set of points which followed the path that the rat travelled through on a typical lap (one without errors). This set of points was interpolated linearly so that the distance between points was 1 pixel ($\sim 0.4 \times 0.4 \text{ cm}$). A set of spatial landmarks on the idealized path was also selected: eight for the Multiple T

(one for each pellet dispenser location, one for each turn, and two that marked where the rat turned to enter and exit the turn sequence, see Figure 3.4) and four for the Take 5 task (one for each pellet dispenser location).

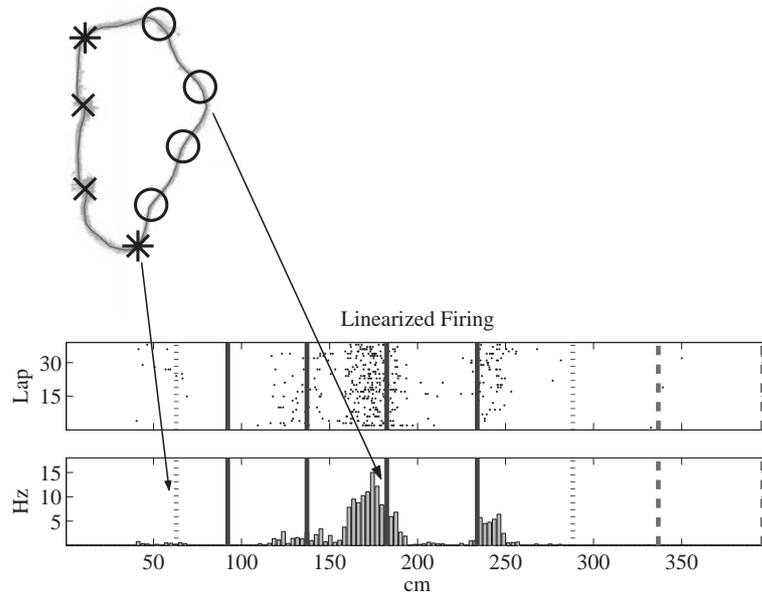


FIGURE 3.4: Linearized path and raster plot. *Top:* Position of the rat during a single session (grey points) and idealized path (solid line). The turn sequence (RRL) is the same as in the schematic in Figure 3.1. Symbols indicate the location of the eight spatial landmarks (circles indicate the four turns, asterisks indicate turns preceding and postceding the turn sequence, and x's indicate the two food delivery sites). *Bottom:* Spatial rastergram and histogram for a dorsal striatal phasic-firing neuron. Dotted lines indicate the start and end of the turn sequence, solid lines indicate turns, and dashed lines indicate the two food delivery sites. Arrows show the correspondence between spatial landmarks (*top*) and the linearized plots (*bottom*). (R010-2001-12-19-TT03-01 Maze = RRL 38 Trials). Figure from Schmitzer-Torbert and Redish (2004)

Linearized spatial plot construction To examine the responses of phasic-firing neurons as rats ran on the maze, a linearized spatial rastergram and histogram were developed. The location of the rat at the time of each spike was mapped to the nearest point on the idealized path, and the location of the spike relative to the idealized path was used to construct linearized rastergrams, and average firing across laps was used to construct linearized histograms (see Figure 3.4).

Warped linearized spatial plots On the Multiple T maze, when rats ran a new maze each day the spatial layout of the turn sequence varied from session to session. For presenting data from a set of PFNs recorded in different sessions (for example, in Figure 5.6), the linearized spatial histograms were warped to 20 bins between each of the eight spatial landmarks. The 20 bins surrounding each spatial landmark (the 10 bins before and the 10 bins after) were taken to represent activity of the PFN near the landmark.

Firing rate To better estimate the instantaneous firing rate of each cell, a measure of continuous firing rate was calculated for each spike train by dividing the session into 10 ms bins, and assigning each spike to one bin. The binned spike train was convolved with a Gaussian ($\sigma = 100$ ms) to create a continuous function of estimated firing rate sampled at discrete intervals of 10 ms. This firing rate

estimate was used in calculations of task responsiveness, in the creation of phasic firing fields, in the temporal versus spatial encoding analysis, and in Bayesian reconstructions (described below).

Task responsiveness The responses of striatal neurons to task parameters were classified using firing rate relative to 1) the time of arrival at each food delivery location, and 2) position on the track. PFNs were classified as *reward-responsive* if they showed a significant increase in firing rate during the five seconds following arrival at either food delivery site. PFNs were classified as *maze-responsive* if they showed a significant increase in firing rate when the rat was running on the maze. To determine if a firing rate was significantly elevated, the mean firing rate of each PFN relative to task events was compared to a distribution of expected mean firing rates created from the same spike train using shuffled event times (i.e. a bootstrap; see Efron, 1982). With a large number of expected mean firing rates, the mean and standard deviation of the distributions of expected mean firing rates can be used as estimates of what the cell's firing rate should be if the cell is not responsive to task parameters. The distributions of expected mean firing rates that were obtained for PFNs frequently exhibited a skew towards positive values, and a square-root transform was applied to normalize the distribution (Sokal and Rohlf, 1995). Under the assumption of normality, estimates of μ and σ from the expected mean firing rate distributions were used to calculate the probability of observing the cell's actual mean firing rate using the inverse of the normal cumulative distribution function. An α of 0.05 was adopted, and a Bonferroni correction for multiple comparisons was applied.

PFNs were classified as *reward-responsive* if the mean firing rate of the PFN in the five seconds following arrival at either food delivery location was significantly larger than the distribution of expected mean firing rates created from 5-second time segments selected randomly from the session. PFNs were classified as *maze-responsive* if the mean firing rate at any location on the maze was significantly larger than the distribution of expected mean firing rates created from similar length time segments selected randomly from the session. For determining maze-responsiveness, the idealized path in each session was divided into 4 regions (Take 5 task) or 8 regions (Multiple T task) using the spatial landmarks described above. On the Multiple T task, if any of the eight regions was more than 1.5 times the average distance between successive turns on the maze, these regions were divided in half.

PFNs which fired very infrequently during the session tended to produce quantized distributions of expected mean firing rates. Such quantized distributions were not normal following the square-root transform. Therefore, all PFNs which fired less than 100 spikes were not considered any further in these analyses, because not enough spikes were observed in the session to accurately estimate the cell's responsiveness to task parameters.

Phasic firing fields On the Multiple T task, to examine the size and distribution of maze-responses in PFNs a quantification of each maze-responsive PFN's activity on the maze was defined. For each maze-responsive PFN, phasic firing fields (PFFs) were defined as each set of continuous 5 cm bins on the linearized spatial histogram which exceeded 50% of the PFN's maximum firing rate in any bin.

Spatial versus temporal encoding On the Multiple T task, a correlation analysis was performed to determine if maze-responses were better related to the location of the rat or to temporal events. For each maze-responsive PFN, the average firing rate at each position along the maze was determined for each lap using the continuous firing rate measure. The correlation of the firing rate as a function

of position was calculated between every pair of laps, and the average correlation was taken to represent how well related the PFN's activity was to the location of the animal on the maze.

For temporal measures, the firing rate of each maze-responsive PFN was calculated for each lap over two temporal windows: the 20 seconds preceding the arrival at the first food delivery site and the 20 seconds following departure from the second food delivery site. For each measure, the correlation of the PFN's firing rate relative to either arrival or departure was calculated for every pair of laps, and the average correlation served to describe how well related a cell's activity was to the time of arrival at the first food delivery site and the time of departure from the second food delivery site.

Reconstruction. Bayesian reconstruction methods were used to determine how well striatal neural ensembles represented task parameters, and the degree to which these representations changed as a function of experience. Bayesian reconstruction follows from Bayes rule: $P(X|F)P(F) = P(F|X)P(X)$. By dividing both sides by $P(F)$, one derives $P(X|F) = \frac{P(F|X) \cdot P(X)}{P(F)}$. This equation allows the calculation of the expected value of X , given the observed firing rate F , once one knows the tuning curve, $P(F|X)$, the prior probability of seeing behavioral variable X , $P(X)$, and the prior probability of seeing firing rate F , $P(F)$. For simplicity of calculation, we treated the cells as independent and calculated $P(X|F_i)$ for each cell i , and then calculated the expected behavioral variable X as the product of the individual expectations, $P(X|F) = C \cdot P(X) \prod_i \frac{P(F|X_i)}{P(F_i)}$, where C is a constant that causes the outcome to be a probability distribution (Sanger, 1996). For stability of calculations, we calculated the likelihood $L(X|F) = \ln C + \ln P(X) + \sum_i [\ln P(X|F_i) - \ln P(F_i)]$.

To calculate $P(X|F_i)$ for a given cell i , we first approximated the probability distribution functions for $P(F)$ and $P(F|X)$. Then, for each bin, on each lap, the values for $P(F)$ and $P(F|X)$ were found from the look-up table. Importantly, in order to avoid the tautology of reconstructing the same data that produced the tuning curves, for each lap, the probability distribution functions were calculated from all other laps (a leave-one-out approach). The results reported here were also consistent with reconstruction that used all of the data (i.e. the tautology).

From the reconstruction, two values were examined: the *reconstruction quality* and the *reconstructed path*. Reconstruction quality was defined as the likelihood of reconstructing the actual location of the rat based on the firing rate of the neural ensemble (reconstruction quality at any time t was defined as $P(X = X_t|F)$, where X_t is the position of the animal at time t). To make more direct comparisons between changes in neural representation and changes in path correlation, the path correlation of the reconstructed path was also examined. The reconstructed location of the rat at any time was defined as the most probable location based on the current firing pattern of the neural ensemble (i.e. $\arg_X \max P(X|F)$).

Measuring changes in performance and reconstruction. For each measure (errors, path correlation, reconstruction quality and reconstructed path correlation), experience dependent changes were examined using a performance criterion. The number of *laps to criterion* ($LC_{measure}$) in each session was defined as the first lap in which the value of the measure was greater than or equal to the median of the values of the measure over all laps. The laps to criterion measures allow an examination of how fast each measure approached asymptotic performance in each session.

Chapter 4

Classification of striatal neurons

As reviewed in Chapter 2.1.1, the striatum is composed of at least five major groups of neurons: projection neurons (medium-sized spiny neurons, or MSPs), and four groups of interneurons (Kawaguchi et al., 1995). One interneuron is cholinergic, while the other three types are GABAergic (Kawaguchi et al., 1995; Kubota et al., 1993). Because striatal neurons can be identified on the basis of intracellular recordings (Kawaguchi, 1993; Kawaguchi et al., 1995), the responses of striatal neurons in slice, anesthetized animals, and organotypic cell culture systems have been well-studied. These studies have revealed many of the important relationships between different cell types and the activity of single neurons, but the relationship between these experiments and the patterns of striatal network activity in *in vivo* in awake, behaving animals is unknown. A better understanding of the role that each striatal neuron type contributes to behavior must acknowledge the very different place that each neuron occupies in the striatal network. In awake, behaving animals, there is a strong need to identify striatal neuron types on the basis of extracellular recordings. Some success has been had in this area in the identification of striatal projection neurons and cholinergic interneurons in the primate (Kimura et al., 1990; Aosaki et al., 1995) which correspond to respectively to phasically active neurons (PANs) and tonically active neurons (TANs). The classification of primate striatal neurons into PANs and TANs is done on the basis of extracellular firing patterns and properties of the extracellularly recorded action potential, which have been related to the properties of MSPs and cholinergic interneurons recorded intracellularly in slice or anesthetized animals. The logic of such an approach is that different patterns of activity will be observed in the awake, behaving animal that can be related to the properties of striatal neurons recorded intracellularly in paralyzed, anesthetized or reduced preparations.

A similar approach to the identification of striatal neuron types on the basis of extracellular recordings has not been aggressively pursued in the awake, behaving rodent. Also, in neither primates nor rodents have the extracellular correlates of other types of striatal interneuron (PV+, NOS+, CR+) been identified. Therefore, the first questions this thesis will address are 1) can neurons recorded extracellularly in the rodent be separated on the basis of firing patterns? 2) can the rodent equivalents of PANs and TANs be identified? and 3) can other types of striatal neuron firing patterns be identified?

4.1 Results

4.1.1 Data sets

A total of 3341 spike trains were collected from 10 rats implanted with hyperdrives over the dorsal striatum. The neural data used in the present chapter included data recorded during one or more behavioral condition (Rest, Multiple T, Take 5, or Nosepoke), and the subsequent two chapters deal with the behavioral correlates of the Multiple T data sets (Chapter 5) and the Take 5 data sets (Chapter 6). Final tetrode locations were verified to lie in the striatum, as shown in Figure 4.1. The majority of final tetrode locations were located in the dorsal striatum, in a region of tissue that included sites in both the dorsomedial and dorsolateral striatum.

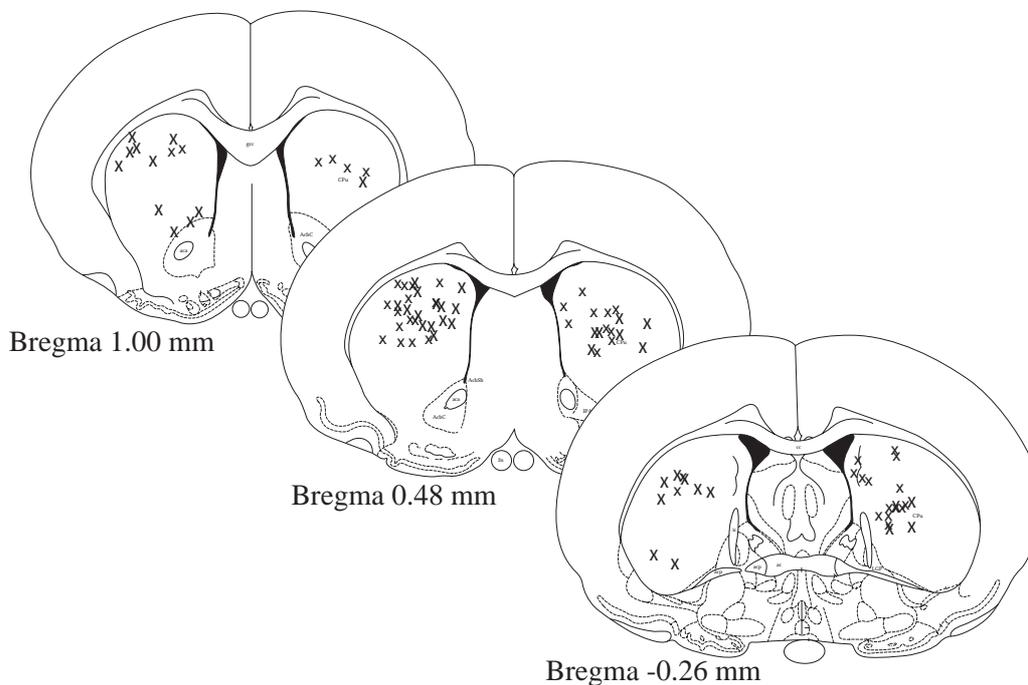


FIGURE 4.1: Recording locations verified histologically. Shown separately are final tetrode locations from ten animals implanted with hyperdrives over the dorsal striatum. Final tetrode positions are marked by x's, and all tetrode locations have been mapped to the nearest of the three coronal sections shown. Tetrodes were observed in a region extending approximately -0.5 to 1.5 mm anterior/posterior relative to bregma. Diagrams adapted from Paxinos and Watson (1998).

4.1.2 Phasic/Tonic separation

Spike trains recorded in the rodent striatum could be separated on the basis of their preference for firing in bouts of activity separated by long ($>1-2$ seconds) pauses (for instance, see the representative spike trains in Figure 4.4). Based on this observation, the proportion of time spent in long interspike-intervals ($\text{Prop}_{\text{ISI}>X}$) identified populations of phasic and tonic neurons in each behavioral condition. $\text{Prop}_{\text{ISI}>X}$ was examined for each behavioral condition over a range of thresholds (X). The range of thresholds which produced a clear separation of phasic and tonic neurons was large and consistent across behavioral conditions. In general, ISI criteria between 500 ms and 10 seconds were able to identify distinct populations of phasic and tonic neurons, independent of the particular behavior rats were engaged in. For the analyses described below, a threshold of 2 seconds was used, and spike

trains which spent less than 40% of the session in ISIs longer than 2 seconds were classified as tonic firing neurons (TFNs), while spike trains that spent greater than 40% of the session in ISIs longer than 2 seconds were classified as phasic firing neurons (PFNs). The $\text{Prop}_{\text{ISI} > 2\text{sec}}$ measure yielded a bimodal distribution of TFNs and PFNs, producing a clear differentiation between these cell types (see Figure 4.2). From the three behavioral tasks, 2420 spike trains were classified as PFNs and 921 spike trains were classified as TFNs.

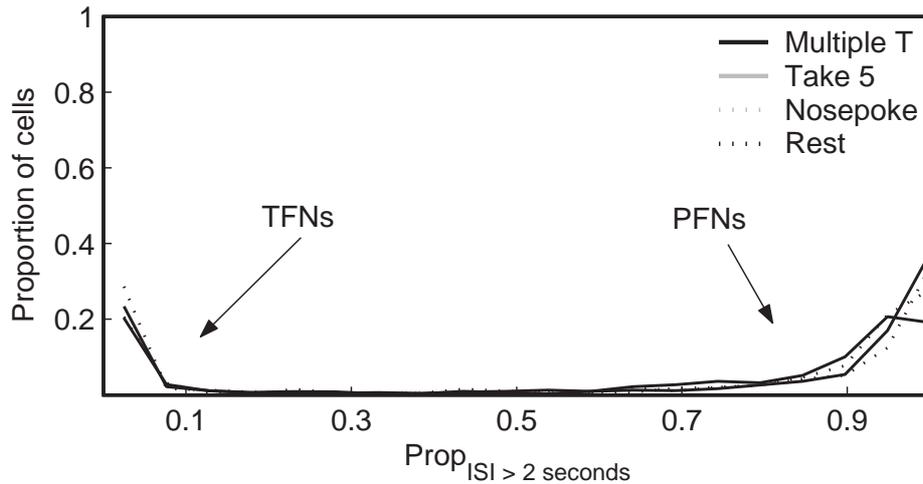


FIGURE 4.2: Separation of phasic and tonic neurons. Plots indicate the distribution of the proportion of time spent in long interspike intervals (>2 seconds, $\text{Prop}_{\text{ISI} > 2\text{sec}}$) for spike trains recorded extracellularly from the rodent striatum in rats performing 4 behaviors: the Multiple T task, the Take 5 task, the Nose poke task, or sitting quietly (Rest). In each task, an ISI criterion of 2 seconds yielded a clean bimodal separation of tonic and phasic neurons.

4.1.3 Multiple subtypes of tonic neurons

TFNs were further differentiated into three subtypes on the basis of their autocorrelations and interspike-interval histograms (see the representative spike trains in Figure 4.4B). Autocorrelations were calculated for each spike train over short (200 ms) and long (10 second) windows following each spike, and principal component coefficients were calculated for each autocorrelation after normalization to a mean of zero and unit standard deviation.

On the basis of the first principal component coefficients of the 10s autocorrelations (Figure 4.3A), TFNs were separated cleanly into two groups, those that tended to “burst” (TFN-1 and TFN-2) and those that did not (TFN-3). Shown in the group averages, in the first 2-3 seconds following the observation of an action potential, TFN-1s and TFN-2s had an increased firing probability, and TFN-3s had a reduced firing probability, relative to their stable autocorrelations. This increased firing probability is taken to be the “burstiness” of the spike train. TFN-1s and TFN-2s were further separated on the basis of the first principal component coefficient of 200ms autocorrelation and the proportion of interspike-intervals shorter than 50 ms (Figure 4.3B), which represents the preference of each cell for high firing rates. TFN-1s and TFN-2s were thus differentiated on the strength of their “burstiness”, with TFN-1s having a higher firing rate in its burst, and a higher proportion of short ISIs than TFN-2s (see also Table 4.1 and Figure 4.4)

Descriptive data for each striatal neuron type is given in Figure 4.4 and Table 4.1. Compared to TFNs, PFNs had low average firing rates (0.5 ± 0.02 Hz). PFNs, TFN-1s and TFN-2s all had a

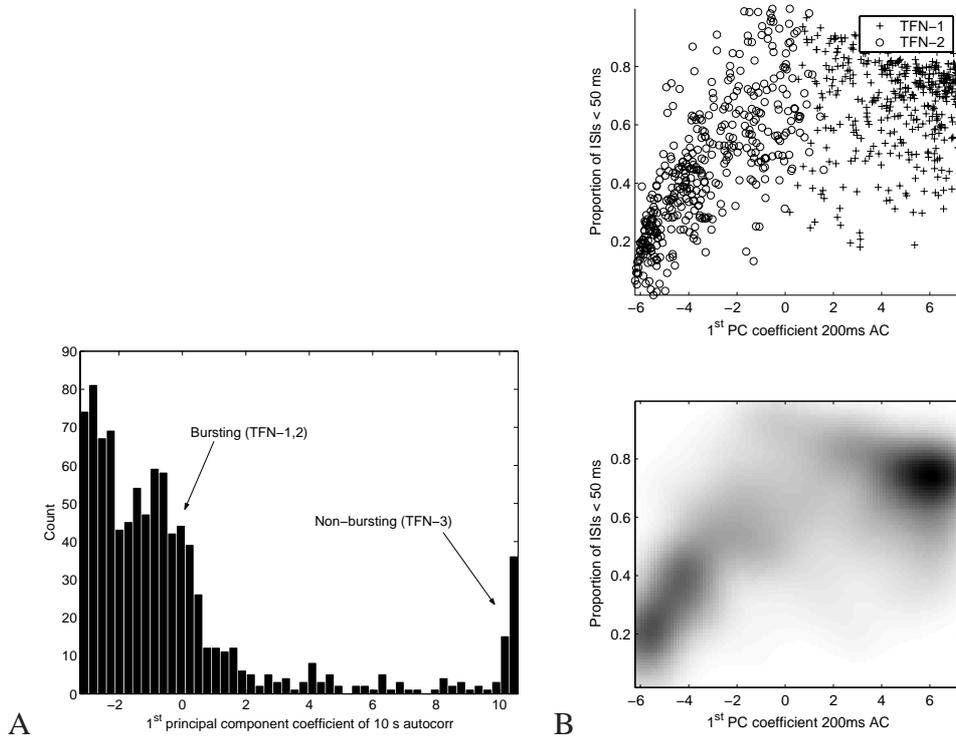


FIGURE 4.3: Identification of TFN subtypes. The contribution of the top 3 autocorrelation principal components of the TFN 200ms and 10 second autocorrelations are shown as a scatterplot (top) and a density plot (bottom). Colored points in the scatterplots indicate the four groups of TFN subtypes identified by the clustering analysis. Principal components were calculated after normalizing each autocorrelation to a mean of zero, and a standard deviation of 1.

tendency to burst which is evident as an early peak in the autocorrelation which decays over several seconds. TFN-1 and -2s had high average firing rates (17.1 ± 1.7 Hz and 14.0 ± 0.9 Hz, respectively), and while both showed a tendency to burst (see Figure 4.4, 10 sec autocorrelations) TFN-1 had a higher burstiness, higher maximal firing rates, and a larger proportion of short ISIs. TFN-3 neurons were slower firing (7.4 ± 1.0 Hz) compared to TFN-1 and -2s, and did not burst. While these four types of firing patterns (PFNs, TFN-1,2, & 3s) were the major groups that could be cleanly identified in these data sets, a fifth, rare type of neuron may have also been present which had low average firing rates (similar to TFN-3s), and very large burstiness.

Cell type	Firing rate (Hz)			% ISIs
	Mean	Max	Min	< 50 ms
PFN	0.50 (0.02)	18.7 (0.3)	0.00 (0.00)	32.1 (0.3)
TFN-1	17.1 (1.74)	64.2 (1.2)	0.79 (0.10)	68.2 (0.8)
TFN-2	14.0 (0.86)	51.7 (1.2)	0.60 (0.14)	45.5 (1.1)
TFN-3	7.4 (0.96)	18.4 (1.2)	0.67 (0.12)	15.5 (1.8)

TABLE 4.1: Average firing rate data for each striatal cell type. Data shown as Mean (SEM) The % ISIs < 50 ms was calculated for each spike train as the proportion of interspike intervals less than 50 ms, and represents the preference of a spike train for high (>20 Hz) firing rates.

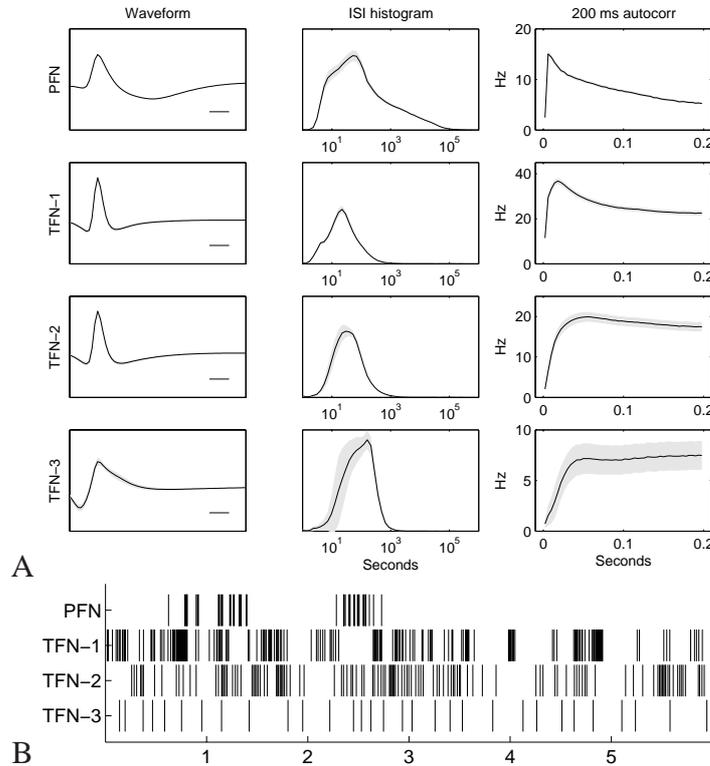


FIGURE 4.4: Cell types. *A*: Average waveform and spiking parameters of PFNs and TFNs. From left to right are shown the average 1) extracellular waveform (2 ms in duration, scale bar = 250 μ s), 2) interspike-interval (ISI) histogram, and 200 ms autocorrelation for each cell type identified. Extracellular waveforms were taken from data sets filtered between 0.3-9 kHz (rats R030, R032, R034, R036, R037, R038), all other data averaged over all spike trains taken from all rats. *B*: Example spike trains selected from each cell type.

4.1.4 Cell classification is stable between behavioral conditions

Within a single session, cell type classifications on the basis of firing patterns remained stable between very different behavioral states. When classifications of spike trains based on firing patterns in the Rest condition were compared to classifications based on firing patterns observed during one of the behavioral tasks (Multiple T, Take 5, or Nosepoke), striatal neurons fell predominantly into the same categories in both behavioral conditions. As shown in Figure 4.5), classifications of spike trains using firing patterns in the Rest condition were maintained during Task performance for each striatal neuron subtype. As the behavior of the rat in the Rest condition (sitting quietly in a terra cotta pot) differed greatly from each of the task conditions (running for food rewards in Multiple T, Take 5, and Nosepoke), these striatal classifications represent firing patterns that are relatively independent of the behavior of the rat over the temporal windows examined (10 to 40 minutes).

Comparing the different types of behavioral tasks (Multiple T, Take 5, and Nosepoke), there were no significant differences in the proportion of striatal neurons falling into each of the four categories for each of the behavioral tasks. As shown in Figure 4.6A&B, similar proportions of PFNs and TFN subtypes were obtained from striatal ensembles in the four behavioral conditions (Rest, Multiple T, Take 5, and Nosepoke). Across the four conditions, there was no effect of behavioral task on the proportion of PFNs which were observed in each rat (ANOVA, $F(3) = 0.36$, n.s.). Also, there was no effect of behavior on the proportions of each type of TFN subtype observed (all ANOVAs, $F(3) < 1.3$, n.s.).

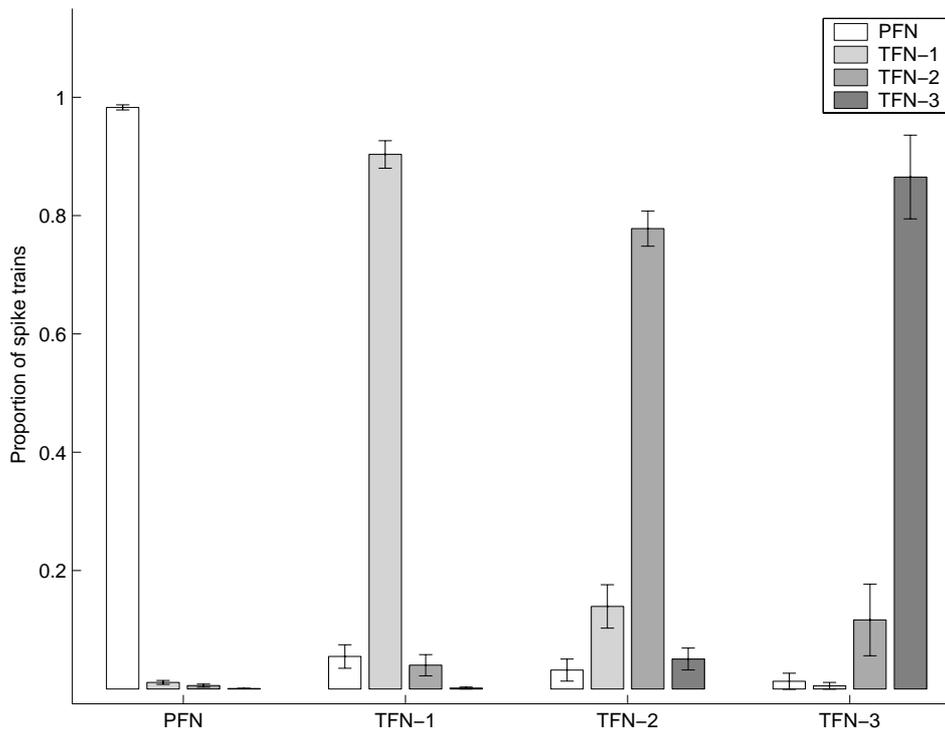


FIGURE 4.5: Stability of classification between rest and behavior. For each type of cell (PFN, TFN-1,2,3), the proportion of cells falling into each category during behavior (Multiple T, Take 5 or Nosepoke) is shown for each category defined during the rest period. Bars represent average across rats, errorbars are 95% confidence intervals.

4.1.5 Spike timing relationships between cell types

In addition to differences in firing patterns, each striatal cell type possessed a unique functional relationship to the firing rates of other striatal cell types. To examine how the firing rate of each cell type depended on the firing of other cell types, cross-correlograms were calculated in each session between every pair of cells recorded on separate tetrodes. Cell pairs from the same tetrode were not included because synchronous spikes recorded on the same tetrode are not easily resolved using this recording system.

Each cross-correlogram reflects the average firing rate of one striatal cell with respect to the action potentials emitted by a second striatal cell. Each cross-correlogram was converted to a probability distribution and cross-correlograms were then averaged within each rat and across rats for each cell type comparison (PFNs vs TFN-1s, TFN-2s vs TFN-3s, etc.). The average data are shown in Figure 4.7. Bins in which the 95% confidence intervals for the mean probability of firing exceeded or fell below chance were taken as evidence of significant relationships between cell types. From this analysis, several strong relationships emerged.

The strongest relationships were obtained between striatal cell pairs in which both cells came from the same cell type category. The firing probability of TFN-3–TFN-3 cell pairs was significantly elevated in a ± 100 ms window centered on TFN-3 spikes. Over a relatively broad window, the activity of TFN-3s was synchronously organized. A similar, but weaker relationship held for TFN-2–TFN-2 cell pairs. TFN-2s were also synchronized, but over a shorter, ± 30 ms window. PFN–PFN cell pairs were also synchronized in the 0 lag cross correlation bin, and in the ± 20 ms bin, but not the ± 10 ms bin. TFN-2–TFN-2 cell pairs did not show significant modulations. Each striatal cell type possessed a unique firing rate relationship, with three cell types firing synchronously specifically

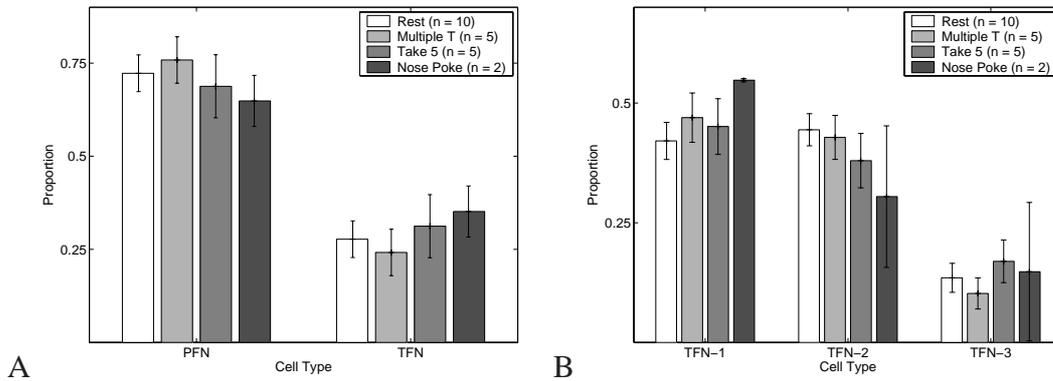


FIGURE 4.6: Stability of cell classification between tasks. A: Proportions of PFNs and TFNs observed in the four behavioral conditions. There were no significant differences in the proportions of PFNs observed in any of the four conditions. B: Proportions of TFN subtypes observed in the four behavioral conditions. There were no significant differences in the proportions of TFN subtypes obtained between conditions. In the legend, n = number of rats. Bars represent average across rats, errorbars are 95% confidence intervals.

with other members of the same cell type. The time scale of this synchronization varied between cell types, with PFNs showing the tightest synchronization, followed by TFN-1s and TFN-3s.

Of the cross-correlations between different types of cells (PFNs with TFN-1s, etc.), the strongest relationship was obtained in TFN-3–PFN cell pairs, which were synchronously active. This relationship was also asymmetrical, with TFN-3s biased to fire before PFN spikes, and PFNs biased to fire after TFN-3 spikes. TFN-2–PFN cell pairs were synchronized at the zero lag cross-correlation bin and TFN-3s were biased to fire 10 ms after TFN-2 spikes.

4.1.6 Relationship of cell type classification to extracellular waveforms

As the shape of intracellularly recorded action potential waveforms can differentiate subpopulations of striatal neurons (Kawaguchi, 1993), extracellular waveforms of striatal neurons were examined using a principal components analysis. The waveform of each cell was represented by the tetrode channel on which the peak of the action potential was largest, and principal component coefficients were calculated after normalizing each average action potential waveform to zero mean and unit standard deviation. Shown in Figure 4.8, the first and second principal component coefficients revealed three typical action potential shapes. These groups were observed in both the 1 ms waveforms from recordings filtered in the 0.6–6kHz range and in the 2 ms waveforms from recordings filtered in the 0.3–9kHz range. The points shown in Figure 4.8 were grouped into three clusters using a K-means algorithm, and the average waveform of the cells included in each group are shown in the figure, and are well-described as being “biphasic”, “triphasic”, or “inverted”. In the inverted group, fewer cells were observed, and it is not clear from this data if there were in fact two types of inverted cells, a biphasic-inverted and a triphasic-inverted group. Therefore, only a single inverted cell type is described here.

Shown in Figure 4.9, PFNs were predominantly biphasic, while TFN-1s & -2s were predominantly biphasic. Of the cell types, TFN-3s were more variable, but tended to fall into the inverted category, especially in recordings which were done in which both positive- and negative-going voltage potentials were used to trigger spike events. On the basis of extracellularly recorded action potentials, PFNs can be clearly differentiated from TFNs, supporting the differentiation of these cell categories on the basis of firing patterns. TFN-1s and TFN-2s had similar extracellular waveforms,

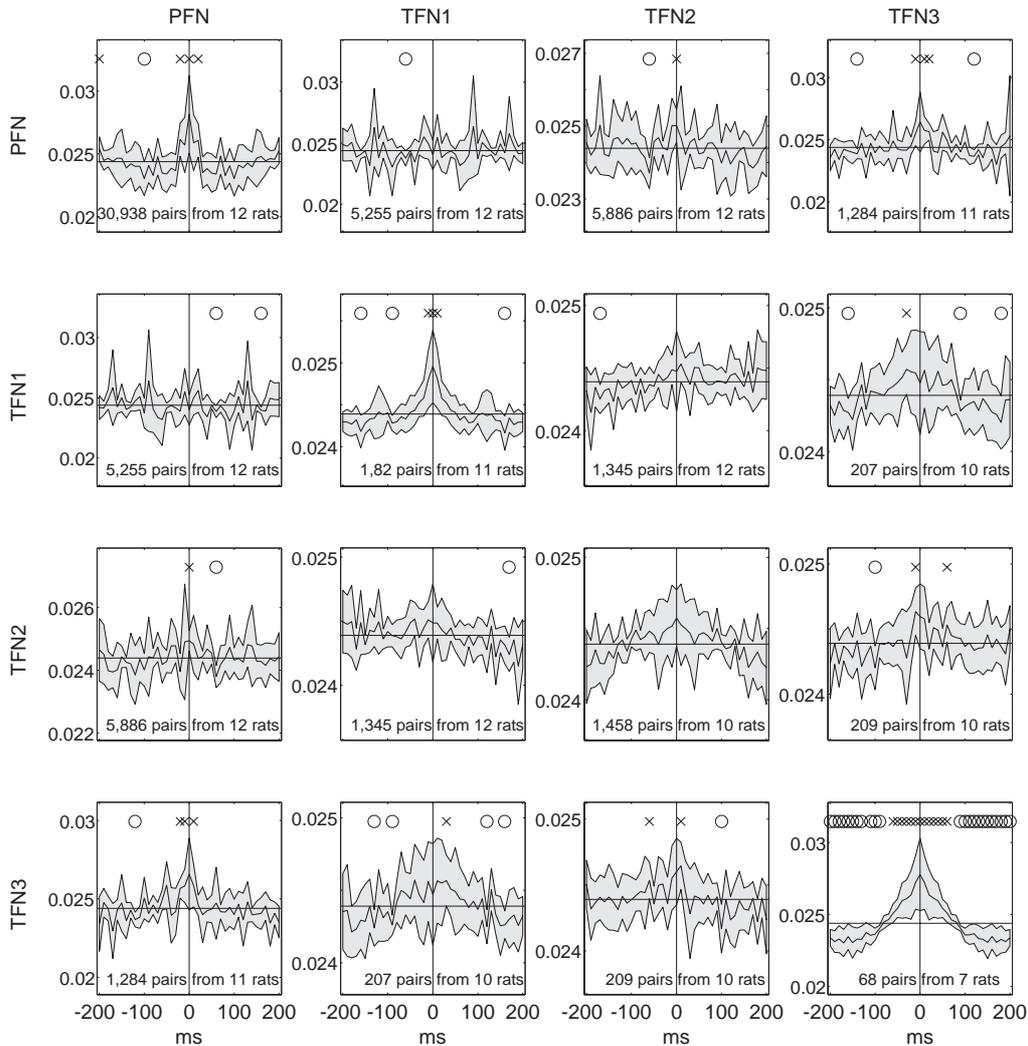


FIGURE 4.7: Average cross-correlations between types of striatal cells. Solid line indicates average over rats, shaded region indicates 95% confidence intervals. Significant relationships were defined as bins in which the lower bound of the confidence intervals exceeded chance (marked by X's) or fell below chance (marked by O's). PFN–PFN cell pairs were synchronized over a window of ± 30 ms, TFN-1–TFN-1 pairs were synchronized over a window of ± 40 ms and TFN-3–TFN-3 cell pairs were synchronized over a window of ± 100 ms. TFN-2–TFN-2 pairs did not demonstrate modulations exceeding chance levels. Relatively weak relationships were obtained between different cell types. TFN-3s were biased to fire before PFN spikes, and PFN firing elevated following TFN-3 spikes.

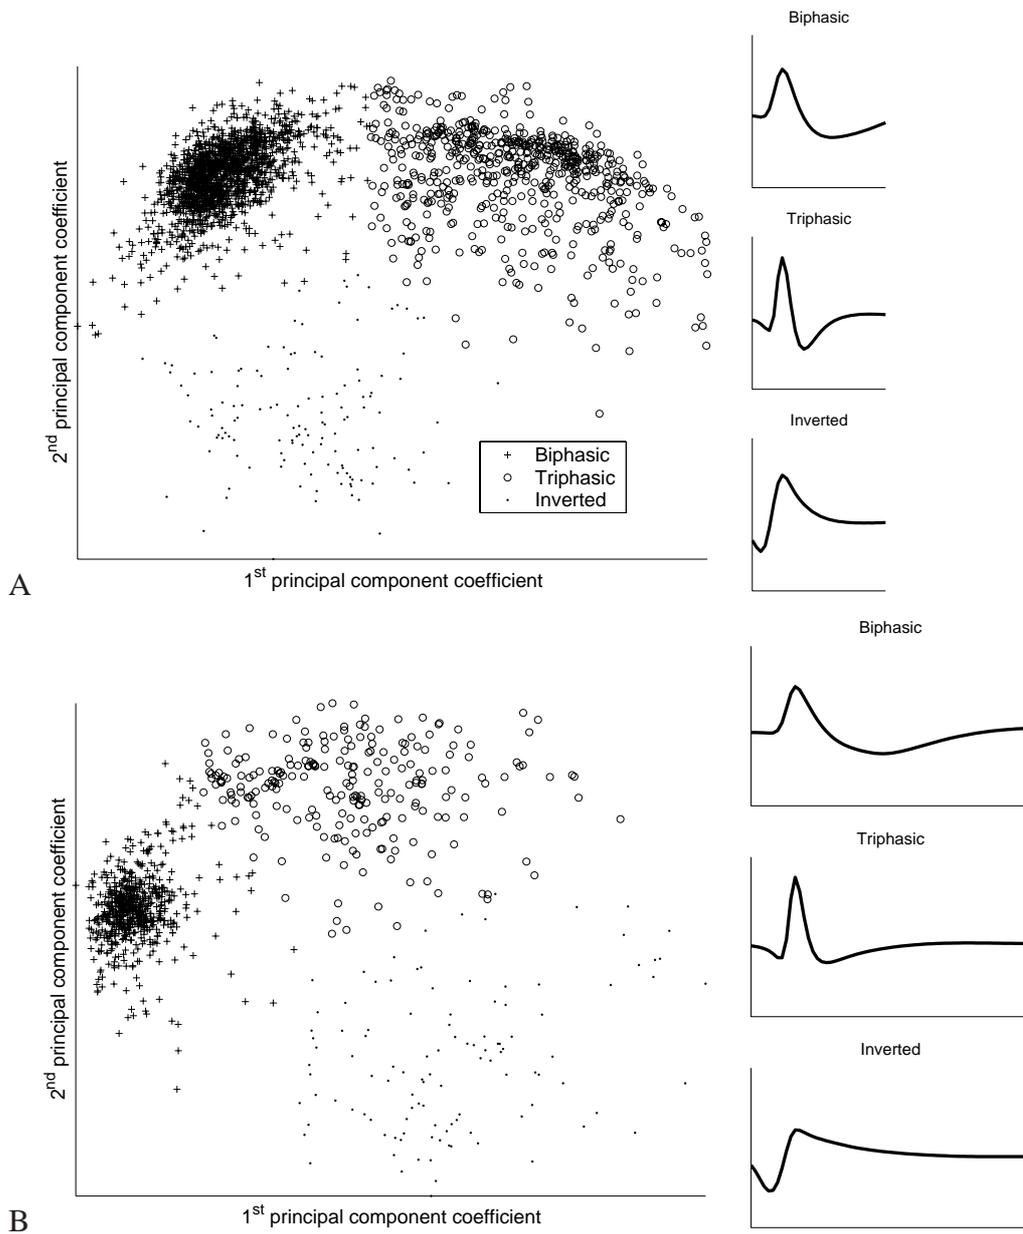


FIGURE 4.8: Principal component analysis of extracellular striatal waveforms. A, B: Scatterplots and average waveforms for energy-normalized extracellular waveforms. Data in A used 1 ms waveforms (filtered between 0.6 and 6 kHz), while data in B used 2 ms waveforms (filtered between 0.3 and 9 kHz). In both data sets, the top two principal components provided a clean identification of three basic waveforms, shown to the right of each scatterplot.

but were both differentiated from TFN-3s, which had a stronger preference for inverted waveform shapes.

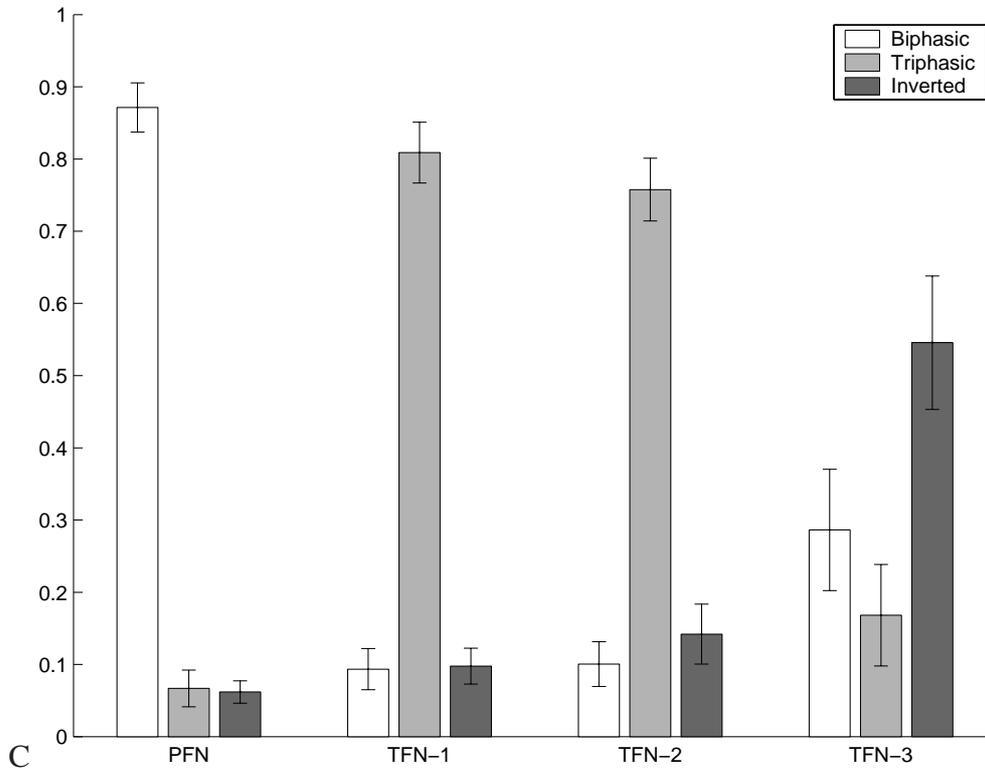


FIGURE 4.9: PFNs were predominantly classified as biphasic waveforms, while two TFN types (1,2) were predominantly triphasic. TFN-3s were more variable, but demonstrated a tendency to be classified as inverted.

Some differences were observed between some extracellular waveform parameters of TFN-1s and TFN-2s. TFN-1s had a significantly smaller spike width, and afterhyperpolarization width compared to TFN-2s, though the differences were small (see Table 4.2).

Cell type	Spike	AHP
	half width	half width
PFN	141 (0.7)	523 (2.6)
TFN-1	87 (2.6)	213 (8.1)
TFN-2	92 (2.3)	245 (8.1)
TFN-3	147 (7.2)	361 (22.2)

TABLE 4.2: Median spike width and afterhyperpolarization duration at half maximum. Measured in milliseconds, values represent Median (SEM).

4.2 Discussion

The data presented in this chapter indicate that in the rodent striatum, four types of striatal firing patterns can be identified on the basis of extracellular recordings. Striatal neurons were first separated into phasic firing (PFNs) and tonic firing (TFNs) neuron types, on the basis of the amount of time spent in long interspike interval (ISIs). Then tonic firing neurons were subsequently separated into

three subtypes on the basis of their autocorrelations and ISI histograms. Predominantly non-bursting neurons (TFN-3s) were separated from bursting neurons (TFN-1 and TFN-2s), and high-firing rate bursting (TFN-1s) were separated from lower-firing rate bursting neurons (TFN-2s). These classifications were stable both across animals performing different tasks, and within single neurons observed during two very different behaviors (rest versus active). Striatal neurons further possessed different extracellular waveforms, and cross-correlation relationships. Together, these data suggest that these striatal firing patterns reflect the existence of different populations of spike generators in the striatum, which may in turn reflect differences between striatal neurons.

4.2.1 Cell type correspondence

Projection neurons versus interneurons In the primate striatum, phasically active neurons (PANs) have been identified as projection neurons, while tonically active neurons (TANs) are thought to be striatal interneurons (Kimura et al., 1990). The phasic, or bursty, firing patterns of PANs correspond well with the properties of identified MSPs reported from intracellular recordings. MSPs have a bistable membrane potential, and shift from a hyperpolarized “Down” state to a depolarized “Up” state in which they fire action potentials. MSPs are thus “bursty”, firing single action potentials or trains of action potentials when in up states, and quiescent in down states (Wilson and Groves, 1981). This alternation of activity and periods of quiescence is thought to be present in awake, behaving primates (Kitano et al., 2002), and allows for projection neurons to be identified in awake, behaving animals. In the data presented in the present chapter, PFNs are thus likely to correspond to projection neurons, while TFNs are more likely interneurons. PFNs were predominantly silent during the task, spending most of the recording session in long ISIs, while TFNs rarely paused.

Cell Types	Mean FR	“Burstiness”	Correspondence
PFN	<1 Hz	++	medium-sized spiny projection neurons?
TFN-1	17 Hz	+++	PV+ GABAergic interneurons?
TFN-2	14 Hz	+	NOS+ or CR+ GABAergic interneurons?
TFN-3	7 Hz	-	TANs/cholinergic interneurons?

TABLE 4.3: Proposed correspondence of PFN/TFN cell types to known striatal neurons.

TANs Although primate TANs are held to be striatal interneurons, it is not clear that all populations of striatal neurons are found in this category. Many studies define TANs to be striatal neurons which fire between 4 and 15 Hz on average and have broad extracellular waveforms (Raz et al., 1996). Such criteria would exclude both phasic neurons and higher-firing, short duration extracellular units, which correspond to PFNs, TFN-1s and TFN-2s in our studies. In our classification scheme, TFN-3s share many similarities to primate TANs and may represent the rodent TAN equivalent. TFN-3s fired a low, tonic rates, did not burst, and had relatively broad extracellular action potentials relative to TFN-1s and TFN-2s. TFN-3s were also synchronized over a relatively broad window of ± 100 ms, which agrees well with other reports of the cross-correlation relationships between simultaneously recorded primate TANs (see for instance Figure 2 in Raz et al., 1996).

Other striatal interneurons TFN-1s and TFN-2s represented tonic, high-firing neurons relative to PFNs and TFN-3s. These properties are inconsistent with both PANs and TANs, as well as the known properties of MSPs and cholinergic interneurons of the striatum. TFN-1s and -2s may therefore represent the activity of other striatal interneurons, some of which are known to be capable of firing in bursts and at high rates (Kawaguchi et al., 1995). Of the three GABAergic interneurons, the firing properties of parvalbumin-immunoreactive (PV+) and the nitric oxide synthase/somatostatin/neuropeptide Y-immunoreactive (NOS+) neurons have been described, while to my knowledge, no published reports exist of the firing properties of calretinin-immunoreactive (CR+) neurons. PV+ striatal interneurons correspond to fast-spiking (FS) striatal neurons, which are capable of firing at high rates (see Kawaguchi, 1993; Bracci et al., 2002). PV+/FS neurons are connected by gap junctions (Koós and Tepper, 1999), receive multiple synaptic contacts from individual corticostriatal axons (Ramanathan et al., 2002) and exert a powerful inhibitory influence on MSPs (Koós and Tepper, 1999). NOS+ striatal interneurons correspond to low-threshold spiking neurons, which are capable of producing low-threshold spikes in addition to fast spikes, and are thought to be able to fire bursts (Kawaguchi, 1993; Kawaguchi et al., 1995).

Of all the TFN subtypes, TFN-1s had the highest average firing rates, maximal firing rates, and the strongest skew towards short interspike intervals. These properties are consistent with those of the PV+/FS striatal interneurons, which are the fastest firing neuronal type. TFN-1s also had a tendency to fire synchronously at shorter timescales than were observed for TFN-3s. PV+/FS neurons are known to be connected by gap junctions, which are thought to allow synchronous firing in connected neurons (Bennett, 1999), as has been demonstrated in the cortex (Galarreta and Hestrin, 1999; Gibson et al., 1999) and cerebellum (Mann-Metzer and Yarom, 1999). The synchronous firing of TFN-1s may therefore reflect the short-timescale synchronization of PV+/FS neurons in the striatum by electrical synapses.

Like TFN-1s, TFN-2s had high mean firing rates and maximal firing rates. However, TFN-2s had lower firing rates than TFN-1s, were less biased to fire with short interspike intervals, and did not have a tendency to fire synchronously. These properties suggest that TFN-2s may represent a different neural type than that corresponding to TFN-1s. Whether TFN-1s correspond to NOS+/LTS neurons or to striatal CR+ is unknown, but the bursty properties reported for TFN-2s are consistent with the proposal that NOS+/LTS neurons are capable of firing in bouts of activity. Provisionally, TFN-2s can be proposed to be a high firing striatal interneuron that is not a member of the PV+/FS interneurons, and may correspond to striatal NOS+/LTS interneurons.

Chapter 5

The Multiple T Task

5.1 Introduction

As reviewed in Chapter 2.2, extracellular recordings in primates have identified two types of neuron (phasically active neurons, PANs and tonically active neurons, TANs) which correspond respectively to striatal projection neurons and cholinergic interneurons. The data presented in Chapter 4 indicate that striatal neurons in the awake, behaving rat can be categorized along a similar phasic/tonic continuum. It is thus likely that rodent phasic-firing neurons (PFNs) correspond to striatal projection neurons and primate PANs, while rodent tonic-firing neurons (TFNs) correspond to striatal interneurons, with different TFN subtypes likely to correspond to different types of striatal interneurons.

With this classification scheme (PFNs, TFN subtypes), the present chapter examines the activity of single striatal neurons and ensembles of neurons in rats performing a sequential navigation task, the Multiple T maze. As described in Chapter 2.3, navigation and instrumental learning studies in the rodent have demonstrated a striatal involvement in the learning of habitual navigation strategies and stimulus-response behavior in operant tasks. While some studies have shown plasticity in the responses of striatal neurons as animals learn to perform navigation and instrumental tasks (Carelli et al., 1997; Jog et al., 1999), most of these studies are carried out over a period of several recording sessions, and demonstrating changes in individual neurons has remained elusive.

To address this question, rats were trained to perform a sequential navigation task, the Multiple T maze, in which animals were required to navigate through a sequence of 3-5 T maze choices in order to receive food rewards. A key element of the task was that in each session, the sequence of turns presented to rats could be changed. This manipulation allowed for the examination of changes in behavior and neural activity as rats learned to perform a completely novel maze sequence.

5.2 Results

5.2.1 Behavior

Behavior Behavioral data was available from 280 sessions taken from 19 rats running the Multiple T task. These sessions included 12 sessions in which rats ran 3T mazes, 136 sessions in which rats ran 4T mazes and 132 sessions in which rats ran 5T mazes. These sessions consisted predominantly of sessions in which rats were presented with a new maze sequence each day (76% of the 280 sessions). In these sessions, rats ran an average of 55.9 ± 4.2 laps per session (mean \pm SEM over rats).

In each Multiple T session, the probability that rats would run a correct lap (one without errors) was lowest in the first 5-10 trials. As shown in Figure 5.1A, rats rapidly increased their probability of navigating through the maze without making an error, indicating rapid learning of the maze sequence.

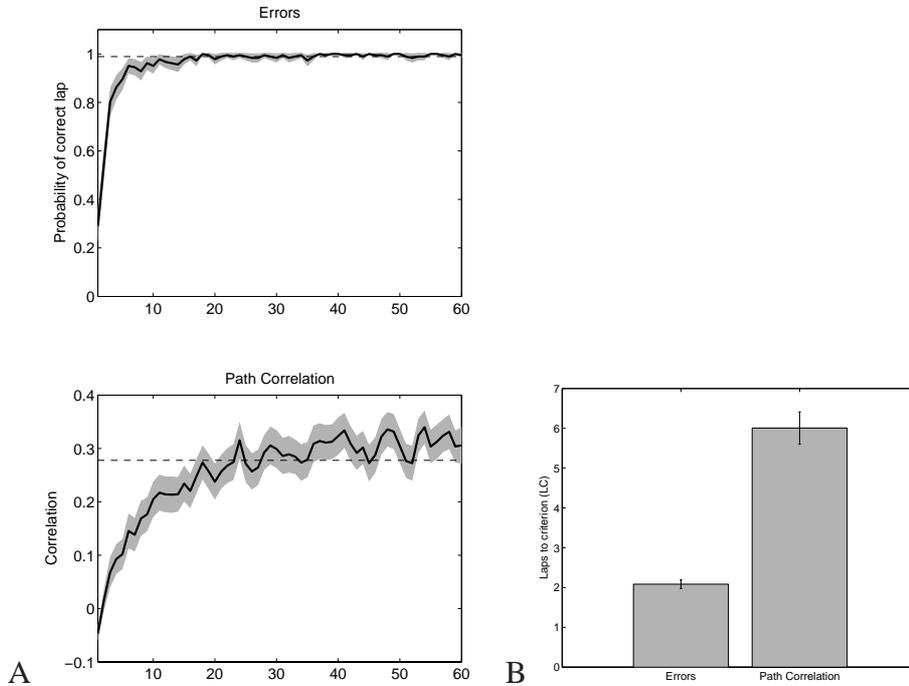


FIGURE 5.1: Error-elimination and path-refinement on the Multiple T maze. *A*: Probability of completing an error-free trial (Errors, *top*) and the correlation of the path taken on trial i with the path taken on trial $i + 1$ (Path correlation, *bottom*). Each measure increased as a function of experience, with the error-elimination proceeding more quickly than path-refinement. Data taken from 15 rats and includes 149 sessions in which rats completed at least 60 trials during the session. Path correlations from each session have been normalized to z-scores to compare across sessions. *B*: Differences between measures, as assessed using the laps to criterion measure (LC).

Similarly, the path rats took through the elevated maze was refined as a function of experience, though at a slower rate than the error measure. As shown in Figure 5.1A, the average normalized path correlation function increased over the first 20+ trials of each session, indicating that the path rats used to navigate through the maze underwent a slow process of refinement as a function of experience. The difference between these two behavioral learning rates (error-elimination and path correlation) was significant. As shown in Figure 5.1B, the number of laps to criterion for the error measure (LC_{error}) was significantly smaller than that of the path correlation measure ($LC_{pathcorr}$, Kruskal-Wallis $\chi^2(1, N = 558) = 221.4, p < 0.001$) indicating that on average, path-refinement continued after the cessation of errors. These data indicate that rats could perform the Multiple T task well, even when presented with maze-sequences which were entirely novel, and could learn to perform novel mazes within a single day. Further, learning to perform the Multiple T maze was reflected by two types of changes in behavioral performance, each with a different timecourse.

How then do these behavioral learning rates depend on maze familiarity? To address this question, behavioral data was examined from six animals that completed a 21 day protocol in which they ran a week of novel mazes (Novel) followed by a week in which the same maze was presented each day (Familiar) followed by a final week in which a new maze sequence was presented each day (Novel). In the 126 sessions in which rats ran 4T mazes in the Novel/Familiar/Novel protocol, rats ran an average of 52.4 ± 8.3 laps per session. As shown in Figure 5.2, LC_{error} was significantly

smaller when rats were running familiar mazes than when rats were running novel mazes (Kruskal-Wallis $\chi^2(1, N = 116) = 17.99, p < 0.001$), indicating that rats eliminated errors more quickly when the same maze sequence was presented repetitively.

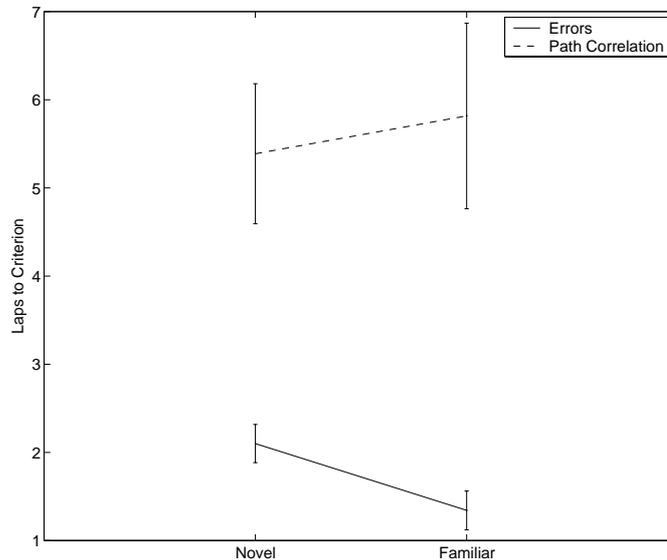


FIGURE 5.2: Effects of maze familiarity on behavioral learning rates. When performance on sequences of turns which were novel (i.e. never before experienced) were compared to those which were familiar (i.e. the same sequence was presented repeatedly), rats made fewer errors on familiar mazes than novel mazes. In contrast, there were no significant differences in the number of laps to criterion for the path correlation measure.

However, maze sequence familiarity had no effect on $LC_{pathcorr}$ (Kruskal-Wallis $\chi^2(1, N = 116) = 1.04, p = 0.31$), as can be seen in Figure 5.2. Thus, in these analyses the rate of path refinement was not observed to differ between sessions in which rats were running Novel or Familiar mazes, while a significant savings in the rate of error-elimination was found when rats ran Familiar mazes.

5.2.2 Neurophysiology

To explore how neural activity in the rodent striatum correlated with task performance, extracellular recordings were made in the dorsal striatum (see Figure 4.1) as rats ran on the Multiple T maze. Neural data was taken from 5 rats which were implanted with hyperdrives and which completed the three week Novel/Familiar/Novel protocol using 4T mazes. From 104 recording sessions, 2,125 spike trains were obtained (21 sessions from 4 rats, 20 sessions from one rat. 425 ± 64 spike trains per rat). Between successive sessions, spike trains recorded on the same tetrode had an average waveform correlation of 0.729 (SD = 0.009). Spike trains which were matched across sessions (i.e. considered to be the same cell in both sessions) had an average waveform correlation of 0.977 (SD = 0.002). As an example of the waveform correlations and stability of cells across sessions, two successive sessions from a representative tetrode are shown in Figure 5.3. As can be seen in the figure, the waveform correlations successfully identified clusters which were stable across the two sessions. 1,144 spike trains were judged to be unique on the basis of the correlation of the spike waveforms between sessions, as described in the Methods (229 ± 28 unique spike trains per rat). 275 spike trains with values of $L_{ratio} > 0.05$ were removed from this set of unique spike trains, leaving a total of 867 spike trains (173 ± 22 spike trains per rat). Analyses were restricted to this set of unique spike trains, but the results reported here were consistent with the entire data set of 2,125

spike trains. From the set of unique spike trains, 589 (68%) were classified as phasic-firing neurons (PFNs). The remaining 278 (32%) were classified as tonic-firing neurons (TFNs). As described in Methods, PFNs with less than 100 spikes had too few spikes to estimate responsiveness to task parameters. Of the 589 PFNs in the unique spike train sample, 194 contained less than 100 spikes and were eliminated from further analysis.

5.2.3 Phasic-firing neurons

Task responses

Of 395 PFNs, 108 (22 ± 9 cells per rat) PFNs were responsive in one or more regions on the maze, and 81 (16 ± 7 cells per rat) PFNs were responsive during the five seconds following arrival at one or both food delivery sites.

Maze-responsive PFNs

Of the PFNs which fired at least 100 spikes a session, 27% were classified as maze-responsive. Figure 5.4 shows an example of a maze-responsive PFN which was active at one location on the maze (as the rat ran between the last turn on the maze and the first food delivery location). Maze-responsive PFNs had between one and six phasic firing fields (PFFs, median number of PFFs per cell = 1), with a median PFF width of 3 bins (~ 15 cm).

Maze-responses were related to spatial cues Maze responses were better related to the position of the rat on the maze than to the time at which the rat arrived at the first food delivery site or the time at which the rat departed the second food delivery site. Across laps, the correlation of each maze-responsive PFN's activity was examined with respect to: 1) the position of the rat on the maze, 2) the 20 seconds preceding arrival at the first food delivery site, and 3) the 20 seconds following departure from the second food delivery site. Higher correlations were obtained for the spatial reference frame than either temporal reference frame. Shown in Figure 5.5 *left* is a PFN with maze-responses that were well related to the location of the rat on the maze and poorly related to both temporal reference frames. The jitter in spike timing observed in the temporal plots is indicative of behavioral variability of the animal across laps. Across rats, the median amount of time rats took from leaving the second food delivery site to arriving at the first food delivery site was 13.5 ± 1.0 seconds. The within-session standard deviation of this departure-to-arrival time was 9.3 ± 2.7 seconds. This within-session variability allowed for the dissociation of responses to spatial and temporal events.

Across the set of 108 maze-responsive PFNs, there were significantly larger correlations in the spatial reference frame than either temporal reference frame (see Figure 5.5 *right*). Paired t-tests across five animals: Space versus Arrival, $t(4) = 3.593$, $p = 0.023$, Space versus Departure, $t(4) = 3.560$, $p = 0.024$).

Maze-responses distribute evenly on the turn sequence To determine if the PFN maze-responses favored specific locations on the maze, the distribution of phasic firing fields (PFFs) over the maze was examined. Figure 5.6 shows the PFFs obtained from maze-responsive PFNs, sorted according to the center of each PFF. Maze-responsive PFNs responded at every location along the length of the turn sequence. The R^2 from a linear regression on the sorted PFF centers was 0.98, indicating that

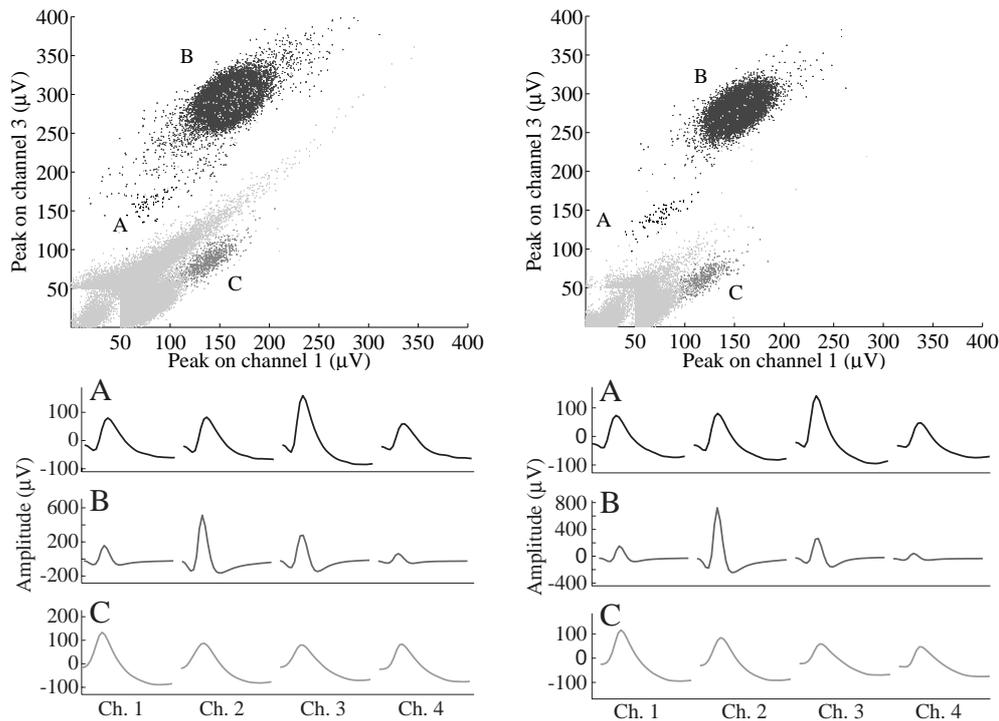


FIGURE 5.3: Example of cells matched across sessions. Data from the same tetrode recorded in successive sessions shows recording stability across days. Three clusters were isolated from each tetrode recording, and were matched across sessions on the basis of the correlation of the average waveforms in each session. *Top:* Peak waveform amplitudes on channel 3 versus channel 1 for both sessions. Points identified as members of one of the three clusters identified in these recordings are plotted black, all other points are shown in gray. *Bottom:* Average waveforms of the three clusters shown above. Waveform correlations of the matched spike trains between sessions (i.e. correlation of the average waveform of cluster A in session 1 with the average waveform of cluster A in session 2, etc.) were > 0.98 . Values of L_{ratio} for these clusters were: *Left:* Cell A = 8.9×10^{-6} , Cell B = 0.00, Cell C = 0.0028 and *Right:* Cell A = 1.3×10^{-4} , Cell B = 0.00, Cell C = 0.0004. Data from R018-2002-09-22-TT01 and R018-2002-09-23-TT01. Figure from Schmitzer-Torbert and Redish (2004)

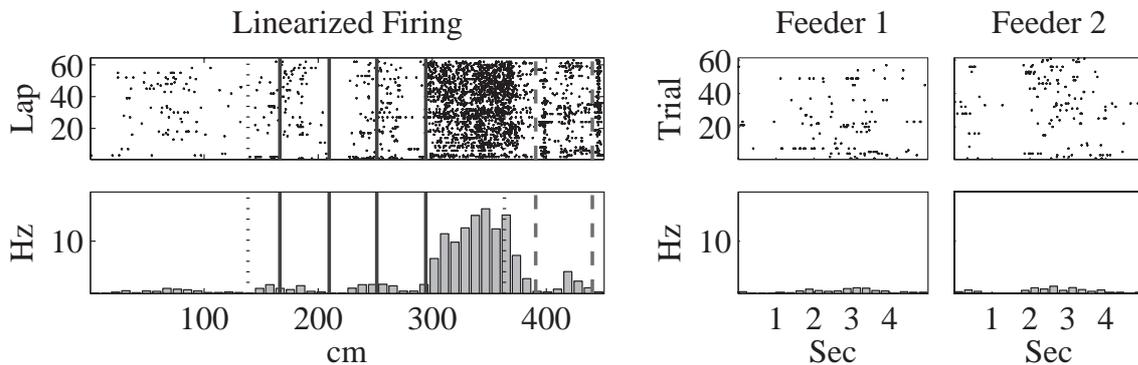


FIGURE 5.4: A maze-responsive PFN. *Left:* Rastergram and histogram of linearized firing on the maze. Symbols indicate the location of the eight spatial landmarks (circles indicate the four turns, asterisks indicate turns preceding and postceding the turn sequence, and x's indicate the two food delivery sites). *Right:* Peri-event time histograms (PETHs) of the firing rate of the cell as the rat arrived at the first and second food delivery sites. This cell was active as the rat ran between the fourth turn and the end of the turn sequence. (R023-2002-08-27-TT07-02 Maze = LRL Trials = 61). Figure from Schmitzer-Torbert and Redish (2004).

the centers of the PFFs were well described by a linear fit. This implies that maze-responses were uniformly distributed on the maze, and did not concentrate at specific landmarks.

Maze-responses did not encode general actions Given that maze-responses did not favor specific locations on the maze, did maze-responses depend on the specific actions that the rats were making? From the example shown in Figure 3.4, we might not expect this to be the case. This PFN had responses at two left turns on the maze: a strong response at the third turn (~ 40 Hz) and a weak response at the fourth turn (~ 10 Hz). This PFN was not equally responsive to all left turns, however. It was silent at two other left turns, where the rat was entering and exiting the turn sequence. To examine how the activity of maze-responsive PFNs as a group depended on the actions that the rats were making, the firing rate in the phasic-firing fields of maze-responsive PFNs was compared to the firing rate of the same PFN at other regions of the turn sequence where the shape of the rat's path was highly similar or dissimilar. Similar and dissimilar paths were defined as locations where the rat's path was well and poorly correlated respectively. In cases where the shape of the paths were very similar, the rat was likely to be making similar motor actions, such as turning in the same direction. In cases where the shapes of the paths were dissimilar, the rat was likely to be making different motor actions, such as turning in opposite directions.

Of the 108 maze-responsive PFNs, 58 cells (from 5 animals, 11.6 ± 6.0 cells per rat) had at least one phasic-firing field (PFF) on the turn sequence. The firing rate and path of the rat in each PFF was compared to other locations on the turn sequence using a sliding comparison window whose size was equal to that of the PFF. For each comparison window, the firing rate in the PFF as a function of position along the idealized path was correlated with the firing rate as a function of position in the comparison window. The rat's path in the PFF was also correlated with the rat's path in the comparison window, to quantify how similar the route taken by the rat in the comparison window was to the route taken in the PFF. If maze-responsive PFNs encoded general actions, then the firing pattern in the PFF should be highly similar to the firing pattern observed at other locations where the rat's route was similar to the route taken through the PFF. Also, we would expect that the firing pattern in the PFF should be poorly correlated with the firing pattern observed at other locations where the rat's route was dissimilar to the route taken through the PFF. In these analyses, a similar

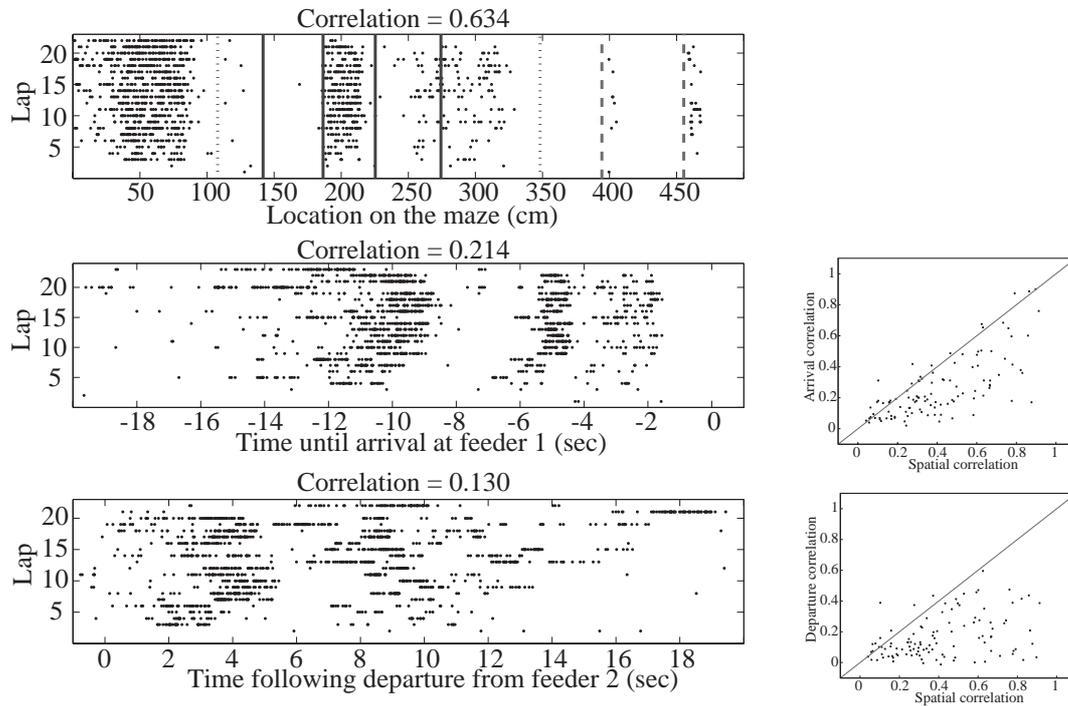


FIGURE 5.5: Maze-responses were better related to spatial than temporal parameters. *Left:* Example of the activity of a maze-responsive PFN with respect to: 1) location on the maze (*top*), 2) in the 20 seconds preceding arrival at the first food delivery site (*middle*), and 3) in the 20 seconds following departure from the second food delivery site (*bottom*). The same set of spikes is shown in each plot. The activity of this PFN was stable with respect to space and variable with respect to the two temporal events. Across laps, this maze-responsive PFN’s activity was better correlated in the spatial reference frame (correlation = 0.63) than in either temporal reference frame (correlations of 0.21 and 0.13). *Right:* Across all maze-responsive PFNs, there was a significant bias towards higher correlations in the spatial reference frame than in the 20 seconds preceding arrival at the first food delivery site (*top*) or in the 20 seconds following departure from the second food delivery site (*bottom*). Figure from Schmitzer-Torbert and Redish (2004).

route was defined as having a path correlation greater than 0.85, and a dissimilar route was defined as having a path correlation less than -0.85. In Figure 5.7 we can see that there was a significant bias towards similar firing patterns in other regions on the maze where the rats’ paths were similar compared to other regions on the maze where the rats’ paths were dissimilar (Kruskal-Wallis $\chi^2(1, N = 116) = 9.98, p = 0.0016$). However, neither group of correlations was biased toward positive correlations, which indicates that even in the Same Route group, there was no tendency for a maze-responsive PFN to fire similarly at other locations on the maze in which the rat’s path was similar to that taken through the PFF. To the extent that the shape of the rat’s path is an indication of the motor activity of the rat, these results indicate that maze-responsive PFNs did not purely encode movements.

Sequence specificity Maze-responses of PFNs were not well described by the actions rats were taking, but they were well described by a combination of action, sensory-context and the specific sequence of turns rats were presented with. 22 maze-responsive PFNs were recorded in at least 2 successive sessions and had at least one phasic-firing field (PFF) on the turn sequence in at least one of these sessions. For each maze-responsive PFN, the correlations between firing pattern and path of the rat were tested. Pairs of sessions were considered in which the center of the PFF and center of the same region in the matching session were in the same location in the environment. Three groups

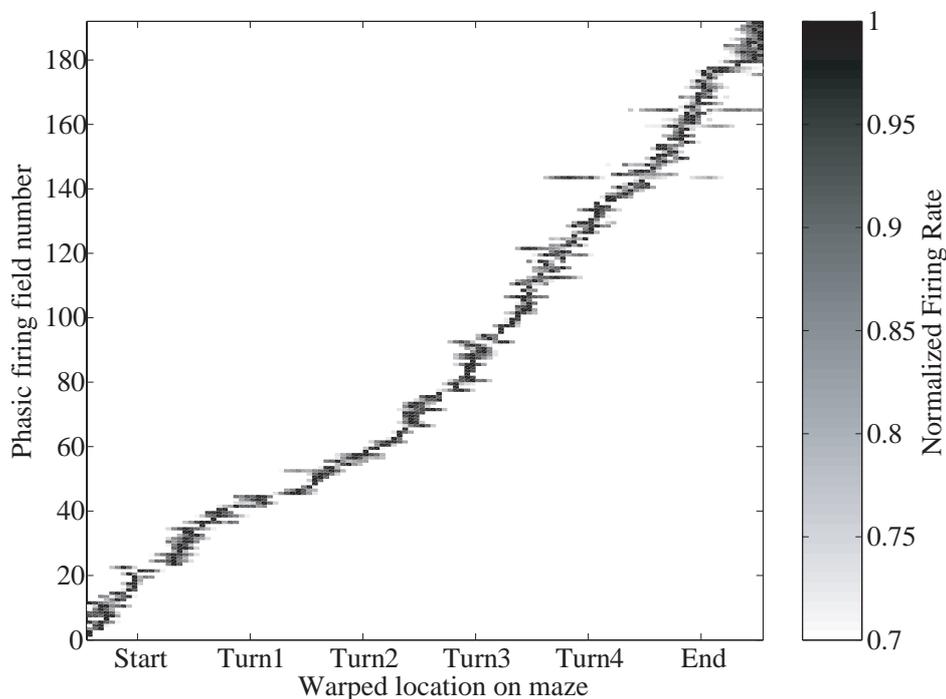


FIGURE 5.6: Distribution of spatial activations on the Multiple T task. To make comparisons across sessions (i.e. across different maze configurations), the response of each PFN on the turn sequence was warped to a fixed number of bins between landmarks (see Methods for a description of the warping process). For each responsive PFN, phasic firing fields (PFFs) were defined as each set of continuous bins on the warped spatial histogram which exceeded 50% of the PFN's maximum firing rate in any bin considered. A PFN could have multiple phasic firing fields. In the plot above, each PFF has been aligned to its center, and normalized to a maximum of 1 for display purposes. The field centers are well fit with a linear regression ($R^2 = 0.98$), implying that PFF centers are distributed evenly across the turn sequence. Figure from Schmitzer-Torbert and Redish (2004).

of correlations were examined: 1) Same Maze: cases where the rat ran the same sequence of turns in both sessions, 2) Same Route: cases where the rat ran a different sequence of turns in each session, but the path taken through the phasic-firing field and the path taken through the same region of the matching session were correlated by at least 0.85, and 3) Different Route: cases where the rat ran a different sequence of turns in each session, but the path taken through the phasic-firing field and the path taken through the same region of the matching session were correlated by less than -0.85. To correct for multiple observations of the same cell, the set of correlations in each group obtained from one maze-responsive PFN and its matching spike trains were averaged.

In the group of Same Maze pairs, correlations from 14 cells were included (from 4 rats, 3.5 ± 1.5 cells per rat). In the group of Same Route pairs, correlations from 4 cells were included (from 3 rats, 1.3 ± 0.4 cells per rat). In the group of Different Route pairs, correlations from 10 cells were included (from 5 rats, 2.0 ± 0.5 cells per rat). Shown in Figure 5.8 are the firing rate correlations for each group. There was a significant difference between groups (Kruskal-Wallis $\chi^2(2, N = 28) = 12.85$, $p = 0.0016$), and pairwise comparisons using Wilcoxon rank sum tests ($\alpha = 0.05/3$) revealed that firing rate correlations in the Same Maze group were significantly higher than the Different Route group ($p = 0.0006$), but did not differ from the Same Route group ($p = 0.574$). The firing rate correlations in the Same Route group were higher than those of the Different Route group, but these differences were not significant ($p = 0.036$). These results indicate that maze-responses were

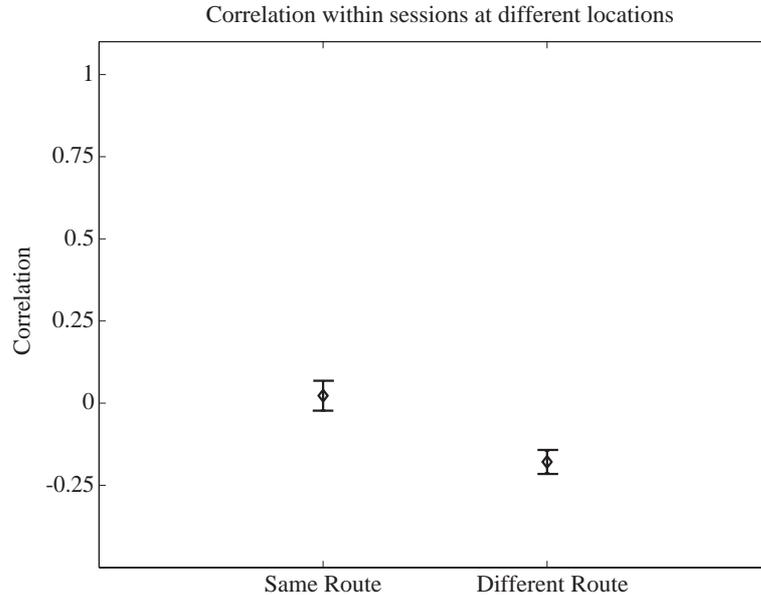


FIGURE 5.7: Distribution of firing rate correlations for similar and dissimilar routes. For 56 maze-responsive PFNs which had phasic-firing fields on the turn sequence, the firing pattern and path of the rat in the PFF was correlated with the firing pattern and path of the rat in windows shifted over the length of the turn sequence. The Same Route group includes firing rate correlations from cases where the rat's path was correlated by at least 0.85, indicating that the rat ran through a very similar route in both windows. The Different Route group includes firing rate correlations from cases where the rat's path was correlated by less than -0.85, indicating that the rat ran through a very dissimilar route in both windows. There was no bias in either group towards positive correlations, indicating that these maze-responsive PFNs did respond entirely on the basis of the shape of the rat's route. Bars represent mean and SEM across cells. Figure from Schmitzer-Torbert and Redish (2004).

highly similar when the same turn sequence was presented, and were not similar when rats took a different route through the same physical location in the environment. Based on the results of the Same Route group, our data also indicate that maze-responses did not purely encode sensory-context/action relationships: when rats ran through similar paths in the same physical locations in the environment, but ran a different turn sequence, correlations were intermediate between both the Same Maze and Different Route groups. The Same Route group contained data from a small number of cells and it could be the case that some PFNs were responding to purely sensory-context/action relationships while other PFNs were further modulated by the specific sequence of turns presented. Maze-responses may thus have reflected a combination of information related to the specific actions performed, the sensory environment those actions were performed in, and in some cases the specific turn sequence presented.

Reward-responsive PFNs

Of the PFNs which fired at least 100 spikes a session, 21% were classified as reward-responsive. Figure 5.9 shows an example of a reward-responsive PFN which was active following arrival at the first food delivery site, but not at the second food delivery site. Of reward-responsive PFNs, 31 (38%) had a significant response only at the first food delivery site, 33 (41%) had a significant response only at the second food delivery site, and 17 (21%) had a significant response at both food delivery sites. Since 79% of the reward-responsive PFNs were responsive at only one of the food delivery sites,

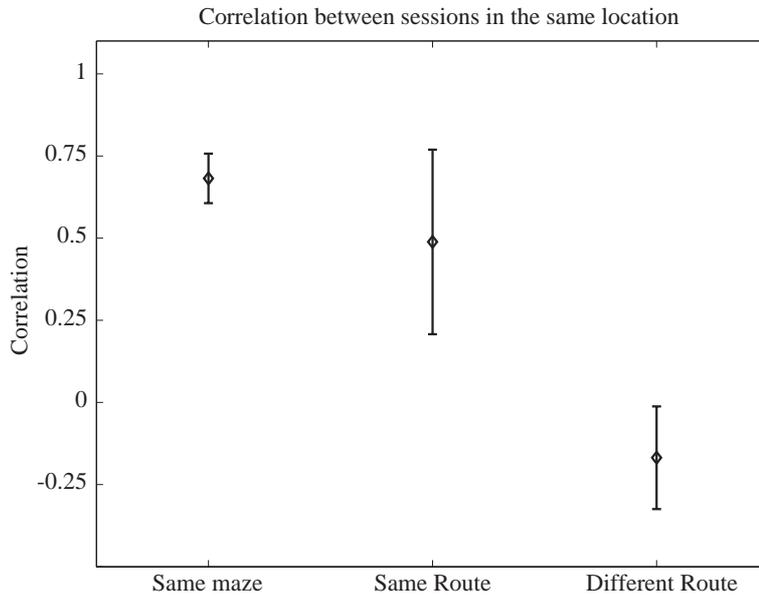


FIGURE 5.8: Correlation of maze-responses between sessions. For 22 maze-responsive PFNs which were observed in at least two sessions and had a phasic-firing field (PFF) on the turn sequence, the correlation of the firing rate observed in the phasic firing field on each session was correlated with the firing rate in the other session at the same location on the maze. These correlations were divided into three groups: 1) Same Maze: correlations obtained from pairs of sessions in which rats ran the same sequence of turns, 2) Same Route: correlations obtained from pairs of sessions in which rats ran different sequences of turns, but had similar paths in the region of the phasic-firing field, and 3) Different Route: correlations obtained from pairs of sessions in which rats ran different sequences of turns, and had dissimilar paths in the region of the phasic-firing field. There was no significant difference between the Same Maze and Same Route groups. The Same Maze condition was significantly higher than the Different Route group, and there was a nonsignificant trend for the Same Route group to be more highly correlated than the Different Route group. Bars represent mean and SEM across cells. Figure from Schmitzer-Torbert and Redish (2004).

these cells did not encode general aspects of food retrieval or consumption (e.g. chewing), which occurred at both food delivery sites.

Maze- and Reward-responsive PFNs are separate populations

In Figures 5.4 and 5.5, the maze-responsive PFNs were not active following arrival at either food delivery location, and in Figure 5.9, the reward-responsive PFN was not active while the rat was running on the maze. This segregation of maze-responses and reward-responses was a characteristic of the entire task-responsive PFN population. Based on the proportions of PFNs which were maze-responsive (27.3%) or reward-responsive (20.5%), we would expect that 5.6% of the PFNs (approximately 22 cells) would have been responsive to both maze and reward if the probability of being a maze-responsive PFN and a reward-responsive PFN was independent. This was not the case: only 1/395 PFNs (0.3%) was classified as both reward-responsive and maze-responsive. This is significantly less than we would expect by chance ($\chi^2(3) = 35.0, p < 0.001$).

Figure 5.10 shows the average normalized firing rate of PFNs following the arrival of the rat at the food delivery sites (*top*) and on the turn sequence (*bottom*). Reward-responsive PFNs were more active following food-delivery than the entire population of PFNs, while maze-responsive PFNs were more active on the maze than the entire population of PFNs. These results follow from the definitions of reward- and maze-responsiveness. However, reward-responsive PFNs were also less active on the maze than either maze-responsive PFNs or the entire PFN population. Likewise, maze-responsive

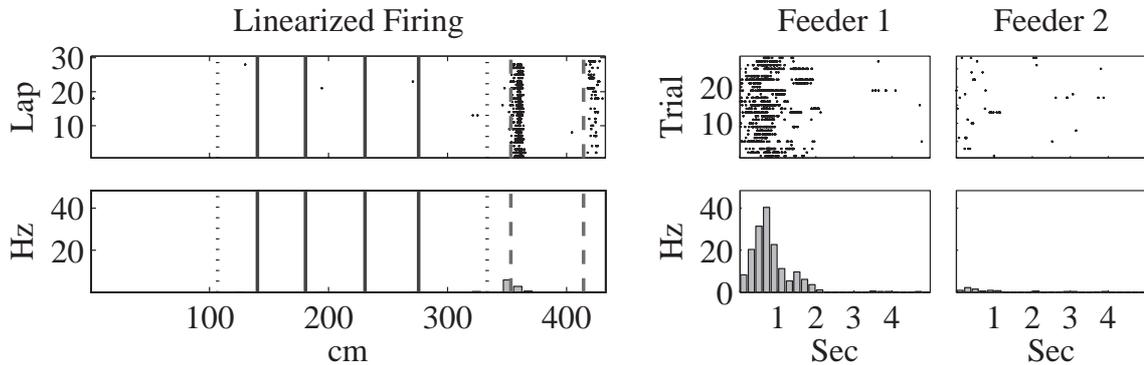


FIGURE 5.9: A reward-responsive PFN. *Left:* Rastergram and histogram of linearized firing on the maze. Key as shown in Figure 5.4. *Right:* Peri-event time histograms (PETHs) of the firing rate of the cell as the rat arrived at the first and second food delivery sites. This PFN had a large response at the first food delivery site, but not at the second food delivery site. (R011-2002-02-16-TT09-03 Maze = LRRR Trials = 28). Figure from Schmitzer-Torbert and Redish (2004).

PFNs were less active following food delivery than either reward-responsive PFNs or the entire PFN population. Our analyses allowed PFNs to be classified as both reward- and maze-responsive, but cells predominantly responded to one or the other parameter, not both. As such, these results further indicate that reward- and maze-responsive PFNs were separate populations of cells.

The differences between groups of PFNs were significant. In the five seconds following food delivery, there was a significant effect across animals of group (All PFNs, Reward-responsive, Maze-responsive, $F_{(2,12)} = 31.2$, $p < 0.001$). Post-hoc comparisons (Tukey-Kramer HSD, $\alpha = 0.05$) revealed that the activity of maze-responsive PFNs was significantly less than the entire population of PFNs. Similarly, on the maze, there was a significant effect of group (All PFNs, Reward-responsive, Maze-responsive, $F_{(2,12)} = 48.6$, $p < 0.001$). Post-hoc comparisons (Tukey-Kramer HSD, $\alpha = 0.05$) revealed that the activity of reward-responsive PFNs was significantly less than the entire population of PFNs.

These results indicate a separation of information processing such that PFNs which responded while the rats ran on the maze did not respond during food receipt and PFNs which responded during food receipt did not respond while the rats ran on the maze.

5.2.4 Tonic-firing neurons

Spatial correlates

While maze-responsive PFNs which were active as rats ran on the Multiple T maze responded in one or more locations, this pattern of response was rare in TFNs. Many TFNs exhibited robust spatial firing rate modulations, with the predominant spatial responses being multiple activations at a relatively constant spatial frequency. These activations were best characterized as *oscillatory*, and could be separated into high frequency (narrow spatial oscillations, 2-3.5 Hz oscillations in the temporal domain, see Figure 5.11) and low frequency oscillations (broad spatial oscillations, < 1 Hz oscillations in the temporal domain, see Figure 5.11). In both examples, the cells' activity was oscillatory as rats were locomoting on the maze, and persisted only briefly following the delivery of food, as rats approached the food delivery location.

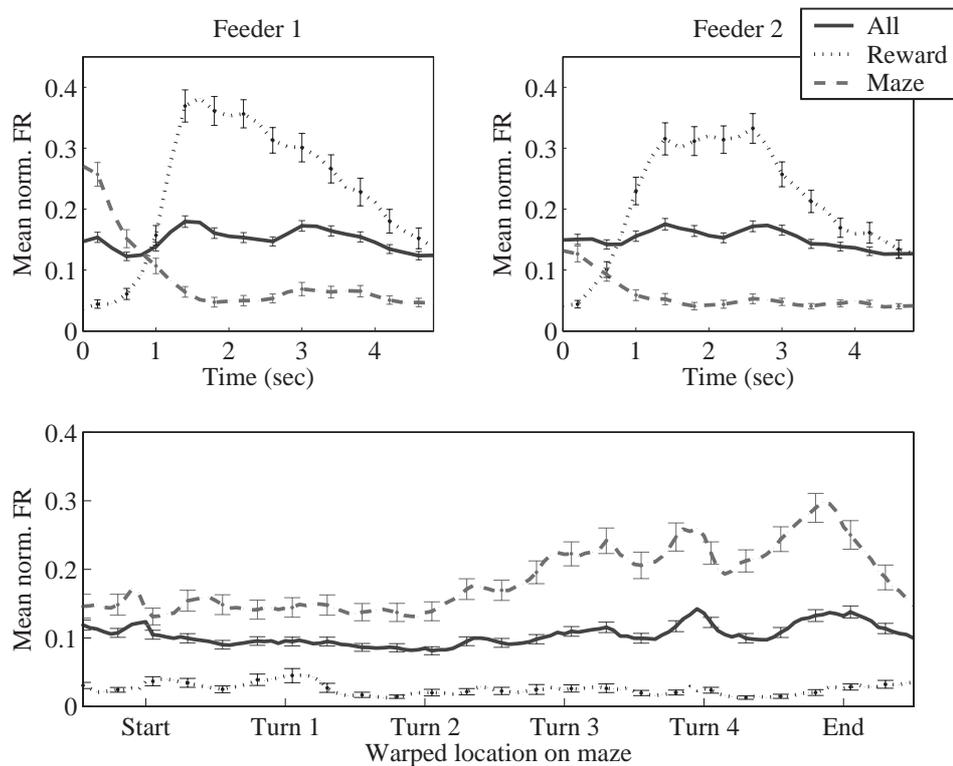


FIGURE 5.10: Segregation of task responsive PFNs. Shown are the average normalized response of PFNs to the delivery of food (*top*) and on the turn sequence (*bottom*). To make comparisons across sessions (i.e. across different maze configurations), the response of each PFN on the turn sequence was warped to a fixed number of bins between landmarks (see Methods for a description of the warping process). For all plots, the entire PFN population is plotted with a solid line, maze-responsive PFNs with a dashed line, and reward-responsive PFNs with a dotted line. Following the delivery of food, maze-responsive PFNs were inhibited relative to reward-responsive PFNs and the entire PFN population. On the turn sequence, reward-responsive PFNs were inhibited relative to maze-responsive PFNs and the entire PFN population. Figure from Schmitzer-Torbert and Redish (2004).

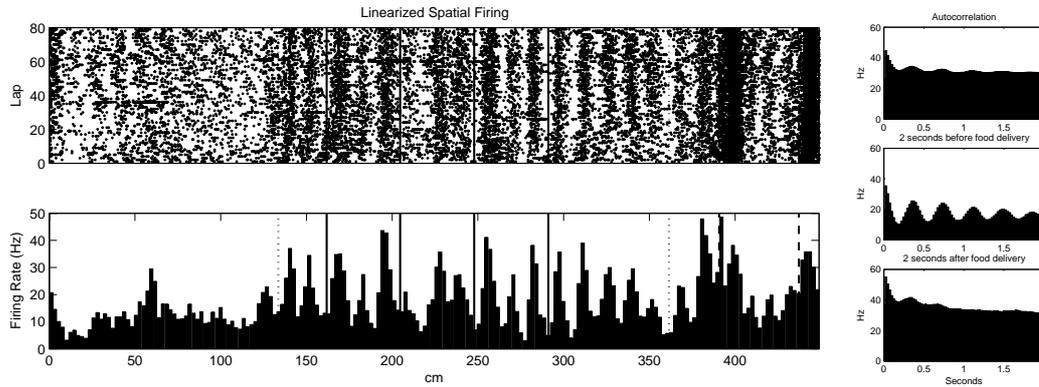


FIGURE 5.11: A TFN-1 with fast spatial oscillations. *Left:* Rastergram and histogram of linearized firing on the maze. Vertical lines indicate landmarks on the maze, including each of the four turns (solid), the two food delivery sites (dashed) and the points at which the rat entered and exited the turn sequence (dotted). Prominent oscillations were observed as the rat was running on the maze. *Right:* The autocorrelation of this cell over the entire session (*top*), in the 12 seconds preceding food delivery (*middle*) and in the 12 seconds following food delivery at the first food delivery site (*bottom*). As suggested in the rastergrams, this cell had oscillatory activity specifically while running on the maze. The residual oscillatory activity seen in the raw autocorrelation and following food delivery is due to contamination by oscillatory activity as the rat is running on the maze. (R023-2002-09-01-TT01-01 Multiple T Maze = LRL Laps = 82 Trials = 80)

To characterize the high-frequency spatial responses of TFNs, autocorrelations and Fourier spectra for each TFN were calculated with spikes taken from 1) the entire session, 2) a window of time preceding food delivery at the first food delivery site, or 3) a window of time following food delivery at the first food delivery site. In general, strong high frequency oscillations were present in TFN-1s as rats were running on the Multiple T maze (see Figure 5.11). High frequency oscillations were not characteristic of TFN-2s or TFN-3s in any temporal window, and TFN-1s did not show high frequency oscillations following food delivery.

To compare power spectra from different TFNs, each spectra was first normalized by the total power. Then, oscillation scores were derived for high and low frequency oscillations by finding the difference in total power in the high frequency (2-3.5 Hz) band relative to the total power in a control band (3.5-5 Hz), and also the difference in total power in the low frequency (0.5-1 Hz) band relative to the total power in a control band (1.5-2 Hz).

As shown in the average Fourier spectra plotted in Figure 5.13, power in the high frequency (2.5-3 Hz) and low frequency (0.5-1 Hz) bands was enriched in the 3 seconds preceding food delivery. An ANOVA comparing the three temporal windows (Preceding food delivery, following food delivery, or the entire session) indicated that power in the 2-3.5 Hz band across all TFNs was specifically enriched preceding food delivery compared to spectra calculated over the entire session or following food delivery ($F_{(2,1761)} = 35.6$, $p < 0.0001$, followed by Tukey-Kramer post-hoc comparisons). Preceding food delivery, power in the 2-3.5 Hz band was significantly greater in TFN-1s than in TFN-2s or TFN-3s ($F_{(2,585)} = 79.8$, $p < 0.0001$, followed by Tukey-Kramer post-hoc comparisons). An ANOVA comparing the three temporal windows (Preceding food delivery, following food delivery, or the entire session) indicated that power in the 0.5-1 Hz band across all TFNs was specifically enriched preceding food delivery compared to spectra calculated over the entire session or following food delivery ($F_{(2,585)} = 36.4$, $p < 0.0001$, followed by Tukey-Kramer post-hoc comparisons). In the 3 seconds preceding food delivery, TFN-2s had significantly more power in the 0.5-1 Hz band

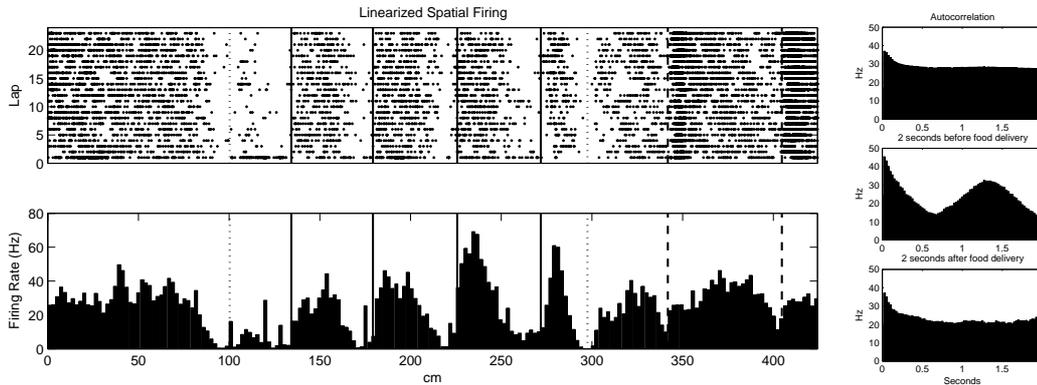


FIGURE 5.12: A TFN-2 with low frequency spatial oscillations. *Left:* Rastergram and histogram of linearized firing on the maze. Key as in Figure 5.11. Prominent modulations in the cell's firing rate were observed as the rat ran on the maze, with strong activations after each turn. Across the session, these oscillations were better related to the location of the rat than temporal parameters, such as the time of food delivery. *Right:* The autocorrelation of this cell over the entire session (*top*), in the 12 seconds preceding food delivery (*middle*) and in the 12 seconds following food delivery at the first food delivery site (*bottom*). This cell had oscillatory activity specifically while running on the maze. (R010-2001-12-11-TT12-01 Multiple T Maze = RRRR Laps = 28 Trials = 24)

than TFN-1s or TFN-3s ($F_{(2,1761)} = 22.9$, $p < 0.0001$, followed by Tukey-Kramer post-hoc comparisons). These results demonstrate that the patterns of oscillations shown in the example TFNs (Figures 5.11 & 5.12) are typical of the population of TFNs: oscillations were present before food delivery, but were not observed following food delivery or in the total session data. These oscillations were transient events that were only during navigation. The frequency of oscillation was cell type specific, with TFN-1s having narrow spatial oscillations and TFN-2s having broad spatial oscillations.

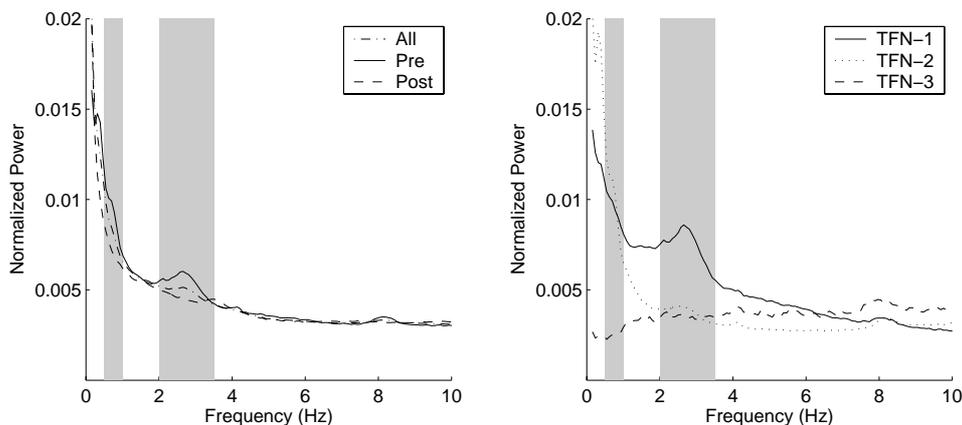


FIGURE 5.13: TFN oscillations. Averaged Fourier spectra were calculated from TFN spike trains and normalized by total power. *Left:* TFNs had increased power in the 2-3.5 Hz range on the Multiple T maze in the 6 seconds preceding (Pre), but not in the 6 seconds following food delivery (Post), or across the entire recording session (All). *Right:* increased power in the 2-3.5 Hz range was restricted to TFN-1s, while increased power in the 0.5-1 Hz range was found in TFN-2s.

Event correlates

In addition to pronounced spatial correlates, many TFN-1s and -2s were responsive following food delivery on the Multiple T maze. In contrast to PFN responses, which consisted of primarily phasic excitations, TFN responses tended to have broad activations or inhibitions (see Figures 5.14 – 5.15 for examples).

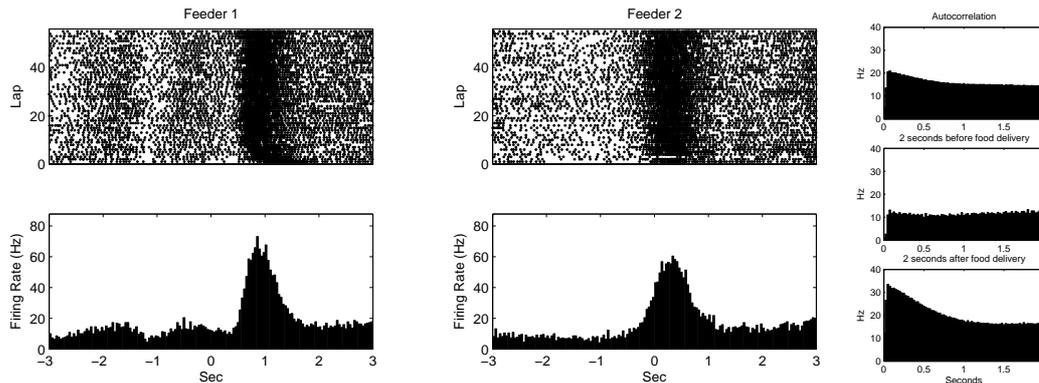


FIGURE 5.14: A TFN-2 with event-related firing on the multiple T task. *Top:* Peri-event time histograms (PETHs) of the firing rate of the cell relative to food delivery at the first and second food delivery sites). *Bottom left:* Rastergram and histogram of linearized firing on the maze. Key as in Figure 5.11. Following food delivery at each food delivery site, this cell had a phasic activation lasting ~ 1 second. *Right:* The autocorrelation of this cell over the entire session (*top*), in the 12 seconds preceding food delivery (*middle*) and in the 12 seconds following food delivery at the first food delivery site (*bottom*). The cell did not show spatial responses or oscillatory activity. (R023-2002-08-20-TT03-04 Multiple T Maze = LLRR Laps = 60 Trials = 57)

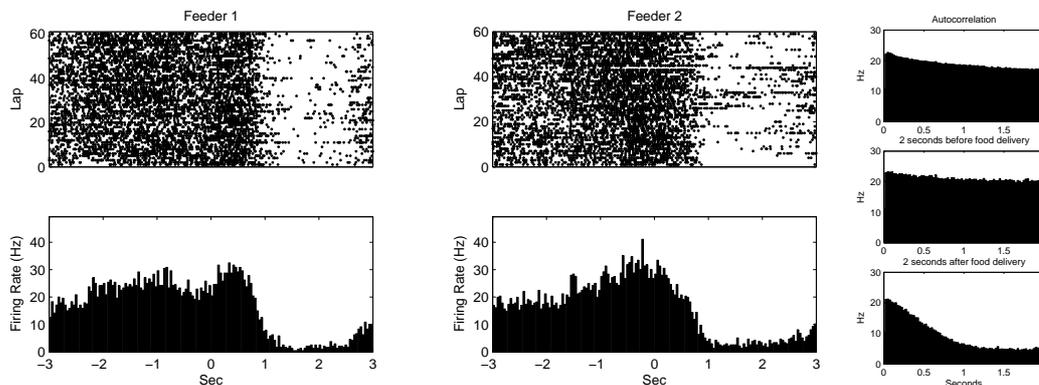


FIGURE 5.15: A TFN-2 with event-related firing on the multiple T task. *Top:* Peri-event time histograms (PETHs) of the firing rate of the cell relative to food delivery at the first and second food delivery sites). *Bottom left:* Rastergram and histogram of linearized firing on the maze. Key as in Figure 5.11. Following food delivery at each food delivery site, this cell had a phasic activation lasting ~ 1 second. *Right:* The autocorrelation of this cell over the entire session (*top*), in the 12 seconds preceding food delivery (*middle*) and in the 12 seconds following food delivery at the first food delivery site (*bottom*). As suggested in the rastergrams, this cell did not display prominent oscillations. (R016-2002-05-07-TT07-01 Multiple T Maze = LRLl Laps = 65 Trials = 61)

5.2.5 Striatal representation of task parameters

In the Multiple T task, both phasic and tonic striatal neurons were modulated by navigation and event-related parameters, such as food delivery and spatial location. With these types of tuning to

task parameters, an obvious question is how well each task parameter was represented by striatal neural activity. To address this question, Bayesian reconstruction techniques were applied to the striatal ensembles recorded as rats performed the Multiple T task. Task parameters of interest came from navigation and food delivery. The primary questions of interest were, given the firing rate of striatal ensembles, how well could each of these task parameters be reconstructed? This question was addressed by examining the probability of reconstructing a task parameter given an ensemble firing pattern ($P(X|F_t)$), where X is the task parameter and F_t is a vector representing the firing rate of the striatal ensemble at time t . A second measure used to assess the striatal representation was the *reconstruction quality* (RQ), which was defined as the probability that the correct value of the task parameter was reconstructed on the basis of the ensemble firing pattern.

During navigation, many striatal neurons had striking patterns of spatial modulation on the Multiple T. These navigation responses were reflected as specific tunings of striatal neurons to spatial location, sequence progress, and/or velocity. How well were each of these variables represented by striatal neural ensembles? Bayesian reconstruction was attempted with respect to spatial location and sequence progress. As shown in Figure 5.16, the spatial location of rats running the Multiple T maze was well represented in striatal ensembles, as is evident in the plot of $P(X|F_t)$, where a clear diagonal in that plot indicates that at each spatial location, the reconstructed location of the animal was preferentially located at the animal's true position.

As many striatal neurons recorded on the multiple T task have spatial tuning, then it should be possible to *reconstruct* the location of the animal from the neural firing patterns. This technique provides a means to test how well the neural population represents the information about its external environment (Rieke et al., 1997). Reconstruction of external information from neural ensembles has been used to reconstruct position in an environment from hippocampal place cell firing, (Wilson and McNaughton, 1993; Brown et al., 1998; Zhang et al., 1998; Jensen and Lisman, 2000) reaching direction from motor cortical neurons, (Georgopoulos et al., 1983; Salinas and Abbott, 1994) and orientation from postsubicular head direction cell ensembles (Johnson et al., 2003).

Reconstruction of spatial location was calculated from dorsal striatal neurons using Bayesian reconstruction techniques. Fig. 5.18 shows the average reconstruction of sequence progress over all animals over all laps. On average, there was a high probability of reconstructing the animal's location on the track on the basis of the neural activity. This indicates that the spatial information contained in striatal ensembles was rich enough to produce a clean reconstruction of the animal's location on the track.

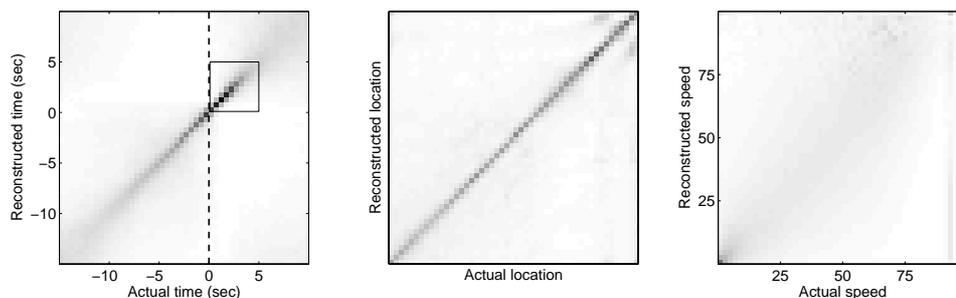


FIGURE 5.16: Reconstruction of task parameters in the Multiple T task. Ensembles of simultaneously recorded striatal neurons were used to reconstruct the time of reward delivery, the spatial location of rats, and the running speed of rats on the Multiple T maze. There was a high-quality representation of the five second interval following the delivery of food, and the spatial location of the rat, but not of running speed.

While there was a high quality representation of the temporal window in which rats received their food rewards on the Multiple T task, the ensemble representation of reward delivery was distinct for food deliveries occurring at each pellet dispenser. On the Multiple T task, there was a significant reduction in the reconstruction quality (RQ) of the five seconds following the delivery of food when Bayesian reconstruction of the time of reward delivery at one pellet dispenser was performed using tuning curves derived from food deliveries which occurred at the other pellet dispenser (see Figure 5.17, *left, middle*). The differences in reconstruction quality between the two conditions (Figure 6.12, *right*) were significant ($t(16) = 4.3$, $p = 0.0006$).

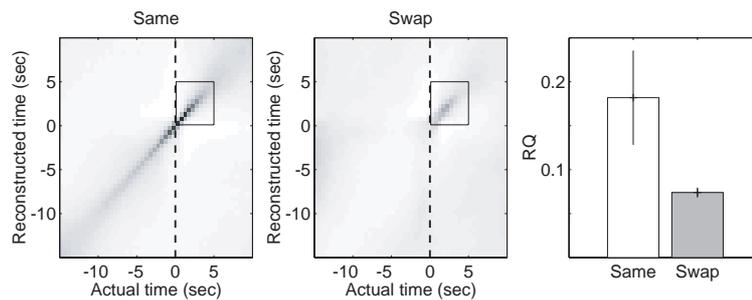


FIGURE 5.17: Similarity of reward-receipt encoding in the Multiple T task. Bayesian reconstruction was used to determine the similarity of the striatal ensemble representation in the 5 second window following the delivery of food occurring at different pellet dispensers. When the tuning curves used in the Bayesian reconstruction were derived from rewards delivered at the same pellet dispenser (*Left, Same*), there was a high quality representation of the five seconds following the delivery of food, as was the case in Figure 5.16. However, when tuning curves were derived from rewards delivered at different pellet dispensers (*Middle, Swap*), there was a sharp reduction in the quality of the reconstruction. There was a significant difference in the reconstruction quality (RQ) between the two conditions (*Right, error bars represent mean and 95% confidence intervals over rats*).

5.2.6 Neural learning correlates

Although the time period in which rewards were received was well-represented in both the Multiple T and the Take 5 tasks, navigation-related variables were only well-represented by striatal ensembles in rats performing the Multiple T. Further, spatial location, but not velocity, of rats performing the Multiple T task was able to be extracted with high quality. As described above for PFNs (see Section 5.2.3), the spatial responses of maze-responsive PFNs were related to not only the spatial position of the animal, but also to the actions performed at that location, and the sequence of turns rats were performing in a given session. Therefore, the representation of spatial location described in the present chapter was evidence of a spatial-action-sequence encoding by striatal ensembles. Given that in many cases rats were performing novel maze sequences, which they had never experienced before, how did the striatal spatial representation develop, and how does the development of the striatal representation relate to the behavior of the animals as they are learning to navigate through the maze? To address this question, the behavioral maze-learning of rats performing the Multiple T maze was compared to the development of the striatal spatial representation.

Reconstruction of spatial location.

Reconstruction techniques allow the exploration of changes in representation as a function of experience (Georgopoulos et al., 1988; Jackson and Redish, 2003; Redish et al., 2000). In particular, measuring the changes in our ability to reconstruct a signal can allow us to estimate the rate at which

the neural representation is changing. To examine how the success of our reconstruction depended on experience, the quality of the reconstruction was examined during task performance on the multiple T Reconstruction quality, as described in the Methods, was defined as the average likelihood of reconstructing the correct position of the animal in each lap of the task. As can be seen in Fig. 5.18, this reconstruction quality changed from lap-to-lap, becoming better with each subsequent lap. This implies that the striatal representation changed across laps, even within a single session.

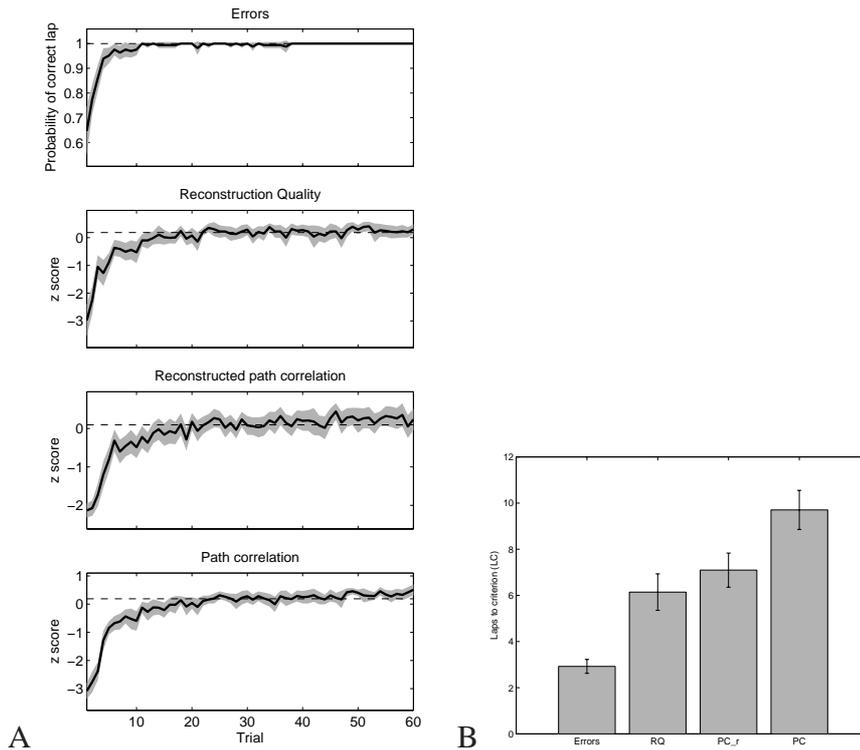


FIGURE 5.18: Comparison of behavioral and neural learning rates.

An improvement in our ability to reconstruct the animal's location could imply a functional remodeling of the responsiveness of striatal neurons during changes in performance. Indeed, the spatial tuning of many striatal neurons changed as a function of experience. Two examples (one phasic neuron and one tonic neuron) are shown in Figure 5.19. Such changes in spatial tuning are likely the basis for an improvement in our ability to reconstruct the position of the animal. An important question to address is whether or not these changes in spatial tuning are influencing the development of our behavioral measures, or are instead produced by changes in behavior. To exclude the possibility that changes in neural responsiveness were somehow confounded with changes in behavior, the number of laps to reach criterion on the reconstruction quality measure (LC_{RQ}) was compared to LC_{error} and LC_{PC} . As with the linearized path correlation, LC_{RQ} for each session was defined as the first lap in which the reconstruction quality was within one standard deviation of its mean value in the last five laps of the session. LC_{RQ} was significantly larger than LC_{error} and smaller than LC_{PC} (From 105 sessions, $F(2,312) = 37.01$, $p < 0.001$, post-hoc comparisons using Tukey-Kramer HSD). These results indicate that the improvement in reconstruction quality followed error-reduction but preceded the development of a stable route.

To make a more direct comparison between the development of a stable route and changes in neural responsiveness, the path correlation measure was compared to *reconstructed path correlation*. The reconstructed position of the rat at any point in time was taken to be the most probable location of

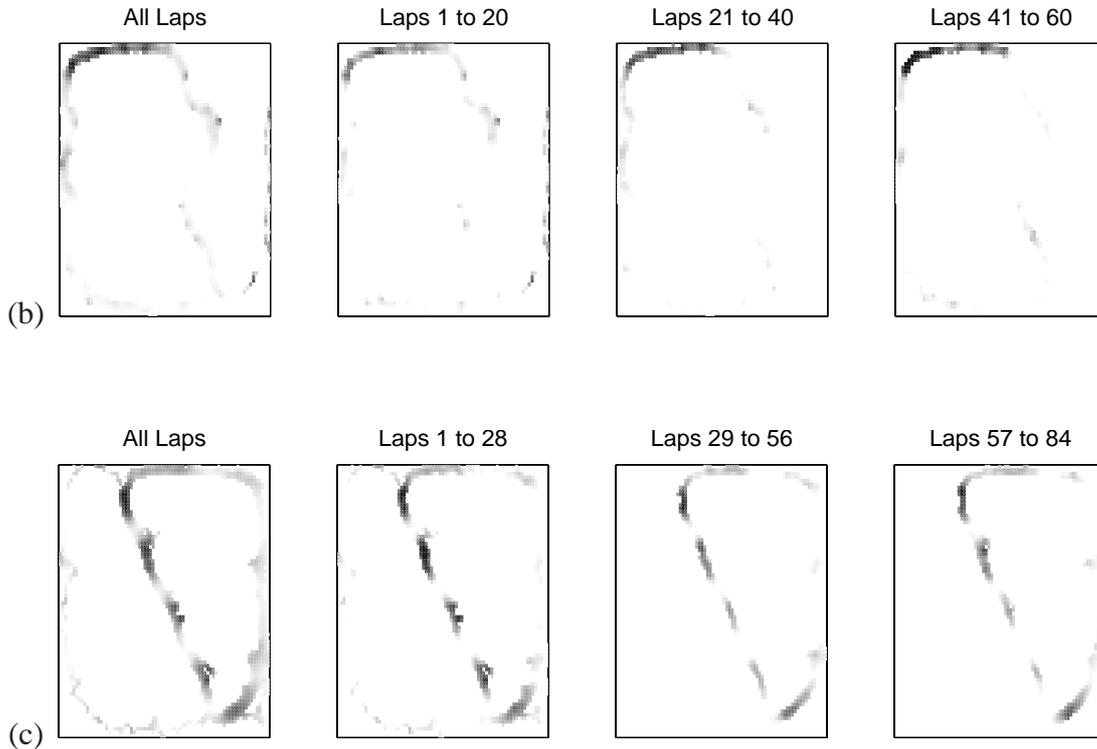


FIGURE 5.19: Changes in spatial tuning as a function of experience. Spatial tuning of striatal neurons changed as a function of experience. Some dorsal striatal neurons became more responsive, (A, example phasic neuron) while others decreased their spatial responsiveness (B, example tonic neuron). Data from R023-2002-08-27-TT07-02 (a) and R010-2001-12-12-TT07-02 (b).

the animal based on the current firing rate of the neural ensemble. The set of reconstructed positions defined a reconstructed path, which could be treated in the same fashion as the actual position data. Figure 5.18 shows the changes in correlation of the actual path and the reconstructed path. Also shown is the linearized path correlation and the linearized reconstructed path correlation, which have been adjusted to similar asymptotes by subtracting the average correlation in the last five laps of the session. Qualitatively, the linearized reconstructed path correlation improved at a faster rate than did the linearized path correlation. To make a more precise comparison, the number of laps to criterion for the linearized reconstructed path correlation (LC_{rPC}) was calculated in the same manner as LC_{PC} . As was the case with reconstruction quality, LC_{rPC} was significantly larger than LC_{error} and smaller than LC_{PC} (From 105 sessions, $F(2,312) = 37.01$, $p < 0.001$, post-hoc comparisons using Tukey-Kramer HSD). These results indicate that the reconstructed path adopted a stable shape before the rat adopted a stable path through the maze. Such a result indicates that changes in neural activity in the dorsal striatum are well-positioned to guide the development of a stable route.

5.3 Discussion

The main results from the Multiple T task are as follows: 1) rats learned the task and performed well even when learning new (i.e. novel) maze sequences, 2) rats demonstrated two changes in behavior in each training session, a fast reduction in the number of errors made and a slower process of path refinement, 3) both phasic and tonic-firing striatal neurons demonstrated robust behavioral correlates, 4) striatal ensembles robustly encoded the location of rats on the maze, and the interval

in which food rewards were received, 5) the development of the striatal representation of space developed with a rate that was intermediate between the behavioral learning rates (error-elimination and path-refinement).

5.3.1 Phasic firing neurons

These data demonstrate that phasic firing neurons (PFNs) in the rodent striatum are active during navigation and food consumption in a sequential navigation task. In the Multiple T task, task-modulated striatal neurons were active either during navigation (maze-responsive) or food receipt (reward-responsive), but not both. These results support a strong segregation of striatal projection neuron activity, such that separate populations of striatal neurons are active during navigation and reward delivery. Both maze- and reward-responsive neurons were often context-dependent: maze-responsive PFNs were active in one or more locations on the maze while reward-responsive PFNs were active at one or both food delivery sites. Maze-responsive neurons which were recorded across successive days (on either the same maze sequence or different maze sequences) demonstrated a preference for actions and spatial location in addition to the specific sequence of turns rats were performing.

Behavioral correlates

Lesions and inactivations of the dorsal striatum in rodents impairs performance of habitual, stimulus–response (S–R) tasks (Kesner et al., 1993; McDonald and White, 1993; Packard et al., 1989; Packard and McGaugh, 1992; Packard, 1999), as well as longer chains of sequential behavior (Berridge and Whishaw, 1992; Cromwell and Berridge, 1996; Matsumoto et al., 1999; Miyachi et al., 1997). One theory of how the dorsal striatum learns and produces S–R behavior is that striatal projection neurons with connections to motor centers encode S–R relationships by responding specifically to complex cortical inputs (Graybiel et al., 1994; Graybiel, 1998). A number of studies have shown that striatal neurons have highly specific responses to task parameters which could encode stimulus–response relationships. Studies in the rat (Carelli et al., 1997; Gardiner and Kitai, 1992), and primate (Kimura, 1986, 1990; Kermadi et al., 1993; Kermadi and Joseph, 1995; Tremblay et al., 1998) have found that the responses of striatal cells often depend on behavioral context. In the rat, for example, Gardiner and Kitai (1992) report that cells in the dorsal striatum which responded to an auditory cue during a movement task usually did not respond to the same cue presented outside of the task, and some cells which responded during head movements during the task did not respond when rats made similar movements outside of the task. Carelli et al. (1997) report that in rats who have learned to barpress for food, dorsolateral striatal cells which responded to movement of the forelimb outside of the instrumental task were not active during lever pressing.

On the multiple T maze, neurons in the dorsal striatum responded as rats navigated the maze and during the delivery of food. Neither maze- nor reward-responses were described by general motor behavior. Less than one third of reward-responsive PFNs responded at similar levels at both food delivery sites, indicating that reward-responses did not simply encode the action of chewing. Maze-responsive PFNs which responded at one location on the maze were not biased to respond similarly at other regions where rats took similar paths, indicating that maze-responses did not simply encode motor activity during navigation. These results are consistent with the studies cited above, which indicated that dorsal striatal neurons correlated with a movement or stimulus during a task are often not active during the same movement or stimulus presentation in a different behavioral context. Within

a session, maze-responses were well related to the spatial location of the animal. However, maze-responses did not encode the spatial position of the animal independent of the animal's actions. First, maze-responses were poorly correlated across sessions when rats took a different path through the same 2-dimensional location in the environment. Secondly, maze-responses were highly correlated across sessions when animals ran the same sequence of turns. Finally, maze-responses were biased toward positive correlations across sessions when animals ran a different sequence of turns but took a similar path through the same 2-dimensional location in the environment. These data indicate that maze-responsive cells were modulated by the location of the animal, what the animal was doing at that location, and to some extent by the specific sequence of actions the rat was performing.

This type of striatal sequence-specificity is consistent with the work done in primates and rats. In primates, Kermadi and colleagues (Kermadi et al., 1993; Kermadi and Joseph, 1995) have shown that striatal neurons in primates preferred specific visuomotor sequences. In rats, Aldridge and Berridge (1998) have shown that dorsal striatal neurons which were active during sequenced grooming were often not active during similar movements occurring outside of grooming sequences. To our knowledge, the data presented here from the multiple T task are the first evidence for sequence-specific striatal activity in rodents performing an arbitrary sequencing task.

Maze-responses were also uniformly distributed over the turn sequence on the maze. If maze-responses encode what actions need to be performed at a particular location/sensory context, then a uniform distribution of maze-responses indicates that the striatal representation is rich enough to specify an action to perform at any point of the task.

Segregation of maze- and reward-responses

Maze-responsive PFNs often responded at multiple locations on the maze and reward-responsive PFNs often responded following arrival at both food delivery sites. However, maze-responsive PFNs did not respond following arrival at either food delivery site and reward-responsive PFNs did not respond while rats were running on the maze. A segregation of maze- and reward-responsive PFNs implies a segregation of information processing in the striatum, and brings up two questions: what is the functional consequence of segregation, and what properties of the striatum produce segregation?

A segregation of maze- and reward-responses may shed light on the computational functions of the striatum. Recent proposals of basal ganglia function suggest that the striatum is involved in selecting appropriate actions in a task by implementing a reinforcement learning algorithm (Barto, 1995; Brown and Sharp, 1995; Daw and Touretzky, 2000; Daw, 2003; Doya, 1999, 2000; Houk et al., 1995; Foster et al., 2000; Montague et al., 1996; Schultz et al., 1997; Sutton and Barto, 1998). In reinforcement learning models of the striatum, the nigrostriatal dopaminergic system provides a reward-prediction error signal, and the striatum implements an actor-critic architecture. The actor is responsible for selecting which action would be appropriate given the current sensory input, while the critic uses the reward-prediction error signal to change the value of sensory inputs so that the most advantageous action will be chosen. Segregation of maze- and reward-responsive PFNs may then reflect the separation of actor and critic components in the striatum.

One possible mechanism of segregation of maze- and reward-responsive PFNs is related to striatal subcompartments. On the basis of markers such as μ -opiate receptor binding, the striatum can be divided into μ -opiate rich *striosomes* (also termed *patch*, Gerfen, 1985) and μ -opiate poor *matrix*, which is rich in acetylcholinesterase (Herkenham and Pert, 1981; Gerfen, 1985; Graybiel and Ragsdale, 1978). Matrix receives inputs from sensorimotor cortex and projects to the substantia nigra pars reticulata and pallidial output nuclei (Gerfen, 1984, 1989; Kawaguchi et al., 1990; Ragsdale and Graybiel, 1984). Striosomes receive input from "limbic" cortex (including infralimbic, prelimbic,

and anterior cingulate cortex) and project to dopaminergic cells in the substantia nigra pars compacta (Gerfen, 1984, 1989; Ragsdale and Graybiel, 1984). With its inputs to the substantia nigra pars compacta, striatal patches are well-suited to be involved in reward-related processing, while matrix is well-suited to be involved in action (Houk et al., 1995; Graybiel, 1998; Kimura, 1995; White, 1989). White and Hiroi (1998) have shown that electrodes placed in striosomes, but not matrix, will support self-stimulation in rats, supporting a relationship between striosomes and reward. Trytek et al. (1996) have shown that motor related neurons tended to be located in the matrix, supporting a relationship between matrix and action. It could be that segregation of maze- and reward-responsive PFNs on the multiple T maze reflects an anatomical segregation of maze-responsive PFNs to the striatal matrix and reward-responsive PFNs to striosomes. The anatomical distributions of maze-responsive and reward-responsive PFNs in the dorsal striatum is an important question to be addressed in future experiments.

5.3.2 Tonic firing neurons

Like PFNs, many tonic firing neurons (TFNs) were modulated by task parameters such as spatial location and reward-delivery. Unlike PFNs, which were highly selective in the locations or events to which they responded, TFNs which were task-modulated often responded over large windows. TFN-1s which were spatially modulated demonstrated narrow, spatially-locked oscillatory activity. In the temporal domain, these spatial oscillations were observed as 2-3.5Hz oscillations observed as rats ran on the maze. Many TFN-2s were also spatially modulated, but demonstrated broad activations, of the size of turns on the maze or larger maze segments. TFN-1s and -2s were also modulated during reward receipt, and the types of modulations observed were more complex than the phasic activations seen in reward-responsive PFNs. TFNs responsive during food delivery demonstrated phasic activations, inhibitions, or patterns of both at one or more food delivery sites.

Of the TFN subtypes, TFN-3s were remarkably unresponsive during the Multiple T task. TFN-3s did not exhibit spatial activations or reliable reward-related responses.

5.3.3 Striatal representation

Beyond the level of individual neurons, ensembles of striatal neurons contained rich behavioral correlates that allowed for the reconstruction of relevant task parameters, such as the spatial location of the rat, and the interval in which food was retrieved and consumed. The ensemble representation of spatial location demonstrated plasticity within a session: the quality of the spatial representation was low early in each session, and grew as a function of the rat's experience on the task. The development of a high quality representation of the rat's spatial location paralleled the development of a stable route through the maze. The changes in striatal representation preceded the development of a stable route, indicating that these changes in neural representation were well-positioned to influence route-stabilization.

Chapter 6

The Take 5 task

6.1 Introduction

In Chapter 5, the results presented from the Multiple T task demonstrated that the firing rate of striatal neurons was well correlated with the animals' behavior in a navigation task. One question that was left unanswered was: *to what degree are maze-responsive PFNs determined by spatial information and sequence information?* On the Multiple T maze, the activations of maze-responsive PFNs were sensitive to not only the spatial location of the rat, the actions performed at that location, but also more global information, such as the specific sequence of turns rats were performing. In the Multiple T maze, information related to the rat's spatial location was always confounded with information related to the progress of the rat through the sequence of actions leading to rewards. Therefore, to explore the encoding of spatial and sequence information, a new task was designed in which these parameters were dissociated. In the Take 5, rats were run on a rectangular track for food rewards that could be delivered to any side of the track (described more fully in Methods, see Figure 3.2). On each trial, rats ran clockwise $\frac{5}{4}$ the length of the track. Each trial consisted of 5 right turns, or 5 pellet dispenser-to-pellet dispenser movements, and the animal's goal location rotated on a trial-by-trial basis. The Take 5 task allowed for a dissociation of *spatial location* from *sequence progress*: the location of the rat in the two-dimensional environment did not accurately predict the next rewarded pellet dispenser unless the rat was able to take into account its progress through the set of 5 right turns leading to reward.

6.2 Results

6.2.1 Behavior

Six rats were trained on the Take 5 task, and all were able to learn the task successfully. Across 58 sessions on the Take 5 task (7 to 15 sessions per rat), rats completed 44.6 ± 4.9 trials per session. Although rats were able to earn a large number of rewards on the task, this performance alone did not indicate the nature of the behavioral strategy rats employed in order to complete trials on the task. For instance, rats could be completing Take 5 trials on the basis of sensory cues (the sounds of food delivery and the tone which signaled food delivery) or rats could have developed a cognitive strategy in order to predict the location of the upcoming reward location were food was expected on each trial. To determine the strategy that rats employed on the Take 5 task, probe trials were included in which there was an omission of either the tone predicting food delivery or food delivery

itself, or both. As shown in Figure 6.1, rats selectively sampled the rewarded food ports in normal trials: Across six animals, there was a high probability that rats would sample the food port that was rewarded in normal trials (Tone+/Food+, in which both the tone predicting food delivery and food delivery itself occurred), but rats had a low probability of sampling other food ports that they ran past on each trial on their way to the rewarded food port. In probe trials, rats were biased to sample the food port in which reward was expected: on every type of probe trial, rats were biased to sample the food port at which food should have been expected. This food port sampling bias was observed even for probe trials in which all sensory cues were eliminated (Tone-/Food-, in which the food prediction tone and food delivery were both omitted) demonstrating that rats were performing this task using a non-sensory strategy to determine where each food delivery was expected. Sensory cues also played a role in the performance of rats on the Take 5 task: when all sensory cues were eliminated (Tone-/Food- probe trials), rats were less likely to sample the correct food port than in normal trials (Tone+/Food+ trials).

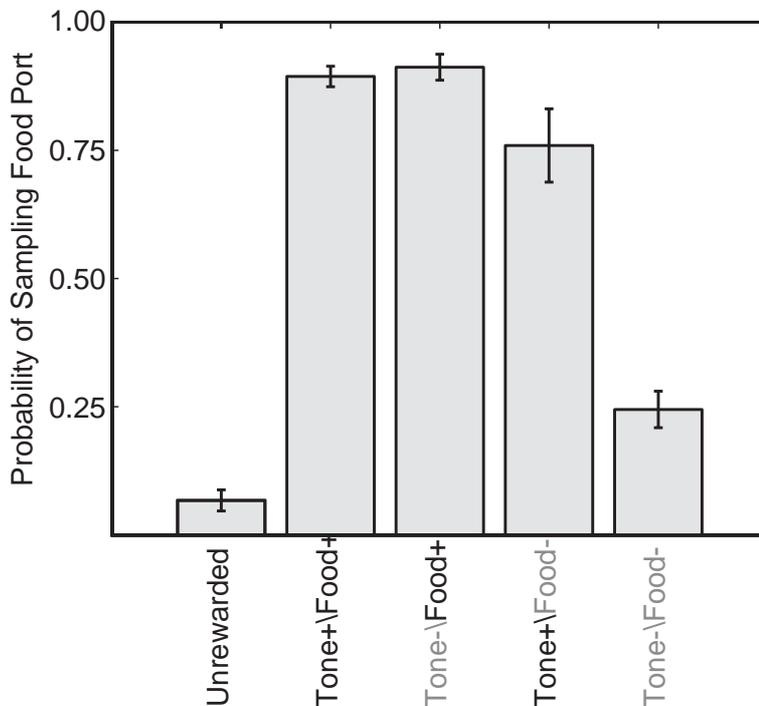


FIGURE 6.1: Food port sampling in the Take 5 task. Bars indicate the average probability of sampling food ports following the rats' arrival at the pellet dispensers. Data was averaged first within, then across, six rats, and bars represent standard error of the mean. Rats rarely sampled the food ports when food was not expected (No reward) and reliably sampled the food ports following normal food delivery (Tone+/Food+). On probe trials, there was an omission of either the tone preceding food delivery (Tone-/Food+), food delivery (Tone+/Food-), or both (Tone-/Food-). On all three types of probe trials, rats sampled the food ports more often than when no food was expected.

6.2.2 Neurophysiology

Neural data was available from 5 of the six animals described above which were implanted with hyperdrives over the dorsal striatum. In the sixth rat, the final tetrode locations included several electrodes placed in or near the border between the striatum and the pallidum, and data from this rat was therefore excluded. From 50 sessions, 784 spike trains were obtained (156.8 ± 41.5 spike trains

per animal). Of the 784 spike trains, 487 (62%) were classified as phasic-firing neurons (PFNs). The remaining 297 (38%) were classified as tonic-firing neurons (TFNs).

6.2.3 Phasic-firing neurons

Task responses

Maze- and reward-responsive PFNs were separate populations. Of the 201 task-responsive PFNs, 10 were classified as both maze-responsive and reward-responsive. This significantly fewer than would be expected based on a random distribution of maze- and reward-responsiveness across the set of PFNs (22 expected based on chance, difference significant by $\chi^2(3) = 10.4$, $p = 0.016$).

Maze-responsive PFNs

A total of 82 of 487 PFNs were responsive at rats ran on the maze (a typical maze-responsive PFN is shown in Figure 6.2). Considering the responses of PFNs with respect to spatial location on the maze, and to the sequence position of the rat in the steps leading to reward on each trial, a small number of PFNs were significantly modulated by either spatial location or sequence progress, but the majority of neurons were modulated by a combination of both spatial location and sequence progress. For instance, the maze-responsive PFN shown in Figure 6.2 responded as the rat ran between each set of food delivery sites, with an obvious preference for movements between the third and fourth food delivery site. In addition to this spatial preference, this cell had a strong dependence on the sequence of actions leading to reward. At the bottom of Figure 6.2, this PFN had a strong increase in activity as the approached the rewarded food delivery site.

Similar patterns of modulation were seen in other maze-responsive PFNs recorded on the Take 5 task. In Figure 6.3, the spatial and sequence tuning of three maze-responsive PFNs are shown, including the PFN shown in Figure 6.2.

Reward-responsive PFNs

A total of 129 of 487 PFNs were responsive in the five seconds following the delivery of food (a typical reward-responsive PFN is shown in Figure 6.4).

6.2.4 Tonic-firing neurons

Spatial correlates

In the Multiple T task, two TFN subtypes exhibited strong spatial modulations as rats navigated through the maze. TFN-1s often exhibited narrow spatial oscillations, which were detected as temporal oscillations in the 2.5-3 Hz band during the 6 seconds preceding food delivery. TFN-2s often exhibited broad spatial oscillations, which were detected as temporal oscillations in the 0.5-1 Hz band during the 6 seconds preceding food delivery. While these patterns of spatial modulations were relatively common on the Multiple T maze, they were not detected in any TFN subtype recorded from rats running the Take 5 task. As was done for TFNs on the Multiple T task, power spectra from TFNs recorded on the Take 5 task were compared in three temporal windows (the 6 seconds preceding food delivery, the 6 seconds following food delivery and across the entire recording session). To compare power spectra from different TFNs, each spectra was first normalized by the total power. Then, as was done for TFNs recorded on the Multiple T task (section 5.2.4), oscillation scores were

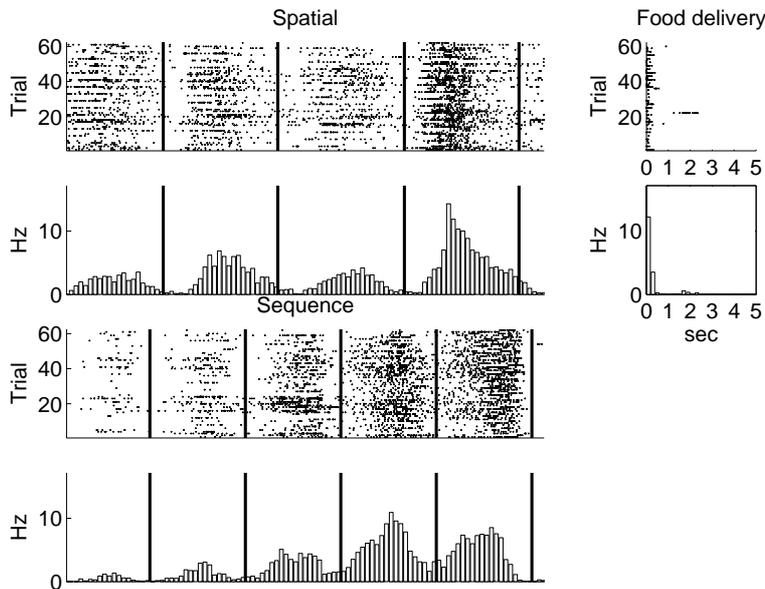


FIGURE 6.2: Maze-responsive PFN. Left: Rastergram and histogram of the activity of a single maze-responsive PFN with respect to linearized spatial location (top). Left bottom: The spatial location of the rat, reordered with respect to the rewarded food location. In both spatial plots, the direction of travel is from left to right. Vertical lines indicate the location of the food delivery sites. Right: Rastergram and peri-event time histogram of the activity of the same PFN with respect to the offset of the tone predicting food delivery. The activity seen shortly following food delivery reflects spiking activity as the rat runs toward the food delivery site, and is easily differentiated from reward-related activity which initiates after the rat arrives at the food delivery site (see Figure 6.4). This PFN was active as the rat ran between each set of food delivery sites (top left) with a strong preference for the location between food delivery site 3 and 4. The cell was also responsive to the location of the animal with respect to the upcoming food delivery location (bottom left), increasing its activity as the rat approached the rewarded food delivery site. The cell was not active as the rat received and consumed its' food reward. Data from R030-2003-05-01 TT10-02.

derived for high and low frequency oscillations by finding the difference in total power in the high frequency (2-3.5 Hz) band relative to the total power in a control band (3.5-5 Hz), and also the difference in total power in the low frequency (0.5-1 Hz) band relative to the total power in a control band (1.5-2 Hz).

As shown in the average Fourier spectra plotted in Figure 6.5, power in the high frequency (2.5-3 Hz) and low frequency (0.5-1 Hz) bands was not enriched in any of the three temporal windows considered. An ANOVA comparing the three temporal windows (Preceding food delivery, following food delivery, or the entire session) indicated no differences between any of the three conditions ($F_{(2,834)} = 0.44$, $p = 0.65$). Similar results were obtained for low frequency oscillations. An ANOVA comparing the three temporal windows (Preceding food delivery, following food delivery, or the entire session) found no differences between any of the three conditions ($F_{(2,834)} = 2.3$, $p = 0.10$, followed by Tukey-Kramer post-hoc comparisons).

In both the Multiple T task and the Take 5 task, rats are running on an elevated track for food rewards. The task-dependence of TFN oscillations indicates that these oscillations are not related to locomotion per se, which rats are performing in both tasks in the 3 seconds preceding food delivery.

Event correlates

While TFN-3s did not show spatial modulations on the Take 5 task, TFN-3s were modulated strongly by events such as food delivery. In an example TFN-3 recorded on the Take 5 task shown in Figure 6.6, the cell did not respond as the rat was running on the maze. Immediately following the tone signalling food delivery, this cell responded with a short period of high firing followed by a brief pause in its activity.

The average population response of TFN-3s recorded on the Take 5 task had a similar response pattern, while TFN-3s recorded on the multiple T showed no such strong response to the delivery of food (see Figure 6.8). As can be seen in the average data, the most common response of TFN-3s was a brief excitation beginning approximately 20 ms after the end of the tone predicting food delivery (~ 120 ms after tone initiation), and lasting about 140 ms. A closer examination of TFN-3 responses revealed that two types of responses, excitations and inhibitions, were present in the TFN-3s recorded on the Take 5 task. These two groups could be differentiated on the basis of their 200ms autocorrelation.

Shown in Figure 6.8A are the 200 ms autocorrelations of the TFN-3s recorded on the Take 5 task, normalized for display purposes by the total number of spikes observed in the 200 ms window. The autocorrelations are sorted on the basis of the strength of the autocorrelation in the first 100 ms, and thus reflect the preference of each cell to fire in the first 100 ms following a spike. Figure 6.8B shows the corresponding responses of the same set of TFN-3s, in the same order as in Figure 6.8B, to the tone predicting food delivery. Excitations following the tone were predominantly present in TFN-3s that had a low firing probability in the first 100 ms of their autocorrelation. Pauses following the tone were frequently seen in TFN-3s which had a higher probability of firing in the first 100 ms of their autocorrelation. Averages for the two groups, split on the basis of the autocorrelations shown in Figure 6.8A, are shown in Figure 6.8C & D. One population of TFN-3s paused for ~ 120 ms while the other population was excited for ~ 160 ms. The population excitations and inhibitions were simultaneous.

In both tasks, neither TFN-1s or -2s showed this pattern of fast excitation and inhibition. Rather, TFN-1s and -2s that were responsive following food delivery or tone presentation in either task tended to have broad activations or inhibitions which were better related to the time at which rats arrived at the food delivery sites than when food was delivered (see Figures 5.14 – 5.15 for examples).

6.2.5 Striatal representation of task parameters

As was done in the Multiple T task (see Chapter 5), Bayesian reconstruction was used to test how well navigation and event-related task parameters were represented by striatal neurons. In addition to the task parameters present in the Multiple T (reward-delivery, spatial location and speed), an additional variable, sequence progress, was tested. The success of the reconstruction was addressed by examining the probability of reconstructing a task parameter given an ensemble firing pattern ($P(X|F_t)$), where X is the task parameter and F_t is a vector representing the firing rate of the striatal ensemble at time t .

To compare the results obtained in the Take 5 task to those of the Multiple T task, a second measure was used to assess the quality of the striatal representation (*reconstruction quality* (RQ), which was defined as the probability that the correct value of the task parameter was reconstructed on the basis of the ensemble firing pattern). In the Take 5 task, the interval in which rewards were received was represented at a level comparable to that of the Multiple T task (shown in Figure 6.10,

compare to results shown in Figure 5.16). However, the reconstruction of spatial location and sequence progress was poor compared to the Multiple T spatial reconstructions, and speed was not well-represented in either task. These results are summarized in Figure 6.11, which shows the average reconstruction quality over all rats for each task parameter. There were no significant differences in the quality of the reconstruction of the five seconds following food delivery in either task ($t(12) = 0.15$, $p = 0.88$). However, in the Multiple T task spatial reconstruction was significantly better than the spatial or sequence reconstruction in the Take 5 task. Across rats, reconstruction quality of space was larger on the Multiple T than the Take 5 ($t(12) = 2.56$, $p = 0.025$), and reconstruction of spatial location on the Multiple T was larger than reconstruction of sequence progress on the Take 5 ($t(12) = 2.62$, $p = 0.022$). Data was also available from four rats tested on both the Take 5 task and subsequently on 3T Multiple T sessions. In these four animals, the same results held for every animal ($RQ_{spatial}$ for Multiple T was larger than $RQ_{spatial}$ or $RQ_{sequence}$ for the Take 5 task for all four rats, and these differences were significant, paired t-tests, $p < 0.025$).

In both the Multiple T task and Take 5 task, striatal ensembles provided a high quality representation of the temporal window following the delivery of food. There were, however, differences between these tasks in how similar the striatal representations were for rewards delivered at separate pellet dispensers. On the Multiple T task, there was a significant reduction in the reconstruction quality (RQ) of the five seconds following the delivery of food when Bayesian reconstruction of the time of reward delivery at one pellet dispenser was performed using tuning curves derived from food deliveries which occurred at the other pellet dispenser (see Figure 5.17). However, on the Take 5 task, a high quality reconstruction of the window in which food rewards were received was obtained even when reconstructions were performed using tuning curves derived from food deliveries at other pellet dispensers. The differences in reconstruction quality between the two conditions (Figure 6.12, *left, middle*) were not significant ($t(8) = 0.99$, $p = 0.35$).

6.3 Discussion

The main results from the Take 5 task are as follows: 1) rats were able to learn the task and demonstrated the adoption of a non-sensory strategy for predicting the upcoming rewarded pellet dispenser, 2) maze-responsive PFNs on the Take 5 were simultaneously modulated by both spatial location and sequence progress, 3) TFN subtypes did not exhibit robust spatial correlates, 4) TFN-3s demonstrated short-latency, event-related responses, and 5) while the tuning curves of maze-responsive PFNs demonstrated firing rate modulations by the spatial location and sequence progress of the rats, there was a complete failure in reconstructing the spatial location or sequence progress of rats on the Take 5 task.

Behaviorally, rats performing the Take 5 task demonstrated an expectation as to where the upcoming food reward would be delivered. The performance of the rats was dominated by a sensitivity to sensory cues, but even in the absence of such cues, rats still were biased to pause at the site where reward should have been delivered. This expectation can be described as evidence for the use of a cognitive strategy guiding behavior on the Take 5 task. The nature of this cognitive strategy is not apparent from the behavior of the animals, however. For example, rats could be adopting a strategy by which they count (or maintain some representation of the total distance travelled between food rewards). Alternatively, rats could adopt a strategy by which, after each food reward, they make one complete revolution of the track, then advance to the next pellet dispenser. In the first case, animals would maintain a representation of the magnitude of their response (the distance travelled since the last reward or the number of pellet dispensers passed since the last reward). In the second

case, rats would on each journey apply two simple rules (return to the last rewarded pellet dispenser, then advance to the next dispenser) in which case they could simply maintain a representation of the next landmark in working memory. While the behavioral data does not discriminate between these possibilities, and other strategies no doubt also could be hypothesized, the performance of rats on the Take 5 task suggests that navigation tasks of this type may be valuable tools for the study of higher level cognition in the rodent.

For the purposes of the present study, the demonstration that the rats possessed some type of cognitive strategy for solving the Take 5 task, then we might ask how does the neural activity in the dorsal striatum support navigation during this behavior? As was found in the Multiple T task (Chapter 5), phasic striatal neurons which were task-responsive were active either during navigation or during food delivery, but not both. In the Take 5 task, maze-responsive PFNs demonstrated less spatial selectivity than was observed in the Multiple T task. For instance, compare the maze-responsive PFN recorded on the Multiple T task shown in Figure 5.4 to the Take 5 maze-responsive PFNs shown in Figures 6.2 & 6.3. On the Take 5 task, many maze responsive PFNs were active during movement between each pair of pellet dispensers, and the fact that each of these movements had a similar trajectory (i.e. in each case the rats were making a right hand turn), may have provided a bias for these maze-responsive PFNs to be less spatially selective. However, maze-responsive PFNs were often more active at one spatial location than others, and also demonstrated a selectivity for the rats' progress in the sequence of actions leading to reward. These results support the observation from the Multiple T task that maze-responsive PFNs were sensitive not only the rat's trajectory, but also the animals spatial location, and the global sequence of actions rats were performing (the sequence of turns rats were presented with). The data from the Multiple T maze demonstrates that maze-responsive PFNs were able to be modulated simultaneously by both spatial and sequence information.

In contrast, while the responses of PFNs on the Take 5 were largely consistent with those observed on the Multiple T maze, the responses of TFNs were dramatically different. On the Take 5 task, TFN-1s and -2s did respond during the receipt of rewards (as was seen for these cell types in the Multiple T task), but both cell types lacked prominent spatial oscillations as rats were running on the maze. Also, while TFN-3s lacked any strong behavioral correlates in the Multiple T task, on the Take 5 task TFN-3s as a population were strongly modulated by the presentation of the tone which signalled food delivery. Examination of the population of TFN-3s indicated the existence of at least two subtypes: those that responded with a phasic increase in firing rate following tone presentation, and those that responded with a phasic decrease in firing rate following tone presentation. The magnitude of each response (the amount of excitation or inhibition) differed between these populations, but they possessed similar timecourses. Also, the autocorrelations of each TFN-3 type were distinctly different. TFN-3s which increased their firing rate in response to the tone predicting food delivery had no tendency to burst, while TFN-3s which decreased their firing rate in response to the tone predicting food delivery had a stronger tendency to fire in bursts. These results may indicate that TFN-3s, as defined in Chapter 4 represent a mixed category composed of two distinct cell types. Also, TFN-3s which paused in their firing following the presentation of the tone predicting food delivery have responses which are similar to those reported for primate tonically active neurons (TANs), which correspond to the cholinergic interneurons of the striatum. While primate TANs also show excitations in response to stimuli which are associated with rewards (reviewed in section 2.2.3), pauses are also present in these TANs, and no TAN responses comparable to the excitations shown in TFN-3s have been reported in primates. It may be that this second type of TFN-3s, which is excited following tone presentation, does not represent the activity of cholinergic interneurons, or that such

cholinergic interneurons in the rodent demonstrate a wider variety of behavioral correlates than those of the primate.

Clear differences between neural activity in the Multiple T and Take 5 tasks were also apparent at the ensemble level. In both tasks, ensembles of striatal neurons cleanly encoded the time in which rats received their food rewards. However, while there was a high-quality representation of spatial location in the Multiple T task, the representation of both spatial location and sequence progress was poor on the Take 5 task. The representation of sequence progress was also not improved by taking into account the spatial dependence of maze-responsive PFNs (shown in Figure 6.3). While both tasks involved navigation in order to retrieve food rewards, one important difference between the Multiple T and Take 5 tasks is the degree to which spatial cues are unambiguously associated with the rats' goals. In the Take 5 task, the spatial location of the rat determines the upcoming goal location (the rewarded pellet dispenser), so long as the rats' sequence progress is known. Both the behavior of the rats on probe trials (in which sensory cues were omitted) and the encoding of sequence progress by maze-responsive PFNs indicate that rats had access to information related to sequence progress. However, at the ensemble level striatal neurons did not provide a clear representation of these navigation parameters. The lack of representation of navigation parameters by striatal ensembles, when single neurons demonstrated tuning to these parameters, is counterintuitive. These results suggest that while maze-responsive PFNs were modulated by both spatial location and sequence progress, either the spatial and sequence tuning of these neurons within the ensemble did not cover the parameter space uniformly, or the responses were too variable at each value of the task parameter to reliably encode spatial location and sequence progress.

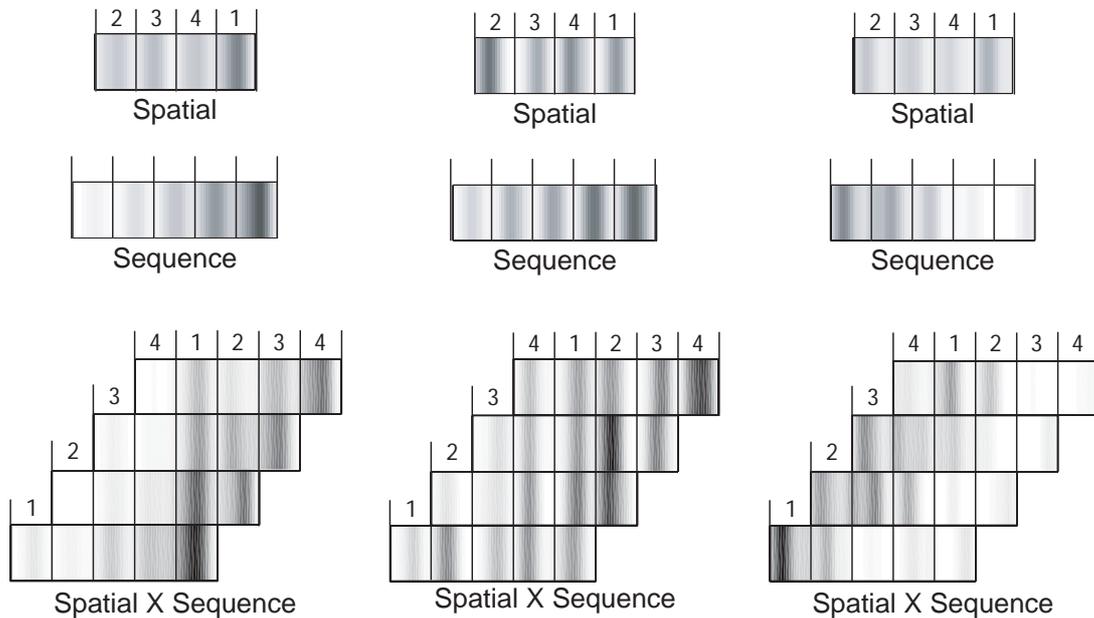


FIGURE 6.3: Encoding of sequence and position by maze-responsive PFNs. Three example maze-responsive PFNs recorded on the Take 5 task. Each PFN was modulated by the animal's location on the maze (Spatial) and in the sequence of actions leading to reward (Sequence). The activity of each PFN was better described by a combination of both spatial location and sequence progress (Spatial \times Sequence). The cell on the left is the same maze-responsive PFN shown in Figure 6.2. The other maze responsive PFNs were R036-2003-08-25 TT06-03 (*middle*) and R032-2003-05-27 TT06-03 (*right*).

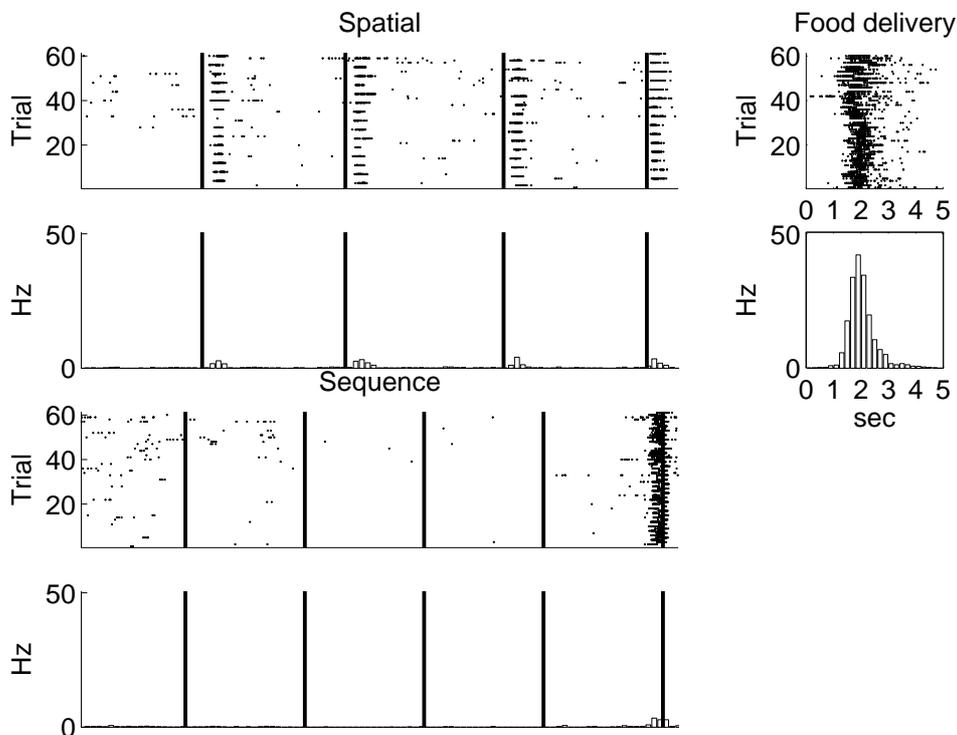


FIGURE 6.4: Reward-responsive PFN. Plotting conventions are the same as in Figure 6.2. The cell was specifically active as the rat received and consumed its food reward. This PFN did not display spatial- or sequence-related firing, except in that its spikes were observed at each food delivery site during reward consumption. Data from R032-2003-06-07 TT03-03.

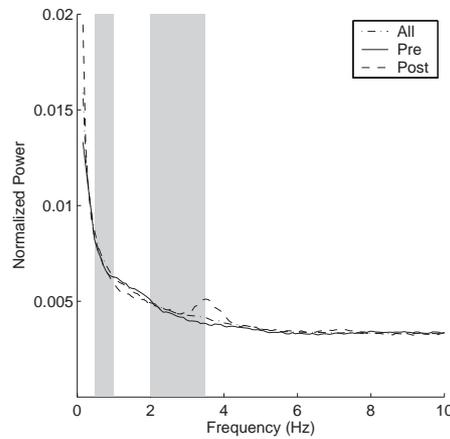


FIGURE 6.5: TFN oscillations in the Take 5 task. Averaged Fourier spectra were calculated from TFN spike trains and normalized by total power. Unlike in the Multiple T task (see Figure 5.13), TFNs on the Take 5 task did not demonstrate increased power in the 2-3.5 Hz or 0.5-1 Hz band in any temporal window (6 seconds preceding [Pre] or following [Post] food delivery, or across the entire recording session [All]).

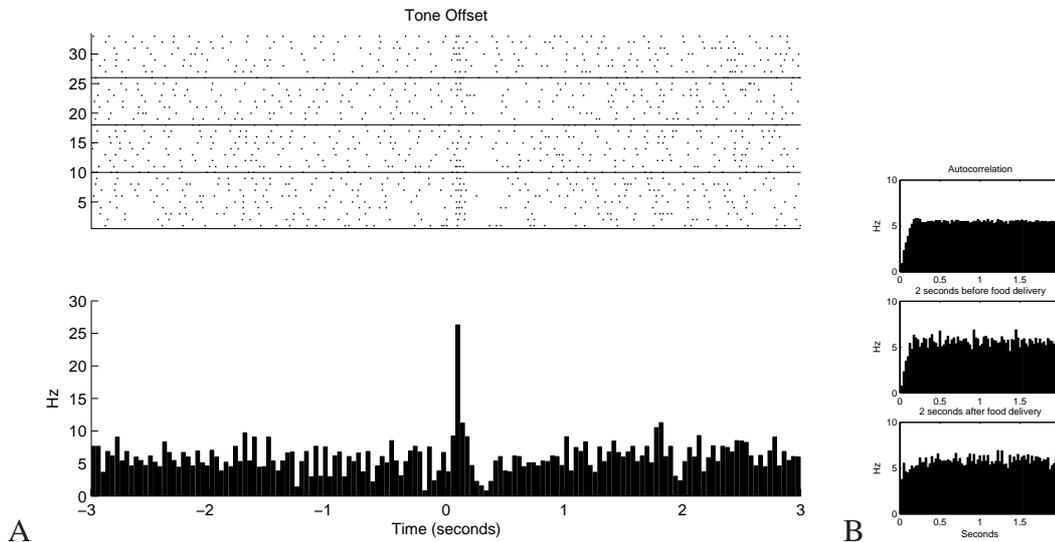


FIGURE 6.6: A TFN-3 with event-related firing on the Take 5 task. *A* Peri-event time rastergram and histogram of the firing rate of the cell relative to the offset of a short (100 ms) tone predicting food delivery. Following the tone that signaled food delivery, this cell had a short period of high firing followed by a brief pause in activity. The cell did not show spatial responses. *B* The autocorrelation of this cell over the entire session (*top*), in the 12 seconds preceding food delivery (*middle*) and in the 12 seconds following food delivery at the first food delivery site (*bottom*). (R037-2003-09-12-TT07-01 Take 5 Trials = 33)

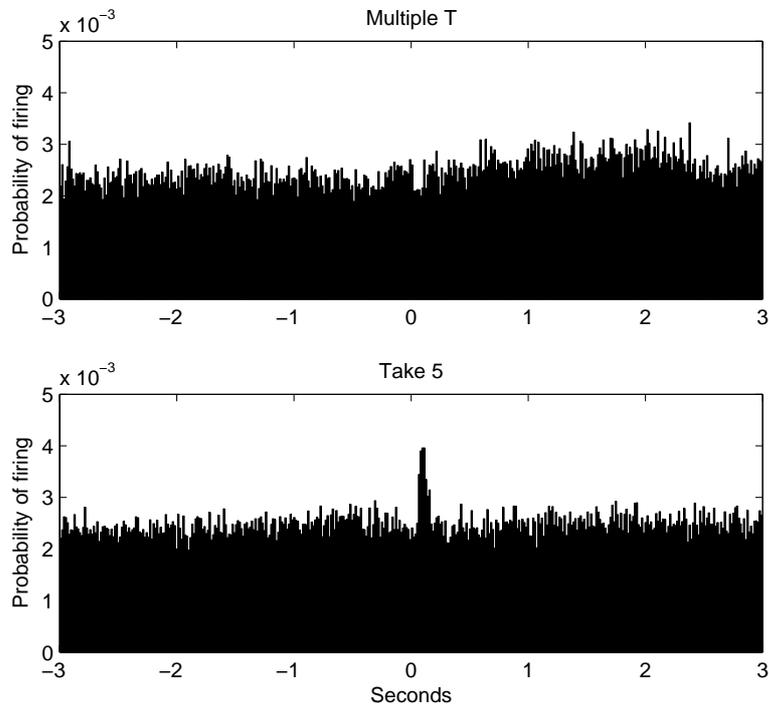


FIGURE 6.7: Average probability of firing of TFN-3s relative to food delivery or tone predicting food delivery. *Top:* On the Multiple T task, TFN-3s as a population were unresponsive to food delivery. *Bottom:* In contrast, on the Take 5 task, TFN-3s as a population responded to the tone signalling food delivery with a fast, time-locked increase in firing probability which lasted for ~ 100 ms. Data averaged over all TFN-3s, bins = 15 ms.

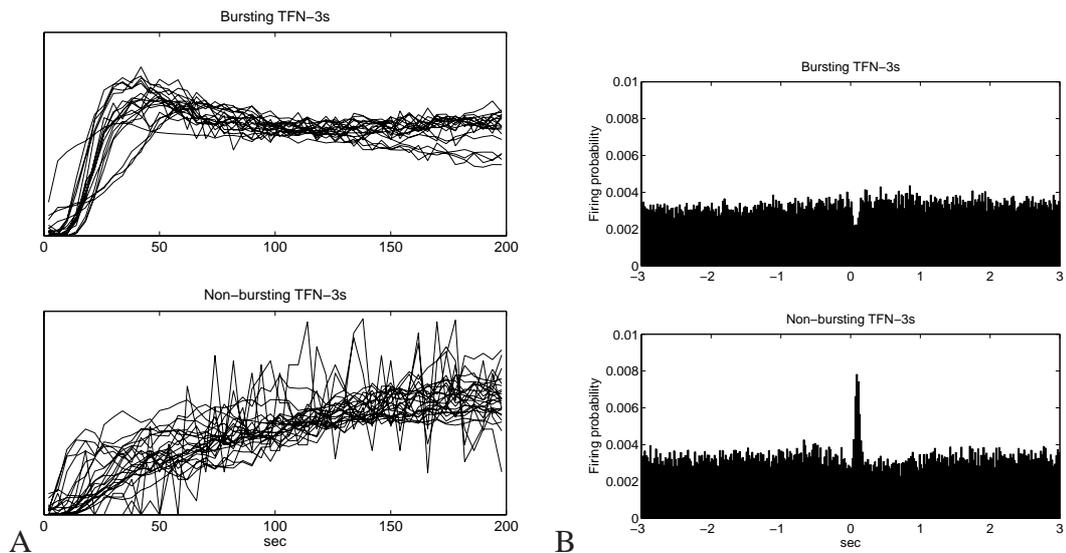


FIGURE 6.8: Two types of TFN-3s could be distinguished on the basis of autocorrelations and behavioral responses. TFN-3s could be sorted into two groups, those with a high or low probability of firing in the first 100 ms of their autocorrelation (A), which corresponded respectively to TFN-3s which were inhibited or excited following tone presentation on the Take 5 task (B).

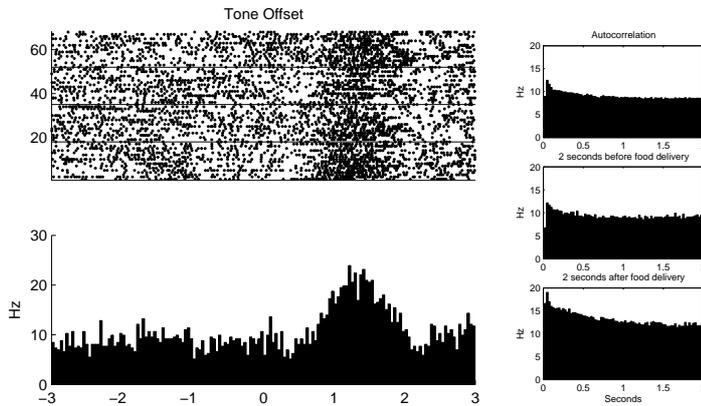


FIGURE 6.9: A TFN-2 with event-related firing on the Take 5 task. *Top:* Peri-event time histograms (PETHs) of the firing rate of the cell relative to food delivery at the first and second food delivery sites). *Bottom left:* Rastergram and histogram of linearized firing on the maze. Key as in Figure 5.11. Following the tone that signaled food delivery, this cell had two activations, the first phasic and lasting ~ 1 second, the second more broad and lasting several seconds. The cell did not show spatial responses. *Right:* The autocorrelation of this cell over the entire session (*top*), in the 12 seconds preceding food delivery (*middle*) and in the 12 seconds following food delivery at the first food delivery site (*bottom*). (R032-2003-06-06 TT04-02 Take 5 Trials = 68)

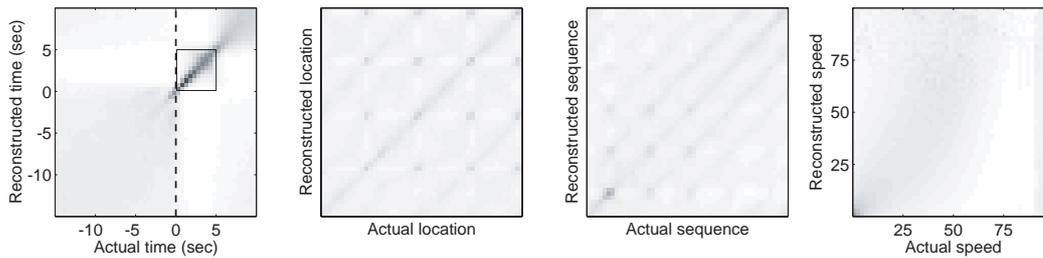


FIGURE 6.10: Reconstruction of task parameters in the Take 5 task. Ensembles of simultaneously recorded striatal neurons were used to reconstruct the time of reward delivery, the spatial location, sequence progress and running speed of rats during Take 5 task performance. As was the case in the Multiple T task, there was a high-quality representation of the five second interval following the delivery of food, but not of running speed. Unlike in the Multiple T task, there was not a high-quality representation of the spatial location or sequence progress of rats running the Take 5 task.

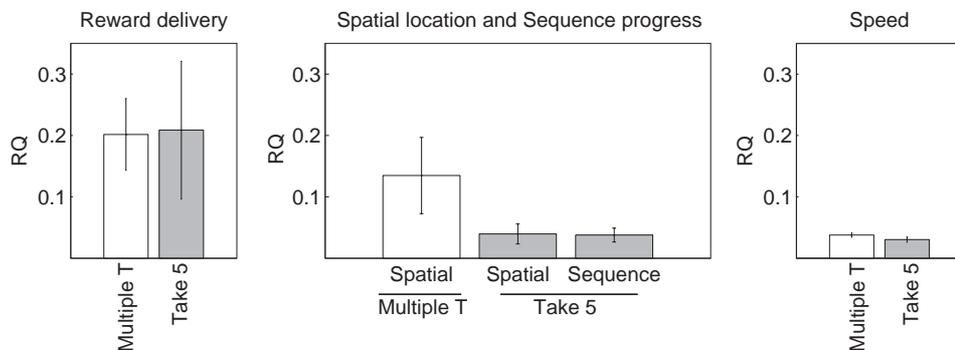


FIGURE 6.11: Reconstruction quality for navigation and event-related parameters in the Multiple T and Take 5 tasks. *Left:* There was no significant difference in the reconstruction quality (RQ, defined in Methods) for the five seconds following food delivery (Multiple T) or the tone signalling food delivery (Take 5). *Middle:* $RQ_{spatial}$ on the Multiple T task was significantly larger than $RQ_{spatial}$ or $RQ_{sequence}$ on the Take 5 task. *Right:* RQ_{speed} was poor in both tasks, but significantly larger in the Multiple T than the Take 5 task. Bargraphs summarize data presented in Figures 5.16 & 6.10. Bars represent mean across rats, errorbars are 95% confidence intervals.

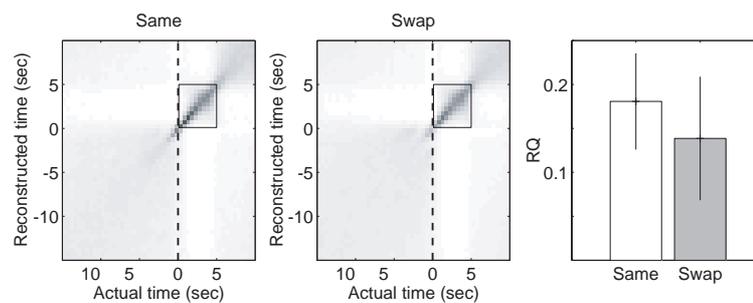


FIGURE 6.12: Similarity of reward-receipt encoding in the Take 5 task. Bayesian reconstruction was used to determine the similarity of the striatal ensemble representation in the 5 second window following the delivery of food occurring at different pellet dispenser. When the tuning curves used in the Bayesian reconstruction were derived from rewards delivered at the same pellet dispenser (*Left, Same*), there was a high quality representation of the five seconds following the delivery of food, as was the case in Figure 6.10. Unlike in the Multiple T task, however, when tuning curves were derived from rewards delivered at different pellet dispensers (*Middle, Swap*), there was also a high-quality representation of the five seconds following reward-delivery. There were no significant differences in the reconstruction quality (RQ) between the two conditions (*Right*, error bars represent mean and 95% confidence intervals over rats).

Chapter 7

Conclusions

Using the data presented in Chapters 4, 5, and 6, we can now return to the questions posed at the start of the thesis. Specifically, these data collected from multiple navigation tasks advances answers to several current issues in basal ganglia function. 1) Can distinct types of striatal neurons (projection neurons, different types of interneurons) be differentiated on the basis of extracellular recordings? 2) How do striatal neurons contribute to the learning and performance of navigation tasks? 3) What type of information is represented by the striatum, and how does the striatal representation change as a function of experience? The remainder of the present chapter summarized the results of the thesis with respect to each of these questions.

In awake, behaving rats performing navigation tasks, putative projection neurons and several putative interneuron types can be identified in the dorsal striatum. On the basis of extracellularly recorded spike trains obtained from neurons in the striatum, phasic and tonic neurons could be easily identified. By comparison to work in awake, behaving primates (Kimura et al., 1990), phasic firing neurons (PFNs) are likely to be striatal projection neurons, while tonic firing neurons (TFNs) are likely to be striatal interneurons. The PFN versus TFN distinction made in rodents is similar to the phasically active neuron (PAN) versus tonically active neuron (TAN) distinction made in primates (Kimura et al., 1984, 1990, 1996). However, while primate TANs correspond to cholinergic interneurons (Aosaki et al., 1995), rodent TFNs are in all likelihood not simply the rodent TAN equivalent. Similarly, while a group of neurons has been described by Berke et al. (2004) which are similar to the rodent TFNs described in Chapter 4, TFNs are not likely to be principally fast-spiking parvalbumin-immunoreactive (PV+) striatal interneurons. TFNs could be divided into several distinct categories on the basis of firing properties such as their autocorrelation and firing rate relationships to other striatal neuron types. These firing patterns were unchanged as animals transitioned from awake, quiet rest to increased activity. Also, similar categorizations were obtained in animals performing very different types of tasks: maze-learning, operant conditioning and a cognitive navigation task. Further, each of these tonic neuron subtypes had a distinct pattern of behavioral correlates, supporting the distinction between tonic neuron subtypes. In all likelihood, TFNs as a population represent several interneuron populations, one of which (TFN-3) has firing properties consistent with primate TANs, and a subset of which has behavioral correlates analogous to primate TANs: a brief pause in activity following the presentation of reward-predictive stimuli. The other categories of striatal TFNs (TFN-1, -2, and other -3s) may represent other interneurons, specifically PV+, NADPH-d+ or calretinin (CR+) GABAergic interneurons. In the case of TFN-3 interneurons which did not pause, but instead demonstrated a phasic activation, these neurons may be an example of cholinergic interneurons with

novel behavioral correlates. The identification of TFNs and TFN subtypes in brain material histologically processed for cell-specific markers (PV, NADPH-d, CR, ChAT) will be an important step in furthering the study of striatal interneurons in awake, behaving rats. While technically difficult, if the firing patterns which characterize PFNs and TFN subtypes persist in anesthetized animals, these cells could be filled with intracellular dyes. Alternatively, use of juxtacellular methods would likely be able to identify these neuron subtypes (Pinault, 1996).

Phasic and tonic neurons of the rodent striatum have distinct patterns of task-dependent, behavioral modulation during rodent navigation. A second question taken up in Chapters 5 & 6 of the thesis was the responses of rodent PFNs to navigation and event-related task parameters in two tasks: the Multiple T and the Take 5. In both tasks, two types of task-related neuron were obtained: those that were active during navigation (maze-responsive) and those that were active following food delivery (reward-responsive). Maze- and reward-responsive PFNs were mutually exclusive groups: PFNs which were active during navigation were not active during food delivery, and vice versa. Such a result may relate to the proposed distinction in actor/critic models between action-related and prediction-related units, with maze-related neurons corresponding to the former and reward-related units related to the latter. Alternatively, the functional segregation could be related to the differences in the types of actions used in each behavior, i.e. locomotion versus food retrieval and consumption. In that case, both maze- and reward-responsive PFNs might be related to actor-units. Each of these two hypotheses would make different predictions as to the location of maze- and reward-responsive PFNs with respect to striatal striosome/matrix compartmentalization. In both cases, maze-responsive PFNs, if they correspond to actor-units, would be expected to be located in the striatal matrix, and project to basal ganglia output nuclei, the globus pallidus and the substantia nigra pars reticulata. Under the hypothesis that reward-responsive PFNs are the critic-units, it would be predicted that reward-responsive PFNs project to the pars compacta division of the substantia nigra. In contrast, the hypothesis that reward-responsive PFNs are action units encoding the actions involved in reward retrieval and consumption, it would be predicted that reward-responsive PFNs also project to basal ganglia output nuclei. An intermediate possibility is that reward-responsive PFNs represent two populations, some of which are involved in action selection, while others are involved in sending reward signals to dopaminergic structures for calculating reward prediction error signals. In this case, it would be expected that maze-responsive PFNs are located in the matrix while reward-responsive PFNs are expected in both compartments. To test these hypotheses, it will be necessary to be able to identify the location of striatal PFNs with respect to striatal subcompartments and projection targets. A more complete understanding of striatal function would be obtained by testing these hypotheses, which was not possible in the experiments presented in this thesis. At least two approaches for future experiments are feasible: 1) the reconstruction of electrode tracks in histological material processed to reveal striatal subcompartments, and 2) the antidromic stimulation of neurons recorded in the striatum with stimulating electrodes placed in projection targets such as the globus pallidus, the SNpc and the SNr.

Chapters 5 & 6 also addressed the nature of striatal responses: to what extent are the responses of striatal PFNs explained by task parameters? Reward-responsive PFNs were phasically active following food delivery: did these responses encode food delivery in general (i.e. independent of other contextual cues)? Maze-responsive PFNs were active in one or more locations on the Multiple T maze, and often in multiple locations on the Take 5 track. Were these responses primarily spatial? Or, did these responses reflect the motor behavior of the rats or more complex variables? In

general, the responses of both maze- and reward-responsive PFNs were sensitive to context: reward-responsive PFNs often differentiated food delivered at pellet dispensers located at different spatial locations, while maze-responsive PFNs typically differentiated spatial locations, different actions occurring at the same physical locations, and the sequence of actions rats were progressing through to receive their food rewards. These results indicate that striatal PFNs were modulated by context, in that reward-responsive PFNs were not active following food delivery at every pellet dispenser, and the responses of maze-responsive PFNs depended on combinations of spatial location, the actions being performed, and the global sequence of actions performed by the rats in order to receive food rewards. This type of contextual modulation is very similar to that described for primate PANs in response to limb movements and saccadic eye movements (reviewed in Chapter 2.2). These results extend on that body of work by demonstrating that many phasic striatal neurons in the rodent have spatially modulated responses which are modulated by the actions taken by the route (the shape of the rat's path on the Multiple T maze), and the sequence of actions being performed (the sequence of turns rats navigated through on the Multiple T maze and the sequence of movements made to receive food rewards on the Take 5 task). When spatial location and sequence progress were dissociated in the Take 5 task, many maze-responsive PFNs demonstrated tuning to both spatial location and sequence progress. While sequence-specific activity has been demonstrated in the primate in visuo-motor sequencing (Kermadi et al., 1993; Kermadi and Joseph, 1995) and in rodents during highly stereotyped grooming movements (Aldridge and Berridge, 1998), the results from the Multiple T and Take 5 tasks are the first demonstrations of sequence-related activity in rodents during arbitrary sequential navigation tasks.

The dorsal striatal representation of task parameters may indicate what types of “events” are salient to structures in the basal ganglia. Together, the literature reviewed in Chapter 2.2 and the correlates of individual striatal neurons presented in Chapters 5 & 6 from rats performing navigation tasks leaves a number of open questions as to the nature of the striatal representation. Specifically, what is the nature of the striatal representation? Given the bewildering variety of striatal behavioral correlates, including limb movements, saccades, sensory stimuli, sequence-information, expectation, reward, spatial location, it seems that striatal neurons in primates and rodents are modulated by almost any experimental parameter that has been tested, although there is some segregation of responsive neurons in specific striatal areas (neurons related to saccadic eye movements are located primarily in the caudate, while those responsive to limb movements are restricted to the putamen (Alexander and DeLong, 1985; Hikosaka et al., 1989a)). If this is true, then is the striatum providing a high-quality representation of each of these variables, including variables which occur simultaneously but are not congruent (for instance, spatial location and sequence progress on the Take 5)?

This question was addressed by employing ensemble-level analyses of simultaneously recorded striatal neurons (see Chapters 5 & 6). Using Bayesian reconstruction, the quality of the neural representation was examined relative to both navigation-related (spatial location, sequence progress, and speed) and event-related (the time of food delivery) task parameters. In both the Multiple T and Take 5 tasks, there existed a high-quality representation of the period of time in which rats retrieved and consumed their food rewards. However, of the navigation-related task parameters, only spatial location was well-represented by striatal ensembles, and only on the Multiple T task: on the Take 5, there was a failure of the reconstruction of spatial location and sequence progress, even when the modulation of sequence tuning in PFNs by upcoming reward location was taken into account. In neither task was speed well-represented. These results indicate a preferential encoding of spatial

location/sequence progress on the Multiple T task over encoding of speed, and a task dependent encoding of spatial location between tasks.

An alluring hypothesis that explains the difference in the encoding of task-parameters between tasks is that the striatal representation may depend on the degree to which these parameters can be unambiguously associated with goals. On both the Multiple T and Take 5 tasks, the presentation of food-predictive cues (the firing of the pellet dispenser on the Multiple T and the tone predicting food delivery on the Take 5) is highly informative of the pattern of actions that need to be taken in order to receive food rewards. Likewise, the location of the rat on the Multiple T maze is well-related to the sequence of actions that must be completed in order to receive the next food delivery. However, on the Take 5 task, the location of the rat is only informative if the rat also maintains a representation of the previously rewarded pellet dispenser. Behaviorally, this is the case, as demonstrated in Chapter 6: even in the absence of sensory stimuli (the tone signalling food delivery and the sound of the pellet dispensers), rats still were biased to pause at the location where reward was expected. This demonstrates that at some level, rats had a cognitive representation of their location in the sequence of actions leading to reward delivery. This representation of sequence progress was also reflected in the tuning of maze-responsive PFNs in the Take 5 to sequence progress. Although striatal PFNs had access to information related to spatial location and sequence progress (reflected by their tuning to these parameters), at the ensemble level there was no evidence for a coherent representation of either of these parameters. This may indicate that while the rats had a behavioral expectation of the upcoming food delivery site in the Take 5 task, and striatal PFNs were tuned to navigation-related parameters, this tuning was not sufficient to differentiate spatial locations and sequence progress, perhaps because of competition between these variables on a trial-by-trial basis for control of maze-responsive PFN firing. Under such an explanation, it is likely then that the striatum, or at least the more dorsal and medial aspect of the striatum which made up the bulk of these recordings, may not participate in action selection in the Take 5 task, but could be involved in navigation-related action selection in the Multiple T task. These differences in spatial encoding in the Multiple T and Take 5 tasks then could be evidence of the task-dependent recruitment of the striatum. If true, then lesions or inactivations of the striatum should have dissociable effects on the behavior of the rats on these tasks. Inactivation of the striatum may impair behavior of rats on the Multiple T while having little effect on performance of the Take 5 task.

Are there no differences between event-related encoding in these two tasks? Again, using Bayesian reconstruction techniques, the representation of the time following food delivery was examined in both tasks by comparing the reconstruction of the time following food delivery at one pellet dispenser using tuning curves derived from trials at that same pellet dispenser, or at one of the other pellet dispensers used in the task (two dispensers were used in the Multiple T task, and four were used in the Take 5 task). Interestingly, while there was a decreased quality of reconstruction in the Multiple T task when tuning curves were derived from a different pellet dispenser, there was no significant decrease in the quality of reconstruction in the Take 5 task in either case. This indicates that in the Multiple T task, striatal ensembles differentiated food delivery at spatially distinct locations, while in the Take 5 task, this was not the case. At the ensemble level, then, the striatum does not necessarily differentiate events occurring in very different sensory contexts. It may be that the difference between these tasks is not the *spatial* context, but the *sequence* context of food delivery. In the Multiple T task food delivery at each of the two pellet dispensers occurred at different locations of the animals' progress through the actions leading to reward: the first pellet dispenser was

located immediately following the navigation sequence, while the second pellet dispenser was located after a short journey from the first. In the Take 5 task, each food delivery occurred in the same location in the animals' progress through the sequence of actions leading to reward: food delivery at each dispenser occurred after the animal ran $\frac{5}{4}$ the length of the rectangular track. This sequence-based hypothesis could easily be tested with a variant of the Take 5 task, in which multiple pellet dispensers were rewarded as the rat progressed through the task sequence: for instance, instead of rewarding only the 5th dispenser-to-dispenser journey, both the 5th and the 4th or 3rd movement could be rewarded. Under the previous explanation, it would be expected that striatal ensembles would differentiate reward-delivery based on sequence progress, but not on the basis of spatial location.

Changes in the striatal representation of space parallel, and precede, the development of a stable route, suggesting a mechanism by which striatal neural activity may guide subsequent behavioral refinement. In the last section of Chapter 5 turns to the issue of learning. Given that rats show two types of changes in behavioral performance on the Multiple T maze, and given that there exists a representation of spatial location, which is presumed to reflect not only space, but also the actions of the rat and the sequence of turns presented (as described in Chapter 5), is there then any evidence for plasticity of the striatal representation of space, and how does the development of the striatal representation of space correspond to the changes in behavior? On the Multiple T maze, rats show two changes in performance: a fast decrease in the number of errors made (incorrect turns on the T maze choices) and a slow increase in the similarity of the path taken through the maze to subsequent paths (small amplitude changes in route). By analogy to the types of striatally-dependent navigation learning that has been demonstrated on the plus maze by Packard and others (as reviewed in Chapter 2.3), the fast decrease in errors may reflect the development of a place strategy dependent on the hippocampus and the posterior dorsomedial striatum, while the slower changes in the route taken through the maze may reflect the development of a response-strategy dependent on the dorso-lateral striatum. As the primary site of the neural recordings described in this thesis was the anterior dorsal striatum, most of the striatum sampled would be expected to be in regions devoted to slowly developing response-based navigation strategies. As such, changes in the activity of striatal neural ensembles in the Multiple T task were expected to be better related to the development of a stable route than to the elimination of errors. Examining the quality of the Bayesian reconstruction of spatial location, this was found to be the case: the striatal spatial reconstruction improved as a function of experience, with a rate that was slower the decrease in the number of errors made, and faster than the development of a stable route. The striatal representation of space was well-positioned, temporally, to influence the development of a stable route, and may reflect the recruitment of areas in the dorsal striatum during automation of navigating complex trajectories. Further experiments are required, however, to establish that the dorsal striatum in this region is involved in performance of the Multiple T tasks, and specifically in the development of a stable route. This evidence could come from studying the behavior of rats performing the Multiple T maze under inactivation of the dorsal striatum or hippocampus with local anesthetics or blockade of neurotransmission. Under the proposed account of the neural basis of performance of the Multiple T task, it would be expected that interference with hippocampal neurotransmission would impair the elimination of errors, while interference with dorsolateral striatal neurotransmission would impair the development of a stable route.

Glossary

Ach Acetylcholine.

ChAT Choline acetyltransferase. Synthesizes the neurotransmitter acetylcholine from choline. ChAT+, neurons which express choline acetyltransferase.

CR Calretinin, a calcium binding protein. CR+, neurons which are immunoreactive for CR.

FS Fast-spiking neurons. Characterizes action potentials of striatal GABAergic, parvalbumin immunoreactive interneurons.

GP Globus pallidus.

LA Long afterhyperpolarization. Characterizes action potentials of striatal cholinergic interneurons.

MPTP 1-methyl-4-phenyl-1,2,3,6-tetrahydropyridine. Used to destroy dopaminergic neurons in primate models of Parkinson's disease.

NADPH-d Nicotinamide adenine dinucleotide phosphate diaphorase. NADPH-d+, neurons which have NADPH-d activity.

NOS Nitric oxide synthase. NOS+, neurons which are immunoreactive for NOS.

PAN Phasically active neuron. In the primate, PANs correspond to medium-sized spiny projection neurons.

PFN Phasic-firing neuron. Striatal neurons in the rodent which spend a majority of their time in long interspike intervals.

PLTS Persistent, low threshold spiking neurons. Characterizes striatal GABAergic, nitric oxide synthase immunoreactive interneurons.

PV Parvalbumin, a calcium binding protein. PV+, neurons which are immunoreactive for PV.

SNpc Substantia nigra pars compacta. Origin of the major dopaminergic input to the dorsal striatum (caudoputamen).

SNpr Substantia nigra pars reticulata.

TAN Tonicly active neuron. In the primate, TANs correspond to the large, aspiny cholinergic striatal interneurons.

TFN Tonic-firing neuron. Striatal neurons in the rodent which rarely have long interspike intervals.

Bibliography

- Albin RL, Young AB, and Penney JB. The functional anatomy of basal ganglia disorders. *Trends in Neurosciences* 12: 366–374, 1989.
- Aldridge JW and Berridge KC. Coding of serial order by neostriatal neurons: A “natural action” approach to movement sequence. *Journal of Neuroscience* 18: 2777–2787, 1998.
- Alexander GE and Crutcher MD. Functional architecture of basal ganglia circuits: Neural substrates of parallel processing. *Trends in Neurosciences* 13: 266–271, 1990a.
- Alexander GE and Crutcher MD. Neural representations of the target (goal) of visually guided arm movements in three motor areas of the monkey. *Journal of Neurophysiology* 64: 164–178, 1990b.
- Alexander GE and Crutcher MD. Preparation for movement: Neural representations of intended direction in three motor areas of the monkey. *Journal of Neurophysiology* 64: 133–150, 1990c.
- Alexander GE and DeLong MR. Microstimulation of the primate neostriatum. II. Somatotopic organization of striatal microexcitable zones and their relation to neuronal response properties. *Journal of Neurophysiology* 53: 1417–1430, 1985. .
- Alexander GE, DeLong MR, and Strick PL. Parallel organization of functionally segregated circuits linking basal ganglia and cortex. *Annual Reviews Neuroscience* 9: 357–381, 1986.
- Aosaki T, Graybiel AM, and Kimura M. Effect of the nigrostriatal dopamine system on acquired neural responses in the striatum of behaving monkeys. *Science* 265: 412–415, 1994.
- Aosaki T, Kimura M, and Graybiel AM. Temporal and spatial characteristics of tonically active neurons of the primate’s striatum. *Journal of Neurophysiology* 73: 1234–1252, 1995.
- Barto AG. Adaptive critics and the basal ganglia. In: *Models of Information Processing in the Basal Ganglia*, edited by Houk JC, Davis JL, and Beiser DG, Cambridge MA: MIT Press, pp. 215–232. 1995.
- Bennett BD and Bolam JP. Synaptic input and output of parvalbumin-immunoreactive neurons in the neostriatum of rats. *Neuroscience* 62: 707–719, 1994.
- Bennett BD and Wilson CJ. Spontaneous activity of neostriatal cholinergic interneurons *in vitro*. *Journal of Neuroscience* 19: 5586–5596, 1999.
- Bennett MVL. Electrical transmission: A functional analysis and comparison to chemical transmission. In: *The handbook of physiology - The nervous system*, edited by Kandel ER, American Physiological Society, Washington, DC, volume 1, chapter 11, pp. 357–416. 1999.

- Berke J, Okatan, Skuski, and Eichenbaum H. Synchronous striatal spindle and gamma oscillations in freely-moving rats. *Society for Neuroscience Abstracts* 2003.
- Berke JD, Okatan M, Skurski J, and Eichenbaum HB. Oscillatory entrainment of striatal neurons in freely-moving rats. *Neuron* 43: 883–896, 2004.
- Berridge KC and Whishaw IQ. Cortex, striatum and cerebellum: Control of serial order in a grooming sequence. *Experimental Brain Research* 90: 275–290, 1992.
- Bevan MD, Booth PAC, Eaton SA, and Bolam JP. Selective innervation of neostriatal interneurons by a subclass of neuron in the globus pallidus of the rat. *Journal of Neuroscience* 18: 9438–9452, 1998.
- Blackwell KT, Czubayko U, and Plenz D. Quantitative estimate of synaptic inputs to striatal neurons during up and down states *in vitro*. *Journal of Neuroscience* 23: 9123–9132, 2003.
- Bonsi P, Florio T, Capozzo A, Pisani A, Calabresi P, Siracusano A, and Scarnati E. Behavioral learning-induced increase in spontaneous gaba_a-dependent synaptic activity in rat striatal cholinergic interneurons. *European Journal of Neuroscience* 17: 174–178, 2003.
- Boussaoud D and Kermadi I. The primate striatum: Neuronal activity in relation to spatial attention versus motor preparation. *European Journal of Neuroscience* 9: 2152–2168, 1997.
- Bracci E, Centonze D, Bernardi G, and Calabresi P. Dopamine excites fast-spiking interneurons in the striatum. *Journal of Neurophysiology* 87: 2190–2194, 2002.
- Bracci E, Centonze D, Bernardi G, and PCalabresi. Voltage-dependent membrane potential oscillations of rat striatal fast-spiking interneurons. *Journal of Physiology* 549: 121–130, 2003.
- Brown EN, Frank LM, Tang D, Quirk MC, and Wilson MA. A statistical paradigm for neural spike train decoding applied to position prediction from ensemble firing patterns of rat hippocampal place cells. *Journal of Neuroscience* 18: 7411–7425, 1998.
- Brown MA and Sharp PE. Simulation of spatial learning in the Morris water maze by a neural network model of the hippocampal formation and nucleus accumbens. *Hippocampus* 5: 171–188, 1995.
- Canales JJ and Graybiel AM. A measure of striatal function predicts motor stereotypy. *Nature Neuroscience* 3: 377–383, 2000.
- Cannizzaro C, Tel BC, Rose S, Zeng B, and Jenner P. Increased neuropeptide y mrna expression in striatum in parkinson's disease. *Molecular Brain Research* 110: 169–176, 2003.
- Carelli RM and West MO. Representation of the body by single neurons in the dosolateral striatum of the awake, unrestrained rat. *Journal of Comparative Neurology* 309: 231–249, 1991.
- Carelli RM, Wolske M, and West MO. Loss of lever press-related firing of rat striatal forelimb neurons after repeated sessions in a lever pressing task. *Journal of Neuroscience* 17: 1804–1814, 1997.

- Centonze D, Bracci E, Pisani A, Gubellini P, Bernardi G, and Calabresi P. Activation of dopamine d₁-like receptors excites its interneurons of the striatum. *European Journal of Neuroscience* 15: 2049–2052, 2002.
- Chang HT and Kita H. Interneurons in the rat striatum: relationships between parvalbumin neurons and cholinergic neurons. *Brain Research* 574: 307–311, 1992.
- Chang Q and Gold PE. Switching memory systems during learning: Changes in patterns of brain acetylcholine release in the hippocampus and striatum in rats. *Journal of Neuroscience* 23: 3001–3005, 2003a.
- Chang Q and Gold PE. Switching memory systems during learning: Changes in patterns of brain acetylcholine release in the hippocampus and striatum of rats. *Journal of Neuroscience* 23: 3001–3005, 2003b.
- Chesselet MF, Gonzales C, Lin CS, Polsky K, and Jin BK. Ischemic damage in the striatum of adult gerbils: Relative sparing of somatostatinergic and cholinergic interneurons contrasts with loss of efferent neurons. *Experimental Neurology* 110: 209–218, 1990.
- Cicchetti F and Parent A. Striatal interneurons in huntington's disease: Selective increase in the density of calretinin-immunoreactive medium-sized neurons. *Movement Disorders* 11: 619–626, 1996.
- Cohen NJ and Squire LR. Preserved learning and retention of pattern-analyzing skill in amnesia: Dissociation of knowing how and knowing that. *Science* 210: 207–210, 1980.
- Cote PY, Sadikot AF, and Parent A. Complementary distribution of calbindin d28k and parvalbumin in the basal forebrain and midbrain of the squirrel monkey. *European Journal of Neuroscience* 3: 1316–1329, 1991.
- Cowan RH and Wilson CJ. Spontaneous firing patterns and axonal projections of single corticostriatal neurons in the rat medial agranular cortex. *Journal of Neurophysiology* 71: 17–32, 1994.
- Cowan RL, Wilson CJ, Emson PC, and Heizmann CW. Parvalbumin-containing gabaergic interneurons in the rat neostriatum. *Journal of Comparative Neurology* 302: 197–205, 1990.
- Cromwell HC and Berridge KC. Implementation of action sequences by a neostriatal site: A lesion mapping study of grooming syntax. *Journal of Neuroscience* 16: 3444–3458, 1996.
- Crutcher MD and Alexander GE. Movement-related neuronal activity selectively coding either direction or muscle pattern in three motor areas of the monkey. *Journal of Neurophysiology* 64: 151–163, 1990.
- Daw ND. *Reinforcement learning models of the dopamine system and their behavioral implications*. Ph.D. thesis, Carnegie Mellon University, 2003.
- Daw NS and Touretzky DS. Behavioral considerations suggest an average reward TD model of the dopamine system. *Neurocomputing* 32-33: 679–684, 2000.
- Dawbarn D, Quidt MED, and Emson PC. Survival of basal ganglia neuropeptide y-somatostatin neurones in huntington's disease. *Brain Research* 340: 251–260, 1985.

- DeLong MR and Georgopolus AP. Motor functions of the basal ganglia. In: *Handbook of Physiology, Sect. 1: The Nervous System, Motor Control, Part 2*, edited by Brookhart JM, Mountcastle VB, Brooks VB, and Geiger SR, Bethesda, MD: American Physiological Society, volume II, chapter 21, pp. 1017–1061. 1981.
- Devan BD and White NM. Parallel information processing in the dorsal striatum: Relation to hippocampal function. *Journal of Neuroscience* 19: 2789–2798, 1999.
- Dickinson A. Actions and habits: the development of behavioural autonomy. *Philosophical Transactions of the Royal Society, London B* 308: 67–78, 1985.
- Dimova R, Vuillet J, Nieoullon A, and Goff LKL. Ultrastructural features of the choilne acetyltransferase-containing neurons and relationships with nigral dopaminergic and cortical afferent pathways in the rat striatum. *Neuroscience* 53: 1059–1071, 1993.
- Doya K. What are the computations of the cerebellum, the basal ganglia, and the cerebral cortex? *Neural networks* 12: 961–974, 1999.
- Doya K. Complementary roles of basal ganglia and cerebellum in learning and motor control. *Current Opinion in Neurobiology* 10: 732–739, 2000.
- Efron B. *The Jackknife, the Bootstrap, and Other Resampling Plans*, volume 38 of *Regional Conference Series in Applied Mathematics*. Philadelphia PA: Society for Industrial and Applied Mathematics, 1982.
- Ferrante RJ, Kowall NW, Beal MF, Jr EPR, Bird ED, and Martin JB. Selective sparing of a class of striatal neurons in huntington's disease. *Science* 230: 561–563, 1985.
- Ferrante RJ, Kowall NW, Beal MF, Martin JB, Bird ED, and Jr EPR. Morphologic and histochemical characteristics of a spared subset of striatal neurons in huntington's disease. *Journal of Neuropathology and Experimental Neurology* 46: 12–27, 1987.
- Figuerdo-Cardenas G, Medina L, and Reiner A. Calretinin is largely localized to a unique population of striatal interneurons in rats. *Brain Research* 709: 145–150, 1996.
- Figueredo-Cardenas G, Morello M, Sancesario G, Bernardi G, and Reiner A. Colocalization of somatostatin, neuropeptide y, neuronal nitric oxide synthase and nadph-diaphorase in striatal interneurons in rats. *Brain Research* 735: 317–324, 1996.
- Finger S. *Origins of neuroscience: A history of explorations into brain function*. New York, New York: Oxford University Press, 1994.
- Fiorillo CD, Tobler PN, and Schultz W. Discrete coding of reward probability and uncertainty by dopamine neurons. *Science* 299: 1898–1902, 2003.
- Foster DJ, Morris RGM, and Dayan P. A model of hippocampally dependent navigation using the temporal difference learning rule. *Hippocampus* 10: 1–6, 2000.
- Fujiyama F, Fritschy J, Stephenson FA, and Bolam JP. Synaptic localization of gaba_a receptor subunits in the striatum of the rat. *Journal of Comparative Neurology* 416: 158–172, 2000.

- Galarreta M and Hestrin S. A network of fast-spiking cells in the neocortex connected by electrical synapses. *Nature* 402: 72–75, 1999.
- Gardiner TW and Kitai ST. Single-unit activity in the globus pallidus and neostriatum of the rat during performance of a trained head-movement. *Experimental Brain Research* 88: 517–530, 1992.
- Georgopoulos AP, Caminiti R, Kalaska JF, and Massey JT. Spatial coding of movement: A hypothesis concerning the coding of movement direction by motor cortical populations. *Experimental Brain Research Suppl.*: 327–336, 1983.
- Georgopoulos AP, Kettner RE, and Schwartz AB. Primate motor cortex and free arm movements to visual targets in three-dimensional space. II. Coding of the direction of movement by a neuronal population. *Journal of Neuroscience* 8: 2928–2937, 1988.
- Gerfen CR. The neostriatal mosaic: Compartmentalization of corticostriatal input and striatonigral output systems. *Nature* 311: 461–464, 1984.
- Gerfen CR. The neostriatal mosaic. I. Compartmental organization of projections from the striatum to the substantia nigra in the rat. *Journal of comparative neurology* 236: 454–476, 1985.
- Gerfen CR. The neostriatal mosaic: (c)ompartmental organization of mesostriatal systems. In: *The Basal Ganglia II*, edited by Carpenter MC and Jayaraman A, New York: Plenum Press, pp. 65–80. 1986.
- Gerfen CR. The neostriatal mosaic: Striatal patch-matrix organization is related to cortical lamination. *Science* 246: 385–388, 1989.
- Gerfen CR, Baimbridge KG, and Miller JJ. The neostriatal mosaic: Compartmental distribution of calcium-binding protein and parvalbumin in the basal ganglia of the rat and monkey. *Proceedings of the National Academy of Sciences* 82: 8780–8784, 1985.
- Gibson JR, Belerlein M, and Connors BW. Two networks of electrically coupled inhibitory neurons in neocortex. *Journal of Neuroscience* 402: 75–79, 1999.
- Graveland GA and DiFiglia M. The frequency and distribution of medium-sized neurons with indented nuclei in the primate and rodent neostriatum. *Brain Research* 327: 307–311, 1985.
- Graybiel AM. The basal ganglia and chunking of action repertoires. *Neurobiology of learning and memory* 70: 119–136, 1998.
- Graybiel AM, Aosaki T, Flaherty AW, and Kimura M. The basal ganglia and adaptive motor control. *Science* 265: 1826–1831, 1994.
- Graybiel AM, Jr CWR, Yoneoka ES, and Elde RP. An immunohistochemical study of enkephalins and other neuropeptides in the striatum of the cat with evidence that the opiate peptides are arranged to form mosaic patterns in register with the striosomal compartments visible by acetylcholinesterase staining. *Neuroscience* 6: 377–397, 1981.
- Graybiel AM and Moratalla R. Dopamine uptake sites in the striatum are distributed differentially in striosome and matrix compartments. *Proceedings of the National Academy of Sciences, USA* 86: 9020–9024, 1989.

- Graybiel AM and Ragsdale CW Jr. Histochemically distinct compartments in the striatum of human, monkey, and cat demonstrated by acetylcholinesterase staining. *Proceedings of the National Academy of Sciences, USA* 75: 5723–5726, 1978.
- Groves PM, Martone M, Young SJ, and Armstrong DM. Three-dimensional pattern of enkephalin-like immunoreactivity in the caudate nucleus of the cat. *Journal of Neuroscience* 8: 892–900, 1988.
- Gutierrez C, Cox CL, Rinzel J, and Sherman SM. Dynamics of low-threshold spike activation in relay neurons of the cat lateral geniculate nucleus. *Journal of Neuroscience* 21: 1022–1032, 2001.
- Guzmán JN, Hernández A, Galarraga E, Tapia D, Laville A, Vergara R, Aceves J, and Bargas J. Dopaminergic modulation of axon collaterals interconnecting spiny neurons of the rat striatum. *Journal of Neuroscience* 23: 8931–8940, 2003.
- Harrington KM and Kowall NW. Parvalbumin immunoreactive neurons resist degeneration in huntington's disease striatum. *Journal of Neuropathology and Experimental Neurology* 50: 309, 1991.
- Harris K. KLUStAKWIK automatic cluster analysis, version 1.0. <http://osiris.rutgers.edu/Buzsaki/software/>, 2002.
- Heimer L, Zahm DS, and Alheid GF. Basal ganglia. In: *The Rat Nervous System*, edited by Paxinos G, New York: Academic Press, pp. 579–628. Second edition, 1995.
- Heindel WC, Salmon DP, Shults CW, Walicke PA, and Butters N. Neuropsychological evidence for multiple memory systems: a comparison of alzheimer's, huntington's and parkinson's disease patients. *Journal of Neuroscience* 9: 582–587, 1989.
- Herkenham M and Pert CB. Mosaic distribution of opiate receptors, parafascicular projections and acetylcholinesterase in rat striatum. *Nature* 291: 415–418, 1981.
- Hikosaka O, Rand MK, Nakamura K, Miyachi S, Kitaguchi K, Sakai K, Lu X, and Shimo Y. Long-term retention of motor skill in macaque monkeys and humans. *Experimental Brain Research* 147: 494–504, 2002.
- Hikosaka O, Sakamoto M, and Usui S. Functional properties of monkey caudate neurons. i. activities related to saccadic eye movements. *Journal of Neurophysiology* 61: 780–798, 1989a.
- Hikosaka O, Sakamoto M, and Usui S. Functional properties of monkey caudate neurons. ii. visual and auditory responses. *Journal of Neurophysiology* 61: 799–813, 1989b.
- Hikosaka O, Sakamoto M, and Usui S. Functional properties of monkey caudate neurons. iii. activities related to expectation of target and reward. *Journal of Neurophysiology* 61: 814–832, 1989c.
- Hollerman JR and Schultz W. Dopamine neurons report an error in the temporal prediction of reward during learning. *Nature Neuroscience* 1: 304–309, 1998.

- Houk JC, Adams JL, and Barto AG. A model of how the basal ganglia generate and use neural signals that predict reinforcement. In: *Models of Information Processing in the Basal Ganglia*, edited by Houk JC, Davis JL, and Beiser DG, Cambridge MA: MIT Press, pp. 249–270, 1995.
- Jackson JC and Redish AD. Detecting dynamical changes within a simulated neural ensemble using a measure of representational quality. *Network: Computation in Neural Systems* 14: 629–645, 2003.
- Jaeger D, Kita H, and Wilson CJ. Surround inhibition among projection neurons is weak or nonexistent in the rat neostriatum. *Journal of neurophysiology* 72: 2555–2558, 1994.
- Jensen O and Lisman JE. Position reconstruction from an ensemble of hippocampal place cells: contribution of theta phase encoding. *Journal of Neurophysiology* 83: 2602–2609, 2000.
- Jog MS, Kubota Y, Connolly CI, Hillegaart V, and Graybiel AM. Building neural representations of habits. *Science* 286: 1746–1749, 1999.
- Johnson A, Seeland KD, and Redish AD. Head-direction ensembles recorded from awake, behaving rats in an open field under cue-conflict situations. *Society for Neuroscience Abstracts* 289.6, 2003.
- Kaneko S, Hikida T, Watanabe D, Ichinose H, Nagatsu T, Kreitman RJ, Pastan I, and Nakanishi S. Synaptic integration mediated by striatal cholinergic interneurons in basal ganglia function. *Science* 289: 633–637, 2000.
- Kawagoe R, Takikawa Y, and Hikosaka O. Expectation of reward modulates cognitive signals in the basal ganglia. *Nature Neuroscience* 1: 411–416, 1998.
- Kawaguchi Y. Large aspiny cells in the matrix of the rat neostriatum in vitro: Physiological identification, relation to the compartments and excitatory postsynaptic currents. *Journal of Neurophysiology* 67: 1669–1682, 1992.
- Kawaguchi Y. Physiological, morphological, and histochemical characterization of three classes of interneurons in the rat neostriatum. *Journal of Neuroscience* 13: 4908–4923, 1993.
- Kawaguchi Y, Wilson CJ, Augood SJ, and Emson PC. Striatal interneurons: chemical, physiological and morphological characterization. *Trends in Neurosciences* 18: 527–535, 1995.
- Kawaguchi Y, Wilson CJ, and Emson PC. Intracellular recording of identified neostriatal patch and matrix spiny cells in a slice preparation preserving cortical inputs. *Journal of Neurophysiology* 62: 1052–1068, 1989.
- Kawaguchi Y, Wilson CJ, and Emson PC. Projection subtypes of rat neostriatal matrix cells revealed by intracellular injection of biocytin. *Journal of Neuroscience* 10: 3421–3438, 1990.
- Kemp JM and Powell TPS. The termination of fibres from the cerebral cortex and thalamus upon dendritic spines in the caudate nucleus: a study with the golgi method. *Philosophical Transactions of the Royal Society of London Series B, Biological Sciences* 262: 429–439, 1971.
- Kermadi I and Joseph JP. Activity in the caudate nucleus of monkey during spatial sequencing. *Journal of Neurophysiology* 74: 911–933, 1995.

- Kermadi I, Jurquet Y, Arzi M, and Joseph J. Neural activity in the caudate nucleus of monkeys during spatial sequencing. *Experimental Brain Research* 94: 352–356, 1993.
- Kesner RP, Bolland BL, and Dakis M. Memory for spatial locations, motor responses, and objects: Triple dissociation among the hippocampus, caudate nucleus, and extrastriate visual cortex. *Experimental Brain Research* 93: 462–470, 1993.
- Kimura M. The role of primate putamen neurons in the association of sensory stimuli with movement. *Neuroscience Research* 3: 436–443, 1986. .
- Kimura M. Behaviorally contingent property of movement-related activity of the primate putamen. *Journal of Neurophysiology* 63: 1277–1296, 1990. .
- Kimura M. Behavioral modulation of sensory responses of primate putamen neurons. *Brain Research* 578: 204–214, 1992. .
- Kimura M. Role of basal ganglia in behavioral learning. *Neuroscience Research* 22: 353–358, 1995.
- Kimura M, Kato M, and Shimazaki H. Physiological properties of projection neurons in the monkey striatum to globus pallidus. *Experimental Brain Research* 82: 672–676, 1990. .
- Kimura M, Kato M, Shimazaki H, Watanabe K, and Matsumoto N. Neural information transferred from the putamen to the globus pallidus during learned movement in the monkey. *Journal of Neurophysiology* 76: 3771–3786, 1996.
- Kimura M, Rajkowski J, and Evarts E. Tonicly discharging putamen neurons exhibit set-dependent responses. *Proceedings of the National Academy of Sciences, USA* 81: 4998–5001, 1984.
- Kincaid AE, Zheng T, and Wilson CJ. Connectivity and convergence of single corticostriatal axons. *Journal of Neuroscience* 18: 4722–4731, 1998.
- Kita H, Kosaka T, and Heizmann CW. Parvalbumin-immunoreactive neurons in the rat neostriatum: a light and electron microscopic study. *Brain Research* 536: 1–15, 1990.
- Kitano K, Câteau H, Kaneda K, Nambu A, Takada M, and Fukai T. Two-state membrane potential transitions of striatal spiny neurons as evidenced by numerical simulations and electrophysiological recordings in awake monkeys. *Journal of Neuroscience* 22: 1–6, 2002.
- Knowlton BJ, Mangels JA, and Squire LR. A neostriatal habit learning system in humans. *Science* 273: 1399–1402, 1996.
- Koós T and Tepper JM. Inhibitory control of neostriatal projection neurons by GABAergic interneurons. *Nature Neuroscience* 2: 467–472, 1999.
- Koós T and Tepper JM. Dual cholinergic control of fast-spiking interneurons in the neostriatum. *Journal of Neuroscience* 22: 529–535, 2002.

- Kowall NW, Ferrante RJ, Beal MF, Jr EPR, Sofroniew MV, Cuello AC, and Martin JB. Neuropeptide y, somatostatin and reduced nicotinamide adenine dinucleotide phosphate diaphorase in the human striatum: a combined immunocytochemical and enzyme histochemical study. *Neuroscience* 20: 817–828, 1987.
- Kubota Y, Mikawa S, and Kawaguchi Y. Neostriatal gabaergic interneurons contain nos, calretinin or parvalbumin. *NeuroReport* 5: 205–208, 1993.
- Lapper S and Bolam JP. Input from the frontal cortex and the parafascicular nucleus to cholinergic interneurons of the dorsal striatum of the rat. *Neuroscience* 51: 533–545, 1992.
- Lapper SR, Smith Y, Sadikot AF, Parent A, and Bolam JP. Cortical input to parvalbumin-immunoreactive neurons in the putamen of the squirrel monkey. *Brain Research* 580: 215–224, 1992.
- Lauwereyns J, Watanabe K, Coe B, and Hikosaka O. A neural correlate of response bias in monkey caudate nucleus. *Nature* 418: 413–417, 2002.
- Ljungberg T, Apicella P, and Schultz W. Responses of monkey dopamine neurons during learning of behavioral reactions. *Journal of Neurophysiology* 67: 145–163, 1992.
- Mann-Metzer P and Yarom Y. Electrotonic coupling interacts with intrinsic properties to generate synchronized activity in cerebellar networks of inhibitory interneurons. *Journal of Neuroscience* 19: 3298–3306, 1999.
- Matsumoto N, Hanakawa T, Maki S, Graybiel AM, and Kimura M. Role of nigrostriatal dopamine system in learning to perform sequential motor tasks in a predictive manner. *Journal of Neurophysiology* 82: 978–998, 1999.
- Matsumoto N, Minamimoto T, Graybiel AM, and Kimura M. Neurons in the thalamic CM-Pf complex supply striatal neurons with information about behaviorally significant events. *Journal of Neurophysiology* 85: 960–976, 2001.
- McDonald RJ and Hong NS. A dissociation of dorso-lateral striatum and amygdala function on the same stimulus-response habit task. *Neuroscience* 124: 507–513, 2004.
- McDonald RJ and White N. Parallel information processing by hippocampal and dorsal striatal memory systems in the Morris water maze. *Society for Neuroscience Abstracts* 19: 362, 1993.
- McDonald RJ and White NM. Hippocampal and nonhippocampal contributions to place learning in rats. *Behavioral Neuroscience* 109: 579–593, 1995.
- McGeorge AJ and Faull RL. The organization of the projection from the cerebral cortex to the striatum in the rat. *Neuroscience* 29: 503–537, 1989.
- Middleton FA and Strick PL. Basal ganglia and cerebellar loops: motor and cognitive circuits. *Brain Research Reviews* 31: 236–250, 2000.
- Mirenowicz J and Schultz W. Importance of unpredictability for reward responses in primate dopamine neurons. *J Neurophysiol* 72: 1024–1027, 1994.

- Miyachi S, Hikosaka O, Miyashita K, Kárádi Z, and Rand MK. Differential roles of monkey striatum in learning of sequential hand movement. *Experimental Brain Research* 115: 1–5, 1997.
- Montague PR, Dayan P, and Sejnowski TJ. A framework for mesencephalic dopamine systems based on predictive hebbian learning. *Journal of Neuroscience* 16: 1936–1947, 1996.
- Mufson EJ and Brandabur MM. Sparing of nadph-diaphorase striatal neurons in parkinson's and alzheimer's diseases. *Neuroreport* 5: 705–708, 1994.
- Nambu A, Tokuno H, and Takada M. Functional significance of the cortico-subthalamo-pallidal 'hyperdirect' pathway. *Neuroscience Research* 43: 111–117, 2002.
- Nisenbaum ES and Berger TW. Functionally distinct subpopulations of striatal neurons are differentially regulated by gabaergic and dopaminergic inputs i. *In vivo* analysis. *Neuroscience* 48: 561–578, 1992.
- Nisenbaum ES, Orr WB, and Berger TW. Evidence for two functionally distinct subpopulations of neurons within the rat striatum. *Journal of Neuroscience* 8: 4138–4150, 1988.
- O'Keefe J and Nadel L. *The Hippocampus as a Cognitive Map*. Oxford: Clarendon Press, 1978.
- Onn SP, Berger TW, and Grace AA. Identification and characterization of striatal cell subtypes using in vivo intracellular recording and dye-labeling in rats: III. Morphological correlates and compartmental localization. *Synapse* 16: 231–254, 1994.
- Onn SP and Grace AA. Dye coupling between rat striatal neurons recorded in vivo: Compartmental organization and modulation by dopamine. *Journal of Neurophysiology* 71: 1917–1933, 1994.
- Packard MG. Glutamate infused posttraining into the hippocampus or caudate-putamen differentially strengthens place and response learning. *Proceedings of the National Academy of Sciences* 96: 12881–12886, 1999.
- Packard MG, Hirsh R, and White NM. Differential effects of fornix and caudate nucleus lesions on two radial maze tasks: Evidence for multiple memory systems. *Journal of Neuroscience* 9: 1465–1472, 1989.
- Packard MG and Knowlton BJ. Learning and memory functions of the basal ganglia. *Annual Reviews Neuroscience* 25: 563–593, 2002.
- Packard MG and McGaugh JL. Double dissociation of fornix and caudate nucleus lesions on acquisition of two water maze tasks: Further evidence for multiple memory systems. *Behavioral Neuroscience* 106: 439–446, 1992.
- Packard MG and McGaugh JL. Inactivation of hippocampus or caudate nucleus with lidocaine differentially affects expression of place and response learning. *Neurobiology of Learning and Memory* 65: 65–72, 1996.
- Parent A, Cicchetti F, and Beach TG. Calretinin-immunoreactive neurons in the human striatum. *Brain Research* 674: 347–351, 1995.
- Parent A and Hazrati LN. functional anatomy of the basal ganglia. I. The cortico-basal ganglia-thalamo-cortical loop. *Brain Research Reviews* 20: 91–127, 1995.

- Paxinos G and Watson C. *The Rat Brain in Stereotaxic Coordinates*. New York: Academic Press, fourth edition, 1998.
- Pennartz CMA, Lee E, Verheul J, Lipa P, Barnes CA, and McNaughton BL. The Ventral Striatum in Off-Line Processing: Ensemble Reactivation during Sleep and Modulation by Hippocampal Ripples. *J Neurosci* 24: 6446–6456, 2004.
- Pinault D. A novel single-cell staining procedure performed in vivo under electrophysiological control: morpho-functional features of juxtacellularly labeled thalamic cells and other central neurons with biocytin or neurobiotin. *Journal of Neuroscience Methods* 65: 113–136, 1996.
- Plenz D and Kitai ST. Up and down states in striatal medium spiny neurons simultaneously recorded with spontaneous activity in fast-spiking interneurons studied in cortex-striatum-substantia nigra organotypic cultures. *Journal of Neuroscience* 18: 266–283, 1998.
- Poldrack RA and Packard MG. Competition among multiple memory systems: Converging evidence from animal and human studies. *Neuropsychologia* 41: 245–251, 2003.
- Prensa L and Parent A. The nigrostriatal pathway in the rat: A single-axon study of the relationship between dorsal and ventral tier nigral neurons and the striosome/matrix striatal compartments. *Journal of Neuroscience* 21: 7247–7260, 2001.
- Ragozzino KE, Leutgeb S, and Mizumori SJ. Dorsal striatal head direction and hippocampal place representations during spatial navigation. *Experimental Brain Research* 139: 372–376, 2001.
- Ragozzino ME, Ragozzino KE, Mizumori SJY, and Kesner RP. The role of the dorsomedial striatum in behavioral flexibility for response and visual cue discrimination learning. *Behavioral Neuroscience* 116: 105–115, 2002.
- Ragsdale CW and Graybiel AM. A simple ordering of neocortical areas established by the compartmental organization of their striatal projections. *Proceedings of the National Academy of Sciences* 87: 6196–6199, 1984.
- Ramanathan S, Hanley JJ, Deniau JM, and Bolam JP. Synaptic convergence of motor and somatosensory cortical afferents onto GABAergic interneurons in the rat striatum. *Journal of Neuroscience* 22: 8158–8169, 2002.
- Rand MK, Hikosaka O, Miyachi S, Lu X, Nakamura, Kitaguchi K, and Shimo Y. Characteristics of sequential movements during early learning period in monkeys. *Experimental Brain Research* 131: 293–304, 2000.
- Rand MK, Hikosaka O, Miyachi S, Lu X, and Miyashita K. Characteristics of a long-term procedural skill in the monkey. *Experimental Brain Research* 118: 293–297, 1998.
- Ravel S, Legallet E, and Apicella P. Responses of tonically active neurons in the monkey striatum discriminate between motivationally opposing stimuli. *Journal of Neuroscience* 23: 8489–8497, 2003.
- Ravel S, Sardo P, Legallet E, and Apicella P. Reward unpredictability inside and outside of a task context as a determinant of the responses of tonically active neurons in the monkey striatum. *Journal of Neuroscience* 21: 5730–5739, 2001.

- Raz A, Feingold A, Zelanskaya V, Vaadia E, and Bergman H. Neuronal synchronization of tonically active neurons in the striatum of normal and parkinsonian primates. *Journal of Neurophysiology* 76: 2082–2088, 1996.
- Raz A, Frechter-Mazar V, Feingold A, Abeles M, Vaadia E, and Bergman H. Activity of pallidal and striatal tonically active neurons is correlated in MPTP-treated monkeys but not in normal monkeys. *Journal of Neuroscience* 21: RC128 (1–5), 2001.
- Redish AD. *Beyond the Cognitive Map: From Place Cells to Episodic Memory*. Cambridge MA: MIT Press, 1999.
- Redish AD, Rosenzweig ES, Bohanick JD, McNaughton BL, and Barnes CA. Dynamics of hippocampal ensemble realignment: Time vs. space. *Journal of Neuroscience* 20: 9289–9309, 2000.
- Redish AD and Schmitzer-Torbert NC. MCLUST spike sorting toolbox, version 3.0. <http://www.cbc.umn.edu/~redish/mclust/>, 2002.
- Redish AD, Schmitzer-Torbert NC, and Jackson JC. Classification of dorsal striatal neurons from extracellular recordings in awake behaving rats. *Society for Neuroscience Abstracts* 28, 2002. Program no. 676.3.
- Reiner A, Albin RL, Anderson KD, D'Amato CJ, Penney JB, and Young AB. Differential loss of striatal projection neurons in huntington's disease. *Proceedings of the National Academy of Sciences* 85: 5733–5737, 1988.
- Reiner A and Anderson KD. The patterns of neurotransmitter and neuropeptide co-occurrence among striatal projection neurons: Conclusions based on recent findings. *Brain Research Reviews* 15: 251–265, 1990.
- Reiner A and Anderson KD. Co-occurrence of gaba, parvalbumin and the neurotensin-related hexapeptide lant6 in pallidal, nigral and striatal neurons in pigeons and monkeys. *Brain Research* 624: 317–325, 1993.
- Reynolds JNJ and Wickens JR. The corticostriatal input to giant aspiny interneurons in the rat: A candidate pathway for synchronising the response to reward-related cues. *Brain Research* 1011: 115–128, 2004.
- Rieke F, Warland D, de Ruyter van Steveninck R, and Bialek W. *Spikes*. Cambridge MA: MIT Press, 1997.
- Rudkin TM and Sadikot AF. Thalamic input to parvalbumin-immunoreactive gabaergic interneurons: organization in normal striatum and effect of neonatal decortication. *Neuroscience* 88: 1165–1175, 1999.
- Saka E, Goodrich C, Harlan P, Madras BK, and Graybiel AM. Repetitive behaviors in monkeys are linked to specific striatal activation patterns. *Journal of Neuroscience* 24: 7557–7565, 2004.
- Saka E, Iadarola M, Fitzgerald DJ, and Graybiel AM. Local circuit neurons in the striatum regulate neural and behavioral responses to dopaminergic stimulation. *Proceedings of the National Academy of Sciences* 99: 9004–9009, 2002.

- Salinas E and Abbott L. Vector reconstruction from firing rates. *Journal of Computational Neuroscience* 1: 89–107, 1994.
- Sanger TD. Probability density estimation for the interpretation of neural population codes. *Journal of Neurophysiology* 76: 2790–2793, 1996.
- Schmitzer-Torbert NC, Jackson JC, Henze D, Harris KD, and Redish AD. Quantitative measures of cluster quality for use in extracellular recordings. *Neuroscience* in press.
- Schmitzer-Torbert NC and Redish AD. Development of path stereotypy in a single day in rats on a multiple-T maze. *Archives Italiennes de Biologie* 140: 295–301, 2002.
- Schmitzer-Torbert NC and Redish AD. Neuronal activity in the rodent dorsal striatum in sequential navigation: Separation of spatial and reward responses on the multiple-T task. *Journal of Neurophysiology* 91: 2259–2272, 2004.
- Schroeder JP, Wingard JC, and Packard MG. Post-training reversible inactivation of hippocampus reveals interference between memory systems. *Hippocampus* 12: 480–484, 2002.
DOI: 10.1002/hipo.10024.
- Schultz W. Predictive reward signal of dopamine neurons. *Journal of Neurophysiology* 80: 1–27, 1998.
- Schultz W, Dayan P, and Montague R. A neural substrate of prediction and reward. *Science* 275: 1593–1599, 1997.
- Schultz W, Romo R, Ljungberg T, Mirenowicz J, Hollerman JR, and Dickinson A. Reward-related signals carried by dopamine neurons. In: *Models of Information Processing in the Basal Ganglia*, edited by Houk JC, Davis JL, and Beiser DG, Cambridge MA: MIT Press, pp. 233–248. 1995.
- Schultz W, Tremblay L, and Hollerman JR. Changes in behavior-related neuronal activity in the striatum during learning. *Trends in Neurosciences* 26: 321–328, 2003.
- Selden N, Geula C, Hersh L, and Mesulam MM. Human striatum: Chemoarchitecture of the caudate nucleus, putamen and ventral striatum in health and alzheimer's disease. *Journal of Neuroscience Methods* 60: 621–636, 1994.
- Sesack SR and Pickel VM. In the rat medial nucleus accumbens, hippocampal and catecholaminergic terminals converge on spiny neurons and are in apposition to each other. *Brain Research* 527: 266–279, 1990.
- Shimo Y and Hikosaka O. Role of tonically active neurons in primate caudate in reward-oriented saccadic eye movement. *Journal of Neuroscience* 21: 7804–7814, 2001.
- Sidibé M and Smith Y. Thalamic inputs to striatal interneurons in monkeys: Synaptic organization and co-localization of calcium binding proteins. *Neuroscience* 89: 1189–1208, 1999.
- Silva MT, Rose S, Hindmarsh JG, and Jenner P. Inhibition of neuronal nitric oxide synthase increases dopamine efflux from rat striatum. *Journal of Neural Transmission* 110: 353–362, 2003.

- Skirboll LR and Bunney BS. The effects of acute and chronic haloperidol treatment on spontaneously firing neurons in the caudate nucleus of the rat. *Life Sciences* 25: 1419–1434, 1979.
- Smith Y, Bevan MD, Shink E, and Bolam JP. Microcircuitry of the direct and indirect pathways of the basal ganglia. *Neuroscience* 86: 353–387, 1998.
- Smith Y and Parent A. Neuropeptide y-immunoreactive neurons in the striatum of cat and monkey: Morphological characteristics, intrinsic organization and co-localization with somatostatin. *Brain Research* 372: 241–252, 1986.
- Smith Y, Raju DV, Pare J, and Sidibe M. The thalamostriatal system: A highly specific network of the basal ganglia circuitry. *TRENDS in Neurosciences* 27: 520–527, 2004.
- Smith-Roe SL, Sadeghian K, and Kelley AE. Spatial learning and performance in the radial arm maze is impaired after N-methyl-D-aspartate (NMDA) receptor blockade in striatal subregions. *Behavioral Neuroscience* 113: 703–717, 1999.
- Sokal RR and Rohlf FJ. *Biometry: The principles and practice of statistics in biological research*. New York: W. H. Freeman, third edition, 1995.
- Stern EA, Kincaid AE, and Wilson CJ. Spontaneous subthreshold membrane potential fluctuations and action potential variability of rat corticostriatal and striatal neurons in vivo. *Journal of Neurophysiology* 77: 1697–1715, 1997.
- Sutton RS and Barto AG. *Reinforcement Learning: An introduction*. Cambridge MA: MIT Press, 1998.
- Swanson LW. Cerebral hemisphere regulation of motivated behavior. *Brain Research* 886: 113–164, 2000.
- Takikawa Y, Kawagoe R, and Hikosaka O. Reward-dependent spatial selectivity of anticipatory activity in monkey caudate nucleus. *Journal of Neurophysiology* 87: 508–515, 2002.
- Thomas TM, Smith Y, Levey AI, and Hersch SM. Cortical inputs to m2-immunoreactive striatal interneurons in rat and monkey. *Synapse* 37: 252–261, 2000.
- Tobler PN, Dickinson A, and Schultz W. Coding of predicted reward omission by dopamine neurons in a conditioned inhibition paradigm. *Journal of Neuroscience* 23: 10402–10410, 2003.
- Tolman EC, Ritchie BF, and Kalish D. Studies in spatial learning. II. Place learning versus response learning. *Journal of Experimental Psychology* 36: 221–229, 1946.
- Tremblay L, Hollerman JR, and Schultz W. Modifications of reward expectation-related neuronal activity during learning in primate striatum. *Journal of Neurophysiology* 80: 964–977, 1998.
- Trytek ES, White IM, Schroeder DM, Heidenreich BA, and Rebec GV. Localization of motor- and nonmotor-related neurons within the matrix striosome organization of rat striatum. *Brain Research* 707: 221–227, 1996.
- Ueda Y and Kimura M. Encoding of direction and combination of movements by primate putamen neurons. *European Journal of Neuroscience* 18: 980–994, 2003.

- VanVleet TM, Heldt SA, Guerrettaz KR, Corwin JV, and Reep RL. Unilateral destruction of the dorsocentral striatum in rats produces neglect but not extinction to bilateral simultaneous stimulation. *Behavioural Brain Research* 136: 375–387, 2002.
- Vuillet J, Kerkerian L, Kachidian P, Bosler O, and Nieoullon A. Ultrastructural correlates of functional relationships between nigral dopaminergic or cortical afferent fibers and neuropeptide y-containing neurons in the rat striatum. *Neuroscience Letters* 100: 99–104, 1989.
- Waelti P, Dickinson A, and Schultz W. Dopamine responses comply with basic assumptions of formal learning theory. *Nature* 412: 43–48, 2001.
- Watanabe K and Kimura M. Dopamine receptor-mediated mechanisms involved in expression of learned activity of primate striatal neurons. *Journal of Neurophysiology* 79: 2568–2580, 1998.
- West MO, Carelli RM, Pomerantz M, Cohen SM, Gardner JP, Chapin JK, and Woodward DJ. A region in the dorsolateral striatum of the rat exhibiting single-unit correlations with specific locomotor limb movements. *Journal of Neurophysiology* 64: 1233–1246, 1990.
- White IM and Rebec GV. Responses of rat striatal neurons during performance of a lever-release version of the conditioned avoidance response task. *Brain Research* 616: 71–82, 1993.
- White NM. A functional hypothesis concerning the striatal matrix and patches: Mediation of S-R memory and reward. *Life Sciences* 45: 1943–1957, 1989.
- White NM and Hiroi N. Preferential localization of self-stimulation sites in striosomes/patches in the rat striatum. *Proceedings of the National Academy of Sciences, USA* 95: 6486–6491, 1998.
- Wiener SI. Spatial and behavioral correlates of striatal neurons in rats performing a self-initiated navigation task. *Journal of Neuroscience* 13: 3802–3817, 1993.
- Wilson CJ and Groves PM. Spontaneous firing patterns of identified spiny neurons in the rat neostriatum. *Brain Research* 220: 67–80, 1981.
- Wilson CJ and Kawaguchi Y. The origins of two-state spontaneous membrane potential fluctuations of neostriatal spiny neurons. *Journal of Neuroscience* 16: 2397–2410, 1996.
- Wilson MA and McNaughton BL. Dynamics of the hippocampal ensemble code for space. *Science* 261: 1055–1058, 1993.
- Yeshenko O, Guazzelli A, and Mizumori SJY. Context-dependent reorganization of spatial and movement representations by simultaneously recorded hippocampal and striatal neurons during performance of allocentric and egocentric tasks. *Behavioral Neuroscience* 118: 751–769, 2004.
- Yin HH, Knowlton B, and Balleine BW. Lesions of dorsolateral striatum preserve outcome expectancy but disrupt habit formation in instrumental learning. *European Journal of Neuroscience* 19: 181–189, 2004.
- Yin HH and Knowlton BJ. Contributions of Striatal Subregions to Place and Response Learning. *Learn Mem* 11: 459–463, 2004.

Zhang K, Ginzburg I, McNaughton BL, and Sejnowski TJ. Interpreting neuronal population activity by reconstruction: Unified framework with application to hippocampal place cells. *Journal of Neurophysiology* 79: 1017–1044, 1998.

Zheng T and Wilson CJ. Corticostriatal combinatorics: The implications of corticostriatal axonal arborizations. *Journal of Neurophysiology* 87: 1007–1017, 2002.