

**UCLA**

**UCLA Electronic Theses and Dissertations**

**Title**

Dynamic population activity in the striatum during associative behavior

**Permalink**

<https://escholarship.org/uc/item/2v958261>

**Author**

Bakhurin, Konstantin Ivanovitch

**Publication Date**

2017

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA

Los Angeles

Dynamic population activity in the striatum during associative behavior

A dissertation submitted in partial satisfaction of the  
requirements for the degree Doctor of Philosophy  
in Neuroscience

by

Konstantin Ivanovitch Bakhurin

2017

© Copyright by

Konstantin Ivanovitch Bakhurin

2017

## ABSTRACT OF THE DISSERTATION

Dynamic population activity in the striatum during associative behavior

by

Konstantin Ivanovitch Bakhurin

Doctor of Philosophy in Neuroscience

University of California, Los Angeles, 2017

Professor Sotirios Masmanidis, Chair

The basal ganglia are a set of subcortical nuclei that are thought to play important roles in motor control, action selection, goal-directed actions, motivation, and non-declarative learning. The striatum is the main input area to the basal ganglia, receiving diverse excitatory input from nearly every cortical area, thalamus, amygdala, and other subcortical structures. The striatum is an important site of plasticity in the basal ganglia, and also contains a complex local microcircuitry. Together these three elements are thought to interact to generate striatal output signals that modulate the rest of the basal ganglia nuclei. Despite decades of study, both the roles that these computations play in generating behavior, and how these dynamics arise, are not well understood. Here, I describe my work using large-scale *in vivo* recordings from populations of striatal neurons in mice performing a Pavlovian learning task. In chapter 1, I study the structure of spontaneous spiking activity to show that striatal populations that are involved in specific behavior are more likely to show significant correlated activity. This suggests that these neurons share specific inputs, most likely as a result of plasticity at glutamatergic synapses. In chapter 2, I focus on task-related population dynamics in striatal activity to address

how these dynamics may be involved in animals' ability to time their actions. I use machine learning techniques to make direct comparisons between the striatum and one of its input areas, the orbitofrontal cortex, showing that the striatum out-performs the OFC in telling time. These results suggest that the representation of time is not uniform in the brain, and that the striatum may have a privileged role in time representation. Lastly, in chapter 3, I use optogenetic inhibition to causally test the necessity of corticostriatal input for generating striatal population dynamics during anticipatory behaviors in our task. Here, I find that suppressing inputs reduces firing rates in the striatum, but does not eliminate the striatum's dynamic properties. These suggest that local network interactions may still play an important role in shaping striatal activity, and that striatal output is driven by a balance of excitation and local microcircuit activity.

The dissertation of Konstantin Ivanovitch Bakhurin is approved.

Michael Fanselow

Michele Basso

Dean Buonomano

Kate Wassum

Sotirios Masmanidis, Committee Chair

University of California, Los Angeles

2017

<b>Table of Contents</b>	<b>Page</b>
List of figures and acronyms.....	vi
Acknowledgements.....	viii
Vita.....	x
Introduction.....	1
Chapter 1.....	19
Chapter 2.....	48
Chapter 3.....	89
General discussion.....	119
Bibliography.....	128

<b>List of Figures</b>	<b>Page</b>
Figure 1.1: Head restrained mice demonstrate single-session discrimination learning	26
Figure 1.2: Large-scale striatal recordings with silicon microprobes.....	28
Figure 1.3: Identification of a cue-discriminating subpopulation of striatal MSNs.....	30
Figure 1.4: Mapping discriminatory activity across the striatal cross-section.....	32
Figure 1.5: Evolution of activity in discriminatory MSNs during training.....	34
Figure 1.6: Correlated resting state activity in the striatum.....	36
Figure 1.7: Discriminating FSIs and MSNs form functional ensembles.....	39
Figure 1.8: Detecting short-latency interactions among striatal unit classes.....	41
Figure 2.1: Large-scale recording of OFC and striatal networks during reward- predictive behavior.....	60
Figure 2.2: Schematic of the support vector machine decoding of elapsed time	62
Figure 2.3: Striatal networks encode elapsed time .....	65
Figure 2.4: Striatal networks encode elapsed time better than OFC networks.....	69
Figure 2.5: Population encoding of elapsed time is distributed throughout striatum and OFC.....	71
Figure 2.6: Population coding of elapsed time is specific to CS+ trials and not fully explained by licking behavior.....	73
Figure 2.7: Striatal population coding of elapsed time shows higher sensitivity to lick onset variability than OFC.....	75
Figure 2.8: Striatal networks outperform OFC networks at predicting lick onset time...	77
Figure 2.9: Simultaneous multi-region recordings indicate that striatum encodes elapsed time better than OFC.....	79
Figure 2.10: Simultaneous multi-region recordings show distinct pre-lick dynamics across striatal and OFC networks.....	81
Figure 3.1: M2 inhibition reduces anticipatory licking probability.....	100
Figure 3.2: M2 projects to dorsolateral striatum.....	102
Figure 3.3: M2 and DLS show population activity related to movement preparation....	104
Figure 3.4: M2 and DLS are synchronized prior to cued licking behavior.....	106
Figure 3.5: M2 provides dynamic drive onto DLS population activity.....	108
Figure 3.6: MSNs are maximally dependent on M2 input prior to lick initiation.....	110
Figure 3.7: Dynamic codes for movement initiation after input suppression.....	112
Figure 3.8: Input suppression reveals a diverse impact of M2 projections on MSN firing.....	114
Figure 3.9: R <sup>2</sup> values do not group according to MSN response classification.....	116

## List of acronyms

CV – coefficient of variation  
DA - dopamine  
DLS – dorsolateral striatum  
FSI – fast spiking interneuron  
GABA –  
ISI – inter spike interval  
ITI – inter trial interval  
LTD – long-term depression  
LTP – long-term potentiation  
M2 – premotor cortex  
MSN – medium spiny neuron  
OFC – orbitofrontal cortex  
PV – parvalbumin  
R-O – response-outcome  
S-O – stimulus-outcome  
S-R – stimulus-response  
SVM – support vector machine  
TAN – tonically active neuron

## **Acknowledgements**

Chapter 1 is a version of Konstantin I. Bakhurin, Victor Mac, Peyman Golshani, Sotiris C. Masmanidis. Temporal correlations among functional specialized striatal neural ensembles in reward conditioned mice. *Journal of Neurophysiology* 2016; 115(3): 1521-1532. Konstantin Bakhurin conceived and designed research, performed experiments, analyzed data, interpreted results of experiments, prepared figures, and drafted the manuscript. Victor Mac performed experiments. Peyman Golshani conceived and designed research. Sotiris Masmanidis was the research director on the project. The American Physiological Society has granted permission to reproduce this article.

Chapter 2 is a version of Konstantin I. Bakhurin, Vishwa Goudar, Justin L. Shobe, Leslie D. Claar, Dean V. Buonomano, and Sotiris C. Masmanidis. Differential encoding of time by prefrontal and striatal network dynamics. *Journal of Neuroscience* 2017 37(4): 854-870. Konstantin Bakhurin was co-first author, conceived and designed research, performed experiments, analyzed data, interpreted results of experiments, prepared figures, and drafted the manuscript. Vishwa Goudar was co-first author, conceived and designed research, analyzed data, interpreted results of experiments, prepared figures, and drafted the manuscript. Justin Shobe performed experiments. Leslie Claar performed experiments. Dean Buonomano conceived and designed research, analyzed data, interpreted results of experiments, and drafted the manuscript. Sotiris Masmanidis was the research director on the project. The Society for Neuroscience has granted permission to reproduce this article.

Chapter 3 is a working version of Konstantin I. Bakhurin, Kwang Lee, Leslie D. Claar, Vishwa Goudar, Dean V. Buonomano, and Sotiris C. Masmanidis, External circuit influences on encoding of movement preparation in the striatum, which is in preparation. Konstantin I. Bakhurin conceived and designed research, performed experiments, analyzed data, interpreted results of experiments, prepared figures, and drafted the manuscript. Kwang Lee conceived and designed research, performed experiments, analyzed data, and interpreted results of

experiments. Leslie Claar conceived and designed research, performed experiments, analyzed data, and interpreted results of experiments. Vishwa Goudar analyzed data. Dean Buonomano interpreted results of experiments. Sotiris Masmanidis was the research director on the project.

I would like to thank the people that have worked with me on these projects, from whom I have learned a lot about collaboration, communication, and exciting data analysis methods. I specifically acknowledge Sotiris Masmanidis for letting me be a part of his nascent lab, Dean Buonomano for the opportunity to work with his team on understanding how the brain tells time, and to Justin Shobe for the many great conversations about experiment design and data interpretation. I also acknowledge the National Institutes of Health for providing funding for the last 2 years of my doctoral studies under a Ruth L. Kirschstein National Research Service Award (T32-NS058280), which was administered intramurally by Jack Feldman and the Neural Microcircuits Training Grant Selection Committee. I thank the administrative staff of the Neurobiology department for rapid assistance with the wide-ranging and somewhat erratic needs of a graduate student. I also acknowledge Larry Kruger for his constant advice over the years and for sharing his knowledge about the history of UCLA and neuroscience as a whole.

I would like to thank my friends and family that have supported me throughout the years. Without your help and understanding as I moved through the trials and tribulations of the Ph.D. experience, I would not have been able to complete the effort. Finally, I thank Mariana in Merida, who helped me realize that all of us have special abilities. I hope my efforts in this process become the foundation for my capacity for doing great science.

## Vita

### Education

- Ph.D. Candidate Neuroscience. University of California, Los Angeles, CA (UCLA) 06.2017  
Advisor: Sotirios Masmanidis.  
Dissertation: Dynamic population activity in the striatum during  
associative behavior
- B.S. Biology, University of Michigan, Ann Arbor, MI 05.2009

### Research

- Research Assistant, Harvard Medical School, Belmont, MA 08.2009 – 07.2011  
Advisor: Uwe Rudolph
- Undergraduate Research Assistant, University of Michigan, Ann Arbor, MI 09.2006 – 07.2009  
Advisor: John Faulkner

### Publications

- Shobe JL, Bakhurin KI, Claar LD, Masmanidis SC. “Selective modulation of orbitofrontal network activity during negative occasion setting.” *In revision*
- Bakhurin KI\*, Goudar V\*, Shobe JL, Claar LD, Buonomano DV, Masmanidis SC. Differential encoding of time by prefrontal and striatal network dynamics. *Journal of Neuroscience* 2017 37(4): 854-870.
- Bakhurin KI, Mac V, Golshani P, Masmanidis SC. Temporal correlations among functional specialized striatal neural ensembles in reward conditioned mice. *Journal of Neurophysiology* 2016; 115(3): 1521-1532.
- Shobe JL, Claar LD, Parhami S, Bakhurin KI, Masmanidis SC. Brain activity mapping at multiple scales with silicon microprobes containing 1,024 electrodes. *Journal of Neurophysiology* 2015; 114(3): 2043-2052
- Mendias CL, Bakhurin KI, Gumucio JP, Shallal-Ayzin MV, Davis CS, Faulkner JA. Haploinsufficiency of myostatin protects against aging-related declines in muscle function and enhances the longevity of mice. *Aging Cell* 2015; 14(4): 704-706
- Engin E\*, Bakhurin KI\*, Smith KS, Hines RM, Reynolds LM, Tang W, Sprengel R, Moss SJ, Rudolph U. Neural Basis of Benzodiazepine Reward: Requirement for  $\alpha 2$  Containing GABA<sub>A</sub> Receptors in the Nucleus Accumbens. *Neuropsychopharmacology* 2014; 39(8): 1805-1815.
- Mendias CL, Gumucio JP, Bakhurin KI, Lynch EP, Brooks SV. Physiological loading of tendons induces scleraxis expression in epitenon fibroblasts. *Journal of Orthopedic Research* 2012; 30(4): 606-612
- Mendias CL, Bakhurin KI and Faulkner JA. Tendons of myostatin-deficient mice are small, brittle and hypocellular. *Proceedings of the National Academy of Sciences* 2008; 105(1): 388-393.

### Poster Presentations

- Bakhurin KI, Lee K, Claar LD, Goudar V, Buonomano DV, Masmanidis SC. External circuit influences on encoding of movement preparation in the striatum. *International Basal Ganglia Society Meeting*, Merida, Mexico, March 2017
- Bakhurin KI, Goudar V, Shobe JL, Claar LD, Buonomano DV, Masmanidis SC. Prefrontal and striatal network dynamics differentially encode behavior-related time. *Society for Neuroscience Annual Meeting*, San Diego, CA, November 2016

- Bakhurin KI, Goudar V, Shobe JL, Claar LD, Buonomano DV, Masmanidis SC. Prefrontal and striatal network dynamics differentially encode behavior-related time. *Gordon Research Conference, Basal Ganglia*, Ventura, CA, February 2016
- Bakhurin KI, Mac V, Golshani P, Masmanidis SC. Striatal network organization predicts initial learning of Pavlovian associations. *Society for Neuroscience Annual Meeting*, Chicago, IL, October 2015.
- Bakhurin KI, Mac V, Golshani P, Masmanidis SC. Learning-dependent behavioral correlates of striatal functional connectivity. *Computational and Systems Neuroscience*. Salt Lake City, UT, March 2015
- Bakhurin KI, Smith WC, Mac V, Golshani P, Masmanidis SC. Experience dependent dynamics of striatal networks during reward anticipation. *Gordon Research Conference, Basal Ganglia*, Ventura, CA, February 2014
- Bakhurin KI, Smith KS, Engin E, Tang W, Sprengel R, Hines R, Moss, SJ, Rudolph U.  $\alpha 2$ -containing GABA<sub>A</sub> receptors in nucleus accumbens mediate self-administration of midazolam but not locomotor sensitization to cocaine. *Society for Neuroscience Annual Meeting*, Washington D.C., November 2011
- Mendias CL\*, Bakhurin KI\*, Arruda EM, Brooks SV, Faulkner JA, Larkin LM. Scleraxis is expressed in adult tendons and is upregulated in response to mechanical loading. *Experimental Biology, American Physiological Society Annual Meeting*, New Orleans, LA, April 2009

### **Funding**

Neural Microcircuits Training Grant Fellowship, NINDS T32-NS058280      04.2015 – 04.2017  
(stipend and tuition)

### **Awards**

Charles H. Sawyer Memorial International Travel Award Neurobiology Department, UCLA	10.2016
Graduate Student Travel Award Brain Research Institute and Semel Institute for Neuroscience UCLA	10.2013
Undergraduate Research Summer Research Fellowship American Physiological Society	2008
Undergraduate Research Opportunities Summer Fellowship University of Michigan	2007

## **Introduction**

The basal ganglia are a series of non-laminar structures situated subcortically in the cerebrum. These nuclei that comprise the basal ganglia can be divided into four major substructures, the striatum, the globus pallidus, the subthalamic nucleus, and the substantia nigra. While the anatomical interconnections of the basal ganglia are well characterized, a large amount of speculation remains as to the actual function of the basal ganglia in the brain. Here, I focus on the main input area of the basal ganglia, the striatum. I outline two problems that I have attempted to address in my graduate studies: 1) How learning and experience shapes striatal networks, and 2) The function and origin of complex striatal dynamics in cognition. My work focuses on using *in vivo* recording approaches to study striatal networks during the generation of predictive behaviors in mice. In addition to employing novel experimental approaches to these problems, these studies are also based on applying large-scale statistical analysis methods on the resulting data.

### **Problem 1: Organization of the striatal microcircuit by external inputs**

The striatum receives excitatory input from a diverse number of brain areas. Inputs from distinct brain regions are thought to converge and interact at the level of the striatal microcircuit, forming the basis of associative learning and contributing to motor output. However, how these interactions ultimately shape striatal output and contribute to behavior is a problem that is only beginning to be addressed. In this section, I will describe two cell types of the striatal microcircuit whose activity depends on external excitatory signaling, the anatomy of excitatory corticostriatal inputs, and how plasticity at the corticostriatal synapse is thought to play a role in learning skilled actions. In chapter 1, I describe experiments and analyses that show how plasticity mechanisms may shape striatal network organization in behaving animals.

*Striatal microcircuitry and its glutamatergic drive*

The principal cell class in the striatum is the medium spiny neuron (MSN), which makes up 95% of the neuronal population in the structure (Kita and Kitai, 1988). These neurons are the sole projecting population of the striatum, and are GABAergic (Kita and Kitai, 1988). MSNs display a relatively hyperpolarized resting membrane potential of -70mV (Calabresi et al., 1987), and eliciting action potentials from MSNs with current injections is impeded by outward potassium currents (Nisenbaum et al., 1994; Surmeier et al., 1988). Morphologically, the MSN appears to be a neuron that is especially specialized in receiving inputs: its dendritic arbor, approximately 500  $\mu\text{m}$  in diameter (C. J. Wilson and Groves, 1980) is densely adorned with spines (Cajal, 1899; C. J. Wilson et al., 1983). Most corticostriatal synapses are formed with the dendritic spines of the MSNs, which makes the MSNs the major input recipient of the striatum (Kemp and Powell, 1971; Somogyi et al., 1981).

As the sole output cell type of the striatum, MSNs perform a transformation on incoming excitatory signals to form spiking activity. Even though the MSNs perform both input and output functions, it is believed that these signals are shaped by local microcircuit interactions. Two major types of interaction are thought to exist: feedback and feedforward inhibition. With the development of intracellular filling techniques, it was found that in addition to sending long-range axonal projections out of the striatum, MSNs also produced extensive collateral arborizations that overlap and often extend past the territory of their dendrites (Kawaguchi et al., 1990; C. J. Wilson and Groves, 1980). The predominant targets of these GABAergic collaterals are the dendrites and dendritic spines of other MSNs (Somogyi et al., 1981).

These feedback projections were proposed to mediate lateral inhibition. While initial *in vitro* experiments examining the existence of such lateral inhibition revealed no evidence for inhibition (D. Jaeger et al., 1994), the presence of relatively sparse connections that are sensitive to GABA antagonists are now generally accepted (Taverna et al., 2004; Tunstall et al., 2002). Interestingly, both studies determined that MSN-MSN inhibition was largely unidirectional, which, together with the sparse coupling probability (~30% pairwise connection)

and a high failure-rate of synaptic transmission, still calls in to question the importance of these interactions in driving striatal output (Tepper et al., 2008). Thus, while the striatal microcircuit must contain numerous feedback connections, their functional role in striatal function is not well understood.

With the lack of strong experimental evidence for viable feedback inhibition in the striatum, research focus on local microcircuit interactions has shifted to the relatively stronger source of inhibition in the striatum: local GABAergic interneurons. While there are several types of interneurons in the striatum (Silberberg and Bolam, 2015), the best studied is the parvalbumin (PV)-expressing interneuron (Cowan et al., 1990). In contrast with the MSNs, the PV interneurons make up a far smaller fraction of the striatal neuron population at 0.7% (Tepper et al., 2008). They receive multiple excitatory synapses from individual corticostriatal projections that also contact MSNs (Ramanathan et al., 2002). It is believed that these cells are more sensitive to weak stimulation of striatal afferents than MSNs, in that they more readily express immediate early genes than projection neurons (Parthasarathy and Graybiel, 1997). Indeed a later study demonstrated that these neurons generated more action potentials than MSNs to cortical stimulation (Mallet et al., 2005). While they are hyperpolarized in the acute slice preparation, it is less so than MSNs, and sufficient depolarization results in a sustained, high firing state that tends to be organized into bursts (Taverna et al., 2007).

Perhaps most importantly, PV interneurons synapse perisomatically onto MSNs (Koós and Tepper, 1999). The axonal arborizations of PV interneurons extend further than their dendritic arbor (Kawaguchi et al., 1995), and synapse onto neighboring MSNs at a rate of 25% (Koós and Tepper, 1999). This is remarkable considering the infinitesimal population size of these cells relative to the MSNs. MSNs do not form synapses recurrent synapses onto PV interneurons, suggesting that these interneurons play a large role in shaping MSN activity. In addition, PV interneurons synapse onto each other and also show gap-junction coupling (Kita et

al., 1990; Russo et al., 2013). Thus striatal interneurons share many properties with interneuron populations in the cortex (Isaacson and Scanziani, 2011).

The role of feedforward inhibition in the striatum is a topic of great interest in the study of the striatal microcircuit and only recently has PV interneuron activity been studied *in vivo*. This is facilitated by the assumption that these cells can be identified in extracellular recordings by their spiking characteristics. *In vivo* these cells are referred to as fast-spiking interneurons (FSIs), as they display higher firing rates than MSNs, show a shorter-duration spike waveform and bursts in their firing activity (Berke et al., 2004; Mallet et al., 2005). Despite the ability for PV interneurons to exert feed-forward inhibition onto MSNs in slice experiments, *in vivo*, FSIs cooperate with MSN populations rather than simply inhibit them. FSIs show modulations during behavioral tasks that suggest that they operate together with MSNs and not necessarily in opposition to them (Bakhurin et al., 2016; Gage et al., 2010; K. Lee et al., 2017). Together, the MSNs and PV interneurons form the best-understood microcircuit elements in the striatum. However, their interactions in the intact brain and how they operate in the context of glutamatergic inputs to shape the striatal output signal is an area of active study.

### *The organization of excitatory inputs of the striatum*

Anatomical tracing studies of striatal afferents have been instrumental in guiding our understanding of the role the structure may play in brain function. One fundamental principal of corticostriatal inputs is that there exists a topographic organization of projections into the striatum (Flaherty and Graybiel, 1991; Hintiryan et al., 2016; Kemp and Powell, 1970; Künzle, 1977; Oh et al., 2014) and rodents (Hintiryan et al., 2016; McGeorge and Faull, 1989; Oh et al., 2014). Similar functional cortical areas (e.g. associative cortices vs. somatosensory-motor areas vs. limbic cortices) converge together into specialized subdomains of the striatum. Thus several functional pathways, or “loops”, recruit their own domains in the striatum (Alexander et al., 1986), but may potentially interact within the basal ganglia circuitry (Joel and Weiner, 1994).

The organization of corticostriatal inputs to the best studied sensorimotor loop in the putamen exhibits a complex somatotopy (Graybiel et al., 1994). First, the striatal somatotopic representations are arranged such that caudal regions of the body are oriented medially and rostral areas more laterally (Carelli and West, 1991; Cho and West, 1997; Künzle, 1975). Second, within each single body part representation domain however, injections of anterograde tracers into a single somatosensory region of the cortex result in labeling of projections in multiple patches of the putamen in the monkey (Flaherty and Graybiel, 1994; Selemon and Goldman-Rakic, 1985) and to a lesser extent in the rodent (L. L. Brown, 1992). Third, using two distinct anterograde tracers, it has been shown that the closer two representations of body parts are in the somatosensory cortex, the closer will be the patches of innervation from the two cortical areas in the striatum (Malach and Graybiel, 1986). Lastly, corresponding regions of motor and sensory cortices that homologous body part representations show partial overlap in the striatum (Hoffer and Alloway, 2001).

Cortical inputs to the striatum have large axonal arborizations that are broadcast within the striatal volume (Zheng and C. J. Wilson, 2002). The origins of these inputs can either be collaterals from descending cortical efferents in the internal capsule (Levesque et al., 1996), or are collaterals of intratelencephalic cortical projections that do not terminate in deeper structures (Lei et al., 2004). They are either focal, about 0.5 mm in size, or more distributed, being >1 mm in size (Kincaid and C. J. Wilson, 1996; Levesque and Parent, 1998). The axons contain varicosities that are distributed along the axon with an average spacing of 12.5  $\mu\text{m}$ , with each varicosity containing a single presynaptic density (Kincaid et al., 1998). Because the varicosities do not cluster and are instead widely distributed within the striatal volume, an individual cortical projection neuron makes few connections with any one MSN (Bolam et al., 2000). As a result, an individual MSN receives only a tiny fraction of its input from any one cortical projection, and suggests a large degree of convergence from many different cortical inputs onto each striatal projection neuron: over 5000 different cortical input sources per MSN (Kincaid et al., 1998).

Altogether, the morphology and physiology of the striatal MSN, the sparse receipt of individual corticostriatal inputs, and the partial overlap of functionally similar inputs suggests that the striatum integrates activity from distinct, but related input areas (L. L. Brown, 1992; Kocsis et al., 1977). This allows for the possibility that non-adjacent cortical areas that are involved in a given behavior have the opportunity to interact at the level of the striatum (Goto and Grace, 2005; Hoffer and Alloway, 2001; Reig and Silberberg, 2014). A given MSN activity profile may therefore reflect an integrated transformation of distributed cortical activity (Zheng and C. J. Wilson, 2002). In addition, because of their reticence to generate action potentials, a concerted and coordinated input onto individual MSNs is required for them to activate (Bolam et al., 2000). Within these constraints one can imagine that any coordinated activity that can excite MSNs must have functional relevance. Indeed, these findings have supported ideas that the striatum plays a fundamental role in associative forms of learning.

#### *Striatal based learning and plasticity*

Prior to definitively linking the striatum with non-declarative forms of learning, it had been known for decades that the brain used multiple memory systems as patients with hippocampal lesions were able to learn new motor skills (B. Milner et al., 1998). However, the regions that mediated these forms of learning were not determined until analogous tasks were developed for animals (Morris et al., 1982). Using radial arm-maze tasks, it was shown that rats could keep track of the spatial position of reward-locations and avoid reentering those same locations. Rats would show slower learning rates and more errors after hippocampal lesions (Packard et al., 1989). This was consistent with human patients' impaired declarative memory. In the same study, lesions of the dorsal striatum did not result in as severe an impairment in the spatial navigation task as hippocampal lesions. However, these rats showed impaired learning in tasks that used cues to guide arm-entry. These results linked the striatum with forms of associative learning in which specific stimuli were linked with specific actions. The double dissociation was later

demonstrated in humans with either amnesia or neurodegenerative disorders targeting the basal ganglia (Knowlton et al., 1996).

Decades prior to any of these discoveries, behavioral psychology had delineated a learning strategy-landscape that sought to characterize the various ways that animals learn to link environmental stimuli with actions. This framework has been influential in shaping our understanding of striatal function. Such associative behavior is typically classified as instrumental or Pavlovian. Instrumental learning enables animals to gain control of the environment by using movements to control how events occur around them. Traditionally, instrumental learning was considered to describe responses that are gradually acquired in the presence of specific stimuli after they were reinforced (for example by a sugar reward). This form of learning, labeled Stimulus-Response (S-R) implies that actions are reactions to specific environmental events. However, the logical interpretation of this organization is that any time the Stimulus is presented, animals will perform the Response. To stop performing the response requires experience with the lack of a reward such that the S-R association would weaken. However, it was demonstrated that the association can be weakened by either weakening the direct relationship between the Response and the reward presentation, or by giving overfeeding animals on the reward prior to testing for Responses (Balleine and Dickinson, 1998). In both cases, animals will generate fewer Responses even though the stimulus that should initiate S-R behavior is still present. This suggested that animals make associations between their Responses and the Outcome (the reward) itself, and identifying R-O learning.

In Pavlovian learning, animals learn to anticipate environmental events based on their antecedent events. Often this also results in concomitant behavioral responses to those antecedent events, which acquire meaning for the animal. This is a powerful form of learning and is obviously evolutionarily adaptive, but unlike with instrumental learning, its limitation is that it renders animals subject to the vagaries of their environment. Pavlovian learning is generally

also sensitive to Outcome devaluation, but as it is not dependent on any actions, the association is labeled as S-O.

The striatum because of its convergent inputs and cells that seem to integrate across many diverse inputs is a strong candidate for mediating these different forms of learning, which have been mapped on to distinct striatal subregions. Within the realm of instrumental learning, two subregions of the dorsal striatum have been associated with mediating R-O or S-R responding. The dorsomedial striatum is associated with R-O learning, as this area receives input from predominantly associative cortices, such as the frontal cortex. In addition, lesions of this subregion result in animals learning instrumental actions in an S-R manner in that they are less sensitive to outcome devaluation (Balleine et al., 2007) or action contingency degradation (Yin and Knowlton, 2004). In contrast, the dorsolateral region of the striatum has been linked to S-R forms of learning. This area of the striatum receives input from sensorimotor and infralimbic cortex. While with enough training on a specific action renders animals insensitive to outcome degradation, animals with dorsostriatal lesions never transition into this habit-like state (Yin et al., 2004). This double dissociation of striatal function suggests that its subregions perform distinct functions during the acquisition of instrumental tasks and their eventual transition into habits.

The striatum's role in Pavlovian conditioning has been typically ascribed to ventral areas of the structure, the nucleus accumbens. This region has long been associated with the linkage of motivational drives with actions, as it receives convergent input from a variety of areas that are known to be involved in eliciting movements (Mogenson et al., 1980). One of these inputs are the dopamine projections from the midbrain (Beckstead et al., 1979; Joel and Weiner, 2000), whose activity patterns reflect changes in the cue-reward relationships (Schultz et al., 1997), and fit well with the Rescorla-Wagner model describing Pavlovian learning (Fanselow and Wassum, 2016). Lesions of the ventral striatum in rats has deleterious effects on Pavlovian approach behaviors (Cardinal et al., 2002), and dopaminergic blockade in this region also

affects acquisition and expression of conditioned behaviors in Pavlovian tasks (Parkinson et al., 2002; Saunders and Robinson, 2012). In addition, the nucleus accumbens has been shown to also mediate Pavlovian components of instrumental behaviors in Pavlovian-to-Instrumental Transfer experiments (Corbit et al., 2001).

If the striatum is the site of these various forms of associative learning, it can then be assumed that it is an important site of plasticity in the brain. *In vitro*, high frequency stimulation protocols similar to what had been applied in the hippocampus (Bliss and Lomo, 1973) results in long-term synaptic depression at corticostriatal synapses that is DA dependent (Calabresi et al., 1992a; Lovinger et al., 1993). LTD has been shown to occur on both direct and indirect pathway MSNs, to be dependent on DA activity on cholinergic interneurons (Z. Wang et al., 2006), and require endocannabinoid signaling (Kreitzer and Malenka, 2007). However, another study showed that DA does indeed promote LTD in indirect pathway MSNs, but blocks LTD in direct pathway neurons, consistent with the rate model of the basal ganglia in which DA potentiates the direct pathway (Shen et al., 2008). Evidence in support of the more canonical form of long-term synaptic plasticity, long-term potentiation, also supports this form of learning as occurring in the striatum (Calabresi et al., 1992b; Kreitzer and Malenka, 2008; Shiflett and Balleine, 2011). Despite the challenge to translate the results of LTD and LTP induction protocols performed *in vitro*, several studies have demonstrated evidence for striatal plasticity *in vivo*.

More direct measures of circuit modification have demonstrated that synaptic potentiation is also a valid form of plasticity. *Ex vivo* studies that record MSN sensitivity to corticostriatal stimulation can reveal synaptic potentiation of excitatory synapses in MSNs in the dorsal striatum after animals acquire new skills, suggesting that striatal plasticity can also show forms of LTP (O'Hare et al., 2016; Shan et al., 2014; Yin et al., 2009). In addition, a variety of studies have shown evidence for alterations in striatal activity with learning using extracellular recording in animals. These studies report changes in single neuron firing across time as

learning progresses (Atallah et al., 2014; Costa et al., 2004; Santos et al., 2015), or reorganization of entire striatal populations (Barnes et al., 2005; Thorn et al., 2010).

Lastly, several studies have gone further to demonstrate selective potentiation of corticostriatal circuitry that is task-specific in behaving animals. In one series of studies, synchrony between striatal and motor cortical neurons developed as animals learned to control a BMI interface with M1 cortical neuron activity. Coherence specifically was highest between the specific cells selected to control the machine (Koralek et al., 2013; 2012). Another study showed complementary results in that only corticostriatal projections that were tuned to reward-predicting stimuli showed potentiation at striatal synapses (Xiong et al., 2015). These studies underscore the importance of learning in these circuits.

I have discussed evidence that the striatum is involved in non-declarative forms of learning. What are the consequences of learning on striatal networks? One of the fundamental principles of learning and plasticity is that networks of neurons form assemblies through experience (R. E. Brown and P. M. Milner, 2003; Yuste, 2015). This mechanism has been highly influential in explaining the brain's high level of organization. Assemblies of neurons form in small networks as observed in slices of hippocampus (Bonifazi et al., 2009) and cortex (Perin et al., 2011). At the opposite end of the spatial scale is the existence of long-range organization of functional brain networks (Bullmore and Sporns, 2009; Raichle, 2010). It is presumed that the brain's functional organization at these various levels arises because of plasticity mechanisms.

I have also introduced the wide-ranging, convergent, topographically organized excitatory input to the striatum, it being a massive release site of dopamine, and evidence for plasticity at the corticostriatal synapses. Thus I anticipated that neural assemblies in the striatum might show functional organization as a result of plasticity at corticostriatal synapses. The work I discuss in chapter 1 investigates several functional relationships within the striatal microcircuitry in awake, behaving animals (Bakhurin et al., 2016). Primarily, I show that populations of neurons in the sensorimotor striatum whose activity tracked behavior in a

Pavlovian odor association task show higher likelihoods of statistical association while animals were at rest outside of the task. In addition, I demonstrate that FSI populations closely mimic MSN responses during the task, and preferentially form associations with task-modulated MSNs.

## **Problem 2: The role of striatal population dynamics in timing predictive behavior and their origins**

The brain is fundamentally a prediction machine that receives sensory information and attempts to anticipate future events in the world (Clark, 2015). How it performs this function is an open question in neuroscience. In this section, I will describe a theoretical framework on how nervous systems naturally predict future events, with an emphasis on the neural encoding of time. Within the last decade, it has become clear that many brain areas, including the striatum, in behaving animals generate complex activity patterns during behavior. In chapter 2, I describe my research into the encoding of time in such complex network activity found in the striatum of behaving animals. Although we can readily observe and describe these dynamics, how they arise from the striatal microcircuitry is not well understood. Their patterns are thought to arise from a complex interaction of the excitatory inputs that the striatum receives and local recurrent inhibitory network activity. In chapter 3, I present new work that explores this balance of inputs and local microcircuit activity in the generation of complex striatal output signals.

### *Time as an emergent property of dynamic neural systems*

Nervous systems have a remarkable degree of flexibility and vast learning capacity. Learning and memory functions of the brain are thought to arise from the activity of large populations of neurons, but how activity across many cells leads to actual neural computation is not well understood. As discussed earlier, neurons can use plasticity mechanisms to associate together into functional assemblies (Fregnac, 2003). These assemblies appear to have functional

properties that form via their interactions, such as oscillatory activity in central pattern generators (Prinz et al., 2004). By assembling together into functional units, neurons can achieve new operations that are impossible for single neurons to perform (Hopfield, 1982; Yuste, 2015). While biology has had to approach understanding these assemblies in a bottom-up manner, computational neuroscience has had many insights using models of groups of neurons as to how brains may learn and use memories to make predictions.

Neural networks are often modeled in two general architectures. One class is called a feed-forward network (Rosenblatt, 1958). In these systems, information flows in a unidirectional manner, often between multiple layers. By adjusting the strengths of connections between layers, these systems are well suited to performing sophisticated classification and discrimination tasks. However, to approximate neural computation, a second class of network design better mimics canonical cortical networks that contain mostly excitatory neurons and some inhibitory interneurons that are interconnected using specific rules (Hopfield, 1982). Recent developments in this class of recurrently connected networks have been able to show properties similar to those observed in vivo, namely activity patterns that are continuously changing (Maass et al., 2002). It is the recurrent network that provides us with a compelling framework within which we can understand how the brain may tell time at subsecond timescales.

Time plays a fundamental role in nearly every aspect of natural neural function. For example, time is an essential dimension of speech in that the brain extracts meaning from the temporal relationships between sounds. Skilled movements, such as tying one's shoes or playing a musical instrument, require the brain to keep track of specific actions and keep them in precise temporal register with each other. In all three examples, the time scales of timekeeping are generally shorter than one second. The neural mechanisms that mediate this time scale are poorly understood.

If we consider the brain as a dynamic system that operates using varying activity that is distributed across neural assemblies, it can be proposed that time is an inherent property of neural function itself (Goel and Buonomano, 2014). Such a model suggests that the brain can learn to keep track of time by monitoring the evolution of its own activity. In order to appreciate this perspective, one can consider that any continuously changing system can be used to tell time by keeping track of the evolution of the change in its configuration. A common example of such a system is the patterns of ripples that are caused by dropping pebbles in a still pool of water at specific time intervals. By taking serial photographs of the pool as the ripples evolve, one can make an estimate as to how much time has passed from the configuration of ripple patterns in each photograph. Replacing the pool of water with a network of neurons, the pebbles with stimuli, and the camera with populations of cells that learn to listen to these dynamics, one can begin to understand how time can naturally emerge from network activity.

Two important considerations on neural networks are required in order to understand how they can be good time keeping mechanisms. First, such systems can keep track of serial order between events because of their dynamic properties. Short-term plasticity in these models, such as paired-pulse facilitation, synaptic depression, or long vs. short acting GABAergic signaling can shift or alter the underlying subthreshold membrane activity of a network of neurons. This can create so called 'hidden states' (Stokes, 2015), making the network sensitive to temporal order of events occurring at the sub-second timescale (Buonomano, 2000; Buonomano and Merzenich, 1995; Mauk and Buonomano, 2004). Presenting identical stimuli with different inter-stimulus intervals will therefore generate unique activity patterns across the population of the same network, activity at a given moment in time dependent on the state of the network in the prior moment (Buonomano, 2000; Maass et al., 2002). Thus stimulus A or stimulus B will result in dynamics that are distinct from stimuli presented in quick succession (AB or BA).

Second, dynamic systems contain naturally occurring noise or variability in their activity. This is thought to underlie the variability seen in recordings of real neural systems (D. Lee et al., 1998). Neural networks also generate random activity because of their intrinsic sources of excitation, and thus their state-dependent dynamics are sensitive to this variability to the degree that the same stimulus will rarely replicate the same activity pattern in these neural networks (Buonomano and Maass, 2009; Sompolinsky et al., 1988). This is because moment-to-moment configurations of the network are always dependent on prior states. Such chaotic activity will be challenging to learn for a read-out system. However, by adjusting the strengths of interconnectivity within the recurrent network, these networks can become a reliable source of time (Goel and Buonomano, 2014; Laje and Buonomano, 2013). The network can essentially learn to reproduce activity patterns, and can thus read-out mechanisms can extract meaningful information from a given state. Like the pool of water, the brain can be used to tell time.

Three essential questions remain within the state-dependent models of timing. 1) How are biologically realistic learning rules applied to the network in order for it to become a reliable clock? 2) Are such state-dependent networks implemented in the brain in order to keep track of time? 3) If they are found in the brain, are they employed in some brain regions more than others? The first question will likely require further work with network models (Sussillo and Abbott, 2009). Several studies using *in vivo* recordings provide support for the second question: Dynamic network trajectories have been reported in a wide variety of different brain areas, including the hippocampus (M. A. Wilson and McNaughton, 1993), parietal cortex (Harvey et al., 2012), motor cortices (Churchland et al., 2012), and prefrontal cortex (Fujisawa et al., 2008; Pinto and Dan, 2015). The third question is addressed in chapter 2.

Recently, the striatum has been shown to also generate these kinds of dynamic activity patterns. Dynamic striatal network-level activity has been reported *in vitro* after applying tonic NMDA to striatal slices (Carrillo-Reid et al., 2008). Furthermore, dynamic activity patterns have been reported in rodents during serial-process decision making tasks (Gage et al., 2010), during

navigation toward rewards in a maze (Thorn and Graybiel, 2014). While one could speculate with regard to the computational roles of these dynamics, they have been shown to be involved in tasks requiring temporal processing (Gouvêa et al., 2015; Mello et al., 2015; Rueda-Orozco and Robbe, 2015).

### *Striatal dynamics and time*

The basal ganglia have received extensive attention with respect to the brain's ability to tell time. This focus originated from understanding that the striatum is involved in learning and expressing sequential actions (Graybiel, 2008). Furthermore, experiments showing that pharmacological manipulations of DA could bidirectionally change animals' performance in a peak-interval procedure (Meck, 1983) and some reports of altered temporal processing in patients with Parkinson's disease (Allman and Meck, 2012). In addition, human imaging experiments reveal activation in the striatum under a variety of different timing tasks (Coull et al., 2010).

Most models for timing at sub-second timescales incorporate the basal ganglia. Some of the first models for timing mechanisms in the brain were based on a centralized pacemaker and counter mechanism. In this model, upon perceiving a trigger stimulus that instructs the beginning of an interval, a switch initiates a counter that reads out repeating signals from the pacemaker. A comparator mechanism would compare the number of pulses to a memorized store of previously experienced pulse numbers, thereby estimating the passage of time (Gibbon, 1977). Dopaminergic signaling (Meck, 1983) and basal ganglia oscillators (Plenz and Kital, 1999) known to have pacemaker-like tonic firing patterns were thought to serve as the counter. However, efforts to identify neural correlates of the accumulator component of the model have not been as conclusive (Buhusi and Meck, 2005). An alternative proposal relies more on the dependency of striatal activity on its convergent, excitatory inputs. The striatal beat-frequency model suggests that multiple oscillators operating at different frequencies in the brain,

particularly within corticostriatal interactions can be monitored and used as a time-keeping mechanism. At specific intervals, different oscillators will show unique overlap in phase relationships, and plasticity mechanisms can be used to train readouts to identify those coincident signatures of distinct intervals (Matell and Meck, 2004).

The previous proposals are based on a general hypothesis that specific circuitry evolved in the brain in order to explicitly track time. However, the state-dependent models described earlier provide an elegant solution for subsecond timing in the brain that does not require that new circuitry evolved specifically for that purpose. This framework also suggests how the same neural mechanisms can be recruited for different tasks that are performed under the same temporal constraints. The striatum has been shown to employ these kinds of population dynamics to explicitly encode the passage of time (Gouvêa et al., 2015; Mello et al., 2015). While these studies identify population clock-like dynamics in the striatum, these findings return to the potential scenario that specific brain areas are specialized for encoding time. While in Chapter 1 I show additional evidence for population trajectories in striatal networks, I pursue their relationships to timing of behavior in chapter 2. There I use machine-learning approaches to demonstrate that multiple brain areas are capable of encoding time at the single-trial level and that this function seems to be distributed across multiple brain areas. The striatum and a cortical input area, the orbitofrontal cortex, were differentiated in their reliability and capacity to tell time. To support the argument that striatal networks are related to timing, I demonstrate that its population codes scale with anticipatory movement onset time. As this sensitivity to movement timing was found only in the striatum, I propose that that this brain region may have a privileged role in representing time.

### *Origins of striatal dynamics*

How striatal dynamics arise requires some consideration, particularly since the striatum is an inhibitory network with sparse recurrent feedback, and thus cannot follow the same principles as

the recurrent networks based on cortical connectivity (Tepper et al., 2008). Early network models of the striatum involved creating mutually inhibitory, competing populations that operated in a “Winner-take-all” fashion (Beiser and Houk, 1998; Connolly and Burns, 1993; Fukai and Tanaka, 1997; Wickens et al., 1995). In these frameworks, convergent input onto a specific population of MSNs will lead to a large enough depolarization leading to action potential generation. These MSNs thus inhibit other populations that do not receive the same convergent input or do not reach firing threshold. These mechanisms have been suggested to support the striatum’s role in action selection by arbitrating between signals from the cortex (Mink, 1996; Redgrave et al., 1999). It is conceivable that such a mechanism, if action switching occurred quickly enough, could result in apparent activity dynamics.

The assumptions made by these early models with respect to strong, mutual inhibition among MSNs are not supported by physiological data. Feedback inhibition is currently thought to be weakly influential in generating spiking activity in neighboring neurons, and is largely unidirectional (Taverna et al., 2004; Tepper et al., 2008; Tunstall et al., 2002). More contemporary models of the striatal network have incorporated more accurate MSN interconnectivity with one another and striatal interneurons and incorporate the possibility that there may be modular channels or preferred paths through which feedback inhibition operates (Plenz, 2003). Lastly, correctly employing correct feedback connectivity in network models of the striatum allow striatal networks to form clusters of synchronized neurons that continuously alternate their relative firing rates (Humphries et al., 2009; Ponzi and Wickens, 2010). These models thus provide evidence that inhibitory striatal networks with weak recurrent connections can still generate complex dynamic activity patterns.

While contemporary models of the striatum have made attempts to simulate naturalistic statistics of excitatory input into the striatum, we still have very little understanding as to the nature of these inputs in the intact brain. Furthermore, striatal models have not yet incorporated the topographic organization of striatal inputs, which should provide unique spatiotemporal

inputs in a task-dependent manner. In chapter 3 I provide preliminary data that investigates the relative contributions of glutamatergic input and local striatal microcircuitry to the generation of striatal dynamic activity. Using optogenetic inhibition of a task-relevant glutamatergic signal from cortex to a recipient area of the striatum, I show that striatal output is a complex interaction of local microcircuit activity and excitatory input. This study is among the first to causally test the role of glutamatergic signaling on striatal activity with high temporal resolution.

## **Chapter 1: Temporal correlations among functionally specialized striatal neural ensembles in reward conditioned mice**

Animals must rapidly learn to discriminate environmental cues associated with beneficial outcomes from irrelevant cues. In the vertebrate nervous system, the basal ganglia are a set of interconnected nuclei whose activity has been extensively linked to reward-guided learning and action selection (Graybiel, 2000; Kravitz et al., 2012). Within these circuits, the striatum serves as the primary input structure to the basal ganglia as well as an important site of synaptic plasticity (Gerfen and Surmeier, 2011; Kreitzer and Malenka, 2008). Previous efforts to understand the role of striatal electrophysiological activity in guiding behavior revealed that this region is strongly modulated by primary rewards, stimuli that predict rewards, and action (Nicola et al., 2004; Roitman et al., 2005; Setlow et al., 2003; L. Tremblay et al., 1998). Furthermore, a number of studies have reported changes in striatal activity that accompany learning (Costa et al., 2004; Jog et al., 1999; Koralek et al., 2013; Tang et al., 2009; Thorn and Graybiel, 2014; L. Tremblay et al., 1998; Xiong et al., 2015; Yin et al., 2009). These effects are consistent with a neural circuit that can become tuned to select specific behavioral responses in anticipation of an appetitive outcome. Yet despite a substantial amount of work on neural dynamics in this area, relatively little is known about how striatal neurons are functionally organized at the network level.

In contrast to the organization of cortical microcircuits which contain strong local excitation (Ko et al., 2014), the striatum is a largely inhibitory structure consisting of GABAergic medium spiny projection neurons (MSNs) coupled with a small population of interneurons (Kreitzer and Berke, 2011; Tepper et al., 2008; Tepper and Bolam, 2004). At millisecond timescales, striatal activity is likely to be strongly influenced by shared glutamatergic signaling from cortical, thalamic, and limbic inputs (Cowan and C. J. Wilson, 1994; Kasanetz et al., 2006; Kincaid et al., 1998; Plenz and Kitai, 1998; Reig and Silberberg, 2014; Stern et al., 1998). It is believed that the striatum integrates these convergent streams of information, with the resulting

activity acting on basal ganglia output nuclei connected to association and motor control areas in the cortex (Alexander et al., 1986). Computational and experimental studies suggest that the converging input to the striatum leads to the formation of functionally specialized subsets of MSNs with temporally correlated activity patterns (Adler et al., 2012; Carrillo-Reid et al., 2008; Humphries et al., 2009; Ponzi and Wickens, 2010; Yim et al., 2011). These findings in the striatum, and a large body of work focusing on cortical circuits (Averbeck and D. Lee, 2006; Bair et al., 2001; Cohen and Maunsell, 2009; Mitchell et al., 2009; Shadlen and Newsome, 1998; Zohary et al., 1994), implicate correlated activity in neural computation and behavior. This evidence raises the possibility that temporal correlations preferentially occur among task-related groups of neurons in the striatum.

To study this relationship, we used multielectrode probe recordings to simultaneously monitor activity from over 100 units in head-fixed mice undergoing Pavlovian reward conditioning. The large scale of these measurements enabled systematic analysis of correlations among behaviorally and electrophysiologically identified groups of cells. The results of this study reveal a potential organizational principle for the dynamics of striatal neurons that encode similar features during the behavioral task.

## **Materials and Methods**

### *Animals and surgical procedures*

All procedures were approved by the University of California, Los Angeles Chancellor's Animal Research Committee. Singly housed male C57Bl/6J mice (n=9, 12-16 weeks old, The Jackson Laboratory) were used in the experiments. Animals underwent an initial surgery under isoflurane anesthesia in a stereotaxic apparatus to bilaterally fix stainless steel head restraint bars (10 mm x 7.5 mm, 0.6 g) on the skull. Animals were anesthetized with isoflurane for a second surgery on the recording session day to make a craniotomy for acute microprobe insertion. Rectangular

craniotomies (0.5 mm AP x 2 mm ML) were centered on the following striatal coordinates relative to bregma: AP, 1.25-1.3 mm, ML, 0.95 mm. An additional craniotomy was made over the posterior cerebellum for placement of an electrical reference wire.

### *Behavioral task*

After recovery from the first surgery, animals were food restricted and fed daily after each training session to maintain ~90% of their baseline weight. They received water *ad libitum*. During daily training sessions, animals were mounted on the head bar bracket on the recording rig and stood on a polystyrene spherical treadmill (200 mm diameter, Graham Sweet Studios) that rotated along a single axis during forward/backward ambulation. The treadmill velocity was monitored with an optical mouse. Delivery of the reward solution (5  $\mu$ L, 10% sweetened condensed milk) was from a tube positioned between an infrared lick meter (Island Motion), and was controlled by an audible solenoid valve actuation (Neptune Research). We studied the behavioral and electrophysiological profiles of previously inexperienced mice trained with odors for the first time. Before conditioning, animals were habituated to head fixation by receiving rewards alone (maximum 100 rewards per daily session, 13-21 s inter-trial interval, ITI), and exposed to a constant flow of odorless air (1.5 L/min) through a tube. After animals successfully consumed 90% of delivered rewards for two consecutive days, they underwent surgery for recording and began conditioning with olfactory cues using an olfactometer. Odorants were introduced by bubbling air (0.15 L/min) through aromatic liquids diluted 1:10 in mineral oil (Sigma-Aldrich), and mixing this product with the 1.5 L/min stream of air. The task involved two stimulus conditions consisting of either a 1 s olfactory cue (CS+) followed by a temporal delay of 1.5 s and subsequent delivery of a reward solution, or a different 1 s olfactory cue (CS-) that was not followed by reward (**Fig. 1.1A**). Odors were presented in pseudorandom order (1 s duration, 17-29 s ITI). The CS+ consisted of amyl acetate and the CS- consisted of citral. Correct CS+ hit trials were defined as those containing anticipatory licking activity detected

between  $t=0-2.5$  s from the cue onset, *i.e.*, prior to reward delivery. Correct CS- withholding trials were defined as the absence of any licking activity between  $t=0-5$  s from the cue onset. Mice underwent recording on the first day of odor conditioning. During the recording animals received 100 CS+ trials paired with reward delivered at 2.5 s after cue onset and 100 CS- trials with no reward.

### *Electrophysiological recordings*

Silicon microprobes (Shobe et al., 2015) were fabricated in a silicon microelectromechanical systems foundry (Innovative Micro Technology). Each silicon microprobe contained a total of 256 electrodes ( $10\ \mu\text{m} \times 10\ \mu\text{m}$  electrode dimensions,  $\sim 30\ \mu\text{m}$  electrode spacing) distributed on multiple silicon prongs. Recordings were performed using one of two device designs. The prongs of each device type were arranged to provide high-density electrophysiological measurements across a large area of the striatum. Type I probes had 4 prongs positioned at fixed depths, with 64 electrodes per prong distributed along 1 mm at the tapered tip of each shaft. Type II probes had 5 prongs distributed at different depths, with 50 or 52 electrodes per prong. Recording and spike sorting procedures are described in (Shobe et al., 2015).

### *Striatal unit classification*

All analysis was carried out with custom Matlab scripts. We used spike waveform trough-to-peak ( $t_{\text{tr-pk}}$ ) duration and coefficient of variation (CV) of baseline firing rate to classify units into putative MSNs, fast spiking interneurons (FSIs), and tonically active interneurons (TANs) (Aosaki et al., 1994; Bennett and C. J. Wilson, 1999; Gage et al., 2010; Mallet et al., 2005). FSIs were separated from non-FSIs by their narrow waveform (maximum FSI  $t_{\text{tr-pk}}=0.475$  ms, minimum non-FSI  $t_{\text{tr-pk}} =0.55$  ms, and maximum non-FSI  $t_{\text{tr-pk}} =1.25$  ms). TANs were separated from other non-FSI units by the regularity of their baseline firing (maximum TAN CV=1.5), leaving units that exceeded this CV as putative MSNs. Around 9% of all units were not classified

in any of these three categories and were excluded from further analysis.

### *Histology*

After recording, animals were overdosed with sodium pentobarbital and perfused with 10% formalin solution. Brains were extracted and fixed overnight at 4°C. To confirm the correct targeting of the microprobe, tissue was stained for tyrosine hydroxylase (TH) using sheep anti-TH primary (Millipore, 1:500) and TRITC-conjugated donkey anti-sheep secondary antibodies (Jackson ImmunoResearch, 1:100). Microprobe tracks were determined by locations of DiD fluorescence in images of TH-stained sections. We confirmed that recordings were located at approximately the same coronal section of the striatum (range of AP positions relative to bregma: 1.1 to 1.4 mm). We could subsequently determine the approximate mediolateral and dorsoventral silicon prong positions, and thus cell position.

### *Discriminating cell identification*

Cue-triggered firing rate,  $R(t)$ , was calculated from the average firing rate of CS+ trials with anticipatory responding (correct hits) and CS- trials without licking (correct withholding). The time bin size was 50 ms. The baseline period was defined as the 5 s interval preceding cue presentation, and the average firing in this period was used to calculate the baseline-subtracted change in firing,  $\Delta R(t)$ . The mean baseline subtracted and normalized firing rate was obtained with the expression  $\Delta R(t)/\Delta R_{max}$ . Discriminating units were determined by comparing the distributions of firing rate during correct CS+ hit,  $R_{CS+}(t)$ , and correct CS- withholding,  $R_{CS-}(t)$ , trials using a permutation test on individual time bins (10,000 iterations). For each time bin, we shuffled the labels of firing activity for each trial to create two distributions of firing rates,  $R_{CS+,shuffled}(t)$  and  $R_{CS-,shuffled}(t)$  shuffled, that could be expected by chance for each trial type. We defined a unit as being discriminating if the absolute value of the difference between  $R_{CS+}(t)$  and

$R_{CS}(t)$  for at least two consecutive time bins was higher than the 99<sup>th</sup> percentile of the distribution of differences between  $R_{CS+,shuffled}(t)$  and  $R_{CS-,shuffled}(t)$ .

### *Behavior correlations*

Pearson correlations between spiking activity of individual units and lick rate or treadmill velocity were calculated over the entire recording, in time bins of 50 ms.

### *Signal and resting state correlations*

For signal correlations, we calculated the Pearson correlation coefficient on the  $R_{CS+}(t)$  signal of simultaneously recorded units, using 50 ms time bins from  $t=0$  to 2.5 s after cue onset on correct CS+ hit trials. We defined resting periods as intervals of at least 2 s during which animals did not make any detected movements (running, licking) and were not presented with any external stimuli (cues and reward). To find the resting state spike count correlations, we serially concatenated spiking activity occurring within these epochs to create a continuous time series vector (500 s, 10 ms bins) representing the resting state firing rate,  $R_{rest}(t)$ . We calculated resting state activity for each individual unit, and then obtained the Pearson correlation coefficient using these vectors. To detect significant correlations, we used the permutation test for correlations on each unit pair containing spiking activity in the resting state. This test involved shuffling  $R_{rest}(t)$  for one unit in the pair and recalculating the correlation coefficient (1000 iterations). This resulted in a distribution of possible correlation coefficients that could be expected by chance. We determined significance if the absolute value of the observed correlation exceeded the 99 percent confidence interval of the absolute value of shuffled values. The correlation probability is the fraction of significant pairs detected out of all possible simultaneously recorded pairs separated by 0.025 to 1 mm. We excluded pairs of units closer than 0.025 mm from all temporal correlation analysis, to minimize any effect of possible spike sorting errors.

### *Spike time cross correlations*

We calculated the spike time cross correlogram (CCG) between pairs of units using 1 ms time bins and a lag of  $\pm 25$  ms. We used spike trains from the entire recording session for this analysis. To identify significant cross correlations we recalculated the CCG after adding a random 0-5 ms jitter to the spike train (500 iterations), and determining the confidence interval from this distribution (Fujisawa et al., 2008).

### *Statistical tests*

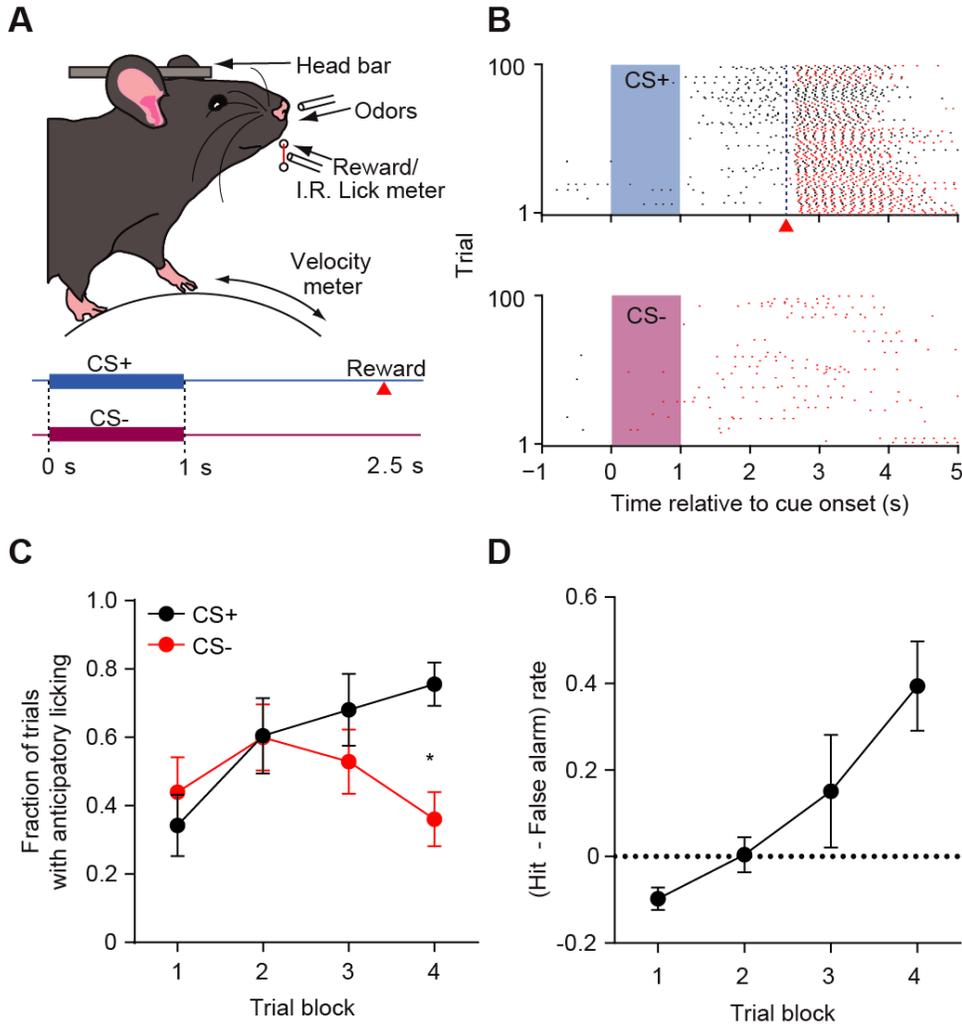
We performed nonparametric permutation tests to determine the significance of linear correlation analysis (Shobe et al., 2015). We used 1,000 to 10,000 shuffles per test. Paired t-test analysis was performed using standard Matlab functions. ANOVAs were performed using GraphPad Prism software.

## **Results**

### *Mice learn a stimulus discrimination task*

To combine large-scale neural recordings with a behavioral assay of stimulus discrimination, we implemented a Pavlovian reward-based odor discrimination task in head-restrained mice (Shobe et al., 2015). Animals typically learned the association between odor presentation and reward delivery within one recording session (**Fig. 1.1B**). To examine how responding to the two trial types changed over time we divided the training session into blocks of 25 trials (**Fig. 1.1C**). On average, hit rate increased relative to false alarm rate during training. A two-way repeated-measures ANOVA revealed a significant interaction between time and trial type ( $p=0.002$ ,  $F_{3,48}=5.94$ ). Moreover, mice were more likely to respond to CS+ trials than CS- trials in the last block ( $p<0.05$ , Sidak's test for multiple comparisons). The difference between hit rate and false

**Figure 1.1**



**Figure 1.1.** Head restrained mice demonstrate single-session discrimination learning. **A:** Experimental setup and trial schematic. Head-fixed mice were placed on a spherical treadmill and were presented with olfactory cues and liquid rewards while licking activity and treadmill velocity were monitored. Trials consisted of either 1 s of odor (CS+) followed by a 1.5 s pause and a reward, or a different 1 s odor (CS-) followed by no outcome. **B:** Licking activity rasters during CS+ and CS- trials for one representative animal. Shaded rectangles represent the olfactory cue presentation period. Black and red tick marks indicate individual licks during trials with correct and incorrect responses, respectively. Red triangles indicate the time of reward delivery. **C:** Learning curves for all animals ( $n=9$ ) showing the mean probability of licking after CS+ (black) and CS- (red) trials in blocks of 25 trials. A two-way ANOVA, repeated measures revealed a significant effect of trial block ( $p=0.003$ ) and a significant interaction between trial types ( $p=0.0016$ ,  $*p<0.05$ , Sidak's test for multiple comparisons). **D:** Evolution of the mean discriminatory behavior rate in blocks of 25 trials ( $p=0.0021$ , one-way, repeated measures ANOVA). Error bars represent SEM.

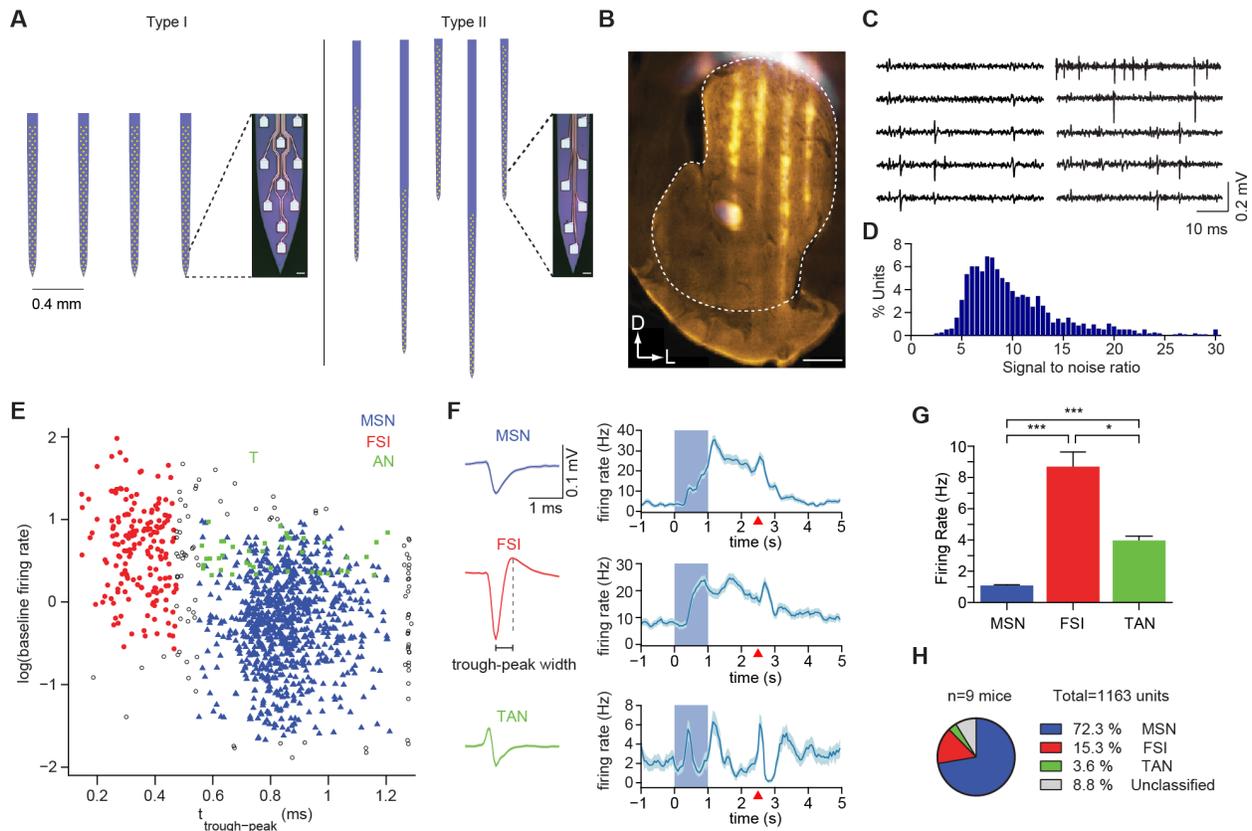
alarm rate was used to quantify discriminatory performance on the task. This measure steadily increased across blocked trials ( $p=0.005$ ,  $F_{3,24}=8.4$ , one-way, repeated measures ANOVA; **Fig. 1.1D**). Together, these results indicate that on average mice were able to develop selective anticipatory responding to CS+ and withholding of licking to CS- in a single training session.

### *Large-scale recordings reveal a population of discriminating MSNs*

To examine the dynamics of large populations of striatal neurons during the stimulus discrimination task, we employed silicon microprobes with 256 recording sites. We used two customized designs that allowed recordings from either dorsal or ventral (nucleus accumbens) striatal subregions, or simultaneously from both. In all cases the recording sites spanned a large extent of the striatum along the mediolateral axis (**Figs. 1.2A, 1.2B**). Recordings captured spiking activity distributed across the microprobe (**Fig. 1.2C**), with each session yielding an average of 115 (range: 51-188) simultaneously measured units. The spike waveforms had a median signal-to-noise ratio of 9 (**Fig. 1.2D**). We used spike waveform width and the CV of baseline firing rate to classify units as putative MSNs, FSIs, and TANs (**Figs. 1.2E, 1.2F**). We found that 52% of TANs showed burst/pause responses to rewards (**Fig. 1.2F**, bottom), which are characteristic firing properties of this cell type (Aosaki et al., 1994). MSNs and TANs displayed wider spike waveforms than FSIs, and MSNs showed lower firing rates (mean $\pm$ SD: 1.1 $\pm$ 1.4 Hz) than both FSIs (mean $\pm$ SD: 8.7 $\pm$ 12.4 Hz) and TANs (mean $\pm$ SD: 4.0 $\pm$ 1.8 Hz,  $p<0.05$ , unpaired t-test, Bonferroni corrected; **Fig. 1.2G**). MSNs and FSIs represented the highest proportion of recorded units (**Fig. 1.2H**), and we therefore focused the majority of our analysis on these two subpopulations.

When visualized at the population level, striatal MSNs exhibited distinct activity patterns during correct CS+ hit and correct CS- withholding trials (pooled data from 9 animals, n=841 MSNs; **Fig. 1.3A**). Since the network appeared to distinguish between the two stimulus conditions, we quantified the fractions of MSNs whose firing rate significantly differed between

**Figure 1.2**

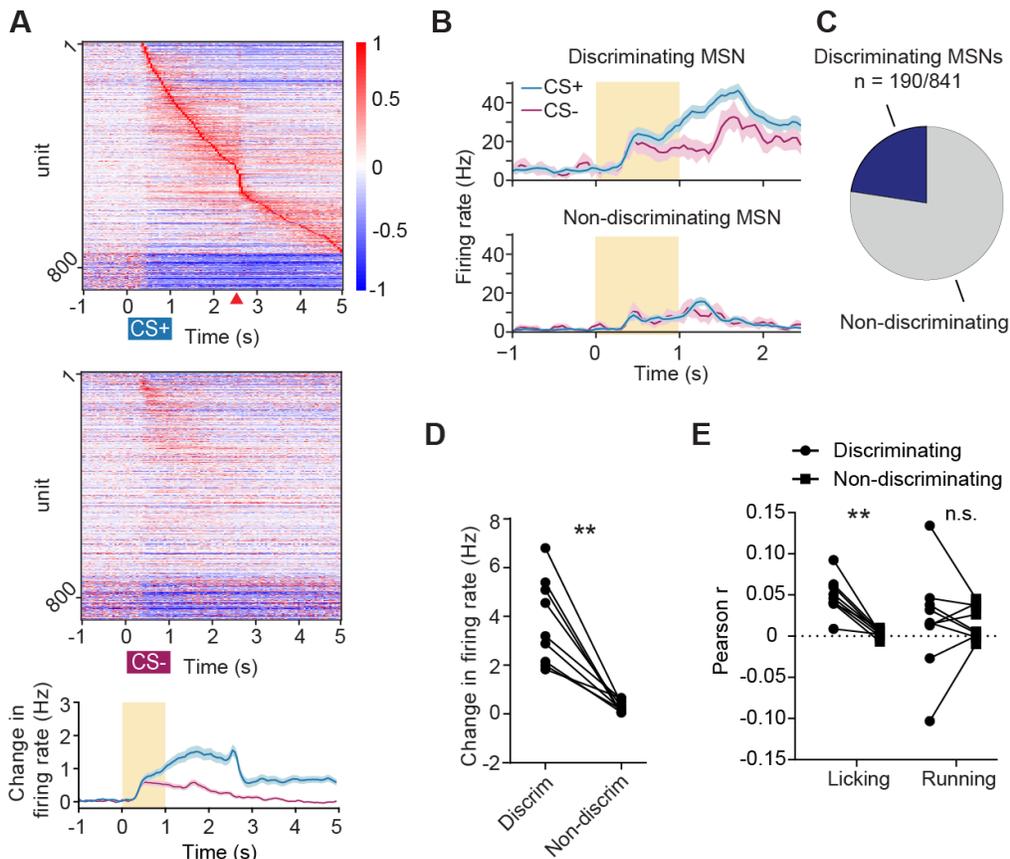


**Figure 1.2.** Large-scale striatal recordings with silicon microprobes. **A:** Illustration of the two 256 electrode silicon microprobe designs used to record in the striatum. Each silicon prong contains a high-density electrode array, with the geometry shown in magnified images of the tips. Short scale bars represent 10  $\mu\text{m}$ . **B:** Fluorescence image of silicon microprobe tracks (white) embedded in a TH-labeled section of the striatum (orange). White outline represents the perimeter of the striatum. Scale bar represents 0.5 mm. **C:** Samples of measured signals from 10 representative recording sites, filtered offline from 600 to 6500 Hz. Columns show simultaneously recorded data from two sets of adjacent recording sites. **D:** Representative waveforms of three putative cell types, medium spiny neurons (MSNs), fast spiking interneurons (FSIs) and tonically active neurons (TANs) identified in this study. **E:** Scatter plot of baseline firing rate versus spike waveform trough-to-peak time with color representing putative cell identity. Gray circles denote unclassified units. **F:** (*left*) Representative waveforms of three putative cell types, medium spiny neurons (MSNs), fast spiking interneurons (FSIs) and tonically active neurons (TANs) identified in this study. (*right*) Corresponding mean firing rate during CS+ trials for each representative unit depicted at left aligned to the cue onset. Shaded rectangle represents odor cue delivery time. Red triangles indicate reward delivery. Rates are averaged over all correctly performed trials. **G:** Mean baseline firing rate across all recorded MSNs, FSIs, and TANs in the study ( $p < 0.05$ , Bonferroni-corrected t-test). **H:** Percentages of each cell class that composed the combined dataset.

correct CS+ hit trials and correct CS- withholding trials between the cue onset ( $t=0$  s) and reward delivery time ( $t=2.5$  s; **Fig. 1.3B**). These units are referred to as discriminating cells. On average, 23% of MSNs were found to have firing rates that discriminated between these trial types (**Fig. 1.3C**). All discriminating MSNs showed higher firing to CS+ relative to CS- cues. In addition to showing different firing rates to the two cues, discriminating MSNs showed overall higher firing rate responses during correct CS+ hit trials as compared to their non-discriminating MSN counterparts ( $p=0.0004$ , paired t-test; **Fig. 1.3D**). This indicates that non-discriminating MSNs tend to be a less active population during CS+ trials. We next inquired whether these populations differentially respond to movements that were concurrently measured in the experiment (licking and running). We correlated spiking activity of discriminating and non-discriminating MSNs with lick rate and treadmill velocity. We found that discriminating MSN firing rate was more correlated to lick rate than non-discriminating MSNs ( $p=0.0002$ , paired t-test; **Fig. 1.3E**). On the other hand, discriminating and non-discriminating MSNs were equally correlated to treadmill velocity ( $p=0.9$ , paired t-test). These observations show that discriminating MSNs were more selective for licking, and were equally selective for running speed relative to non-discriminating MSNs. Thus, it appears that some discriminating MSNs are modulated by both licking and running, which is consistent with a study showing that MSNs can multiplex multiple aspects of behavior (Rueda-Orozco and Robbe, 2015). Furthermore, our definition of discriminating cells does not preclude that these neurons encode other behaviors, including non-motor aspects of the task.

Taking advantage of the high throughput recording capabilities of silicon microprobes, we mapped neural activity across a large extent of one section of the anterior striatum. Using their estimated position, neurons were assigned to one of twelve subregions forming a 4 x 3 compartment grid (**Fig. 1.4A**). The activity of the pooled population in each compartment was then averaged. Mean cue-triggered firing was found to be heterogeneous across the different grid compartments (**Fig. 1.4B**). Most strikingly, the responses to CS+ and CS- trials were most

**Figure 1.3**



**Figure 1.3.** Identification of a cue-discriminating subpopulation of striatal MSNs. **A:** Mean baseline subtracted and normalized firing rates for 841 MSNs obtained from 9 animals during correct CS+ trials (top) and correctly withheld CS- trials (center). Units in both plots are sorted by latency to peak firing during CS+ trials (top plot). Cues are presented between 0 and 1 s, indicated by colored rectangles. Reward delivery during CS+ trials is indicated with the red triangle. Bottom panel shows the mean baseline subtracted firing rate for all neurons depicted in the heat plots. The orange rectangle represents the odor delivery time. **B:** Mean baseline subtracted firing rate for two MSNs during correct CS+ trials (blue) and correctly withheld CS- trials (magenta). Odor delivery time is depicted with the orange rectangle. Top: a representative discriminating MSN defined by differential firing between the CS+ and CS- trial conditions. Bottom: a representative non-discriminating MSN. Discrimination was determined on the interval from 0 to 2.5 s following cue onset. **C:** Pie chart showing the mean fraction of cue-discriminating MSNs. **D:** Mean value per animal of baseline subtracted firing rate between 0 and 2.5 s during correct CS+ trials for discriminating and non-discriminating MSNs ( $p=0.0004$ , paired t-test). **E:** Mean Pearson correlation coefficient between spiking activity of discriminating and non-discriminating MSNs, and lick rate ( $p=0.0002$ , paired t-test) or treadmill velocity ( $p=0.9$ , paired t-test). Each point represents one animal. All error bars represent SEM.

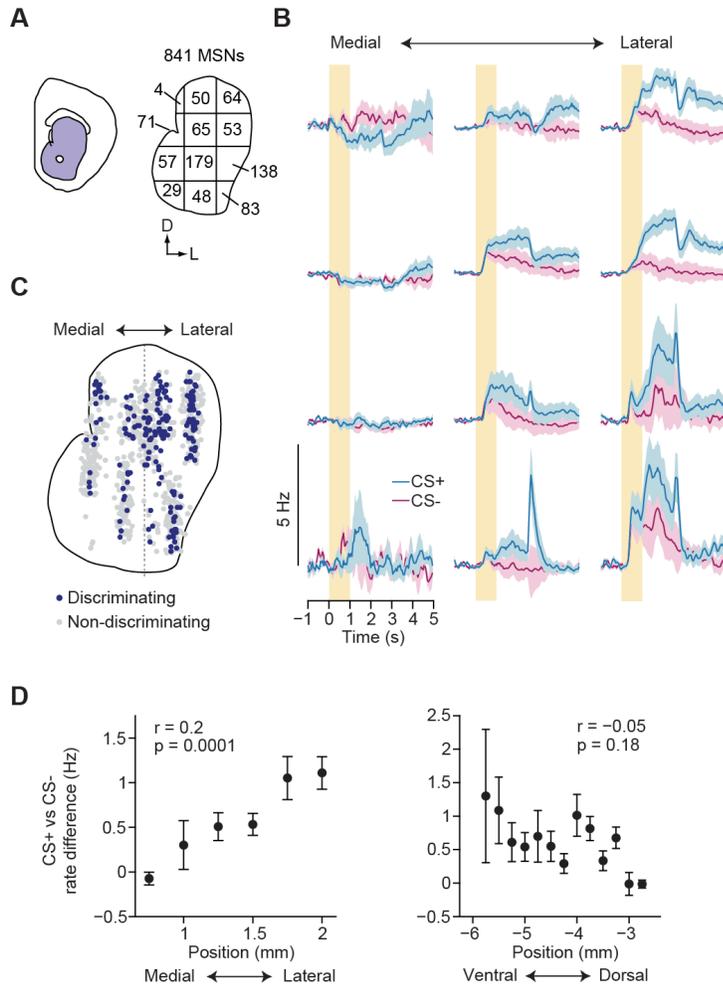
segregated in the lateral portions of the striatum. We also found that discriminating MSNs were more likely to be found on laterally positioned electrodes (discriminating MSN mediolateral position: 1.74 mm, 1 SD=0.15 mm, non-discriminating MSN mediolateral position: 1.5 mm, 1 SD=0.15 mm,  $p=0.02$ ,  $n=9$ , paired t-test; **Fig. 1.4C**). In agreement with these two observations, the mediolateral position of individual MSNs was correlated with greater differences between correct CS+ and CS- firing ( $r=0.2$ ,  $p<0.0001$ , permutation test for correlations; **Fig. 1.4D, left**). We did not detect a corresponding correlation between CS+ and CS- firing and dorsoventral unit position ( $r=-0.05$ ,  $p=0.18$ , permutation test for correlations; **Fig. 1.4D, right**). Altogether, MSN population firing activity in the lateral striatum was more likely to show discriminatory firing, and that discriminating MSNs are more selective for licking than non-discriminating MSNs.

Previous work has shown that striatal neuron firing is altered as training progresses (L. Tremblay et al., 1998), suggesting that similar effects may be present in our recordings. We therefore investigated how discriminating and non-discriminating MSN firing changed over time by measuring the difference between mean CS+ and CS- evoked firing in blocks of 25 trials (**Figs. 1.5A, 1.5B**). We observed a steady change in discriminating MSN firing activity across trial blocks, but did not see this pattern in non-discriminating MSNs. A two-way, repeated-measures ANOVA revealed a significant effect of trial block ( $p<0.0001$ ,  $F_{3,48}=11.3$ ), and MSN population ( $p=0.0003$ ,  $F_{3,48}=21.83$ ) and showed an interaction between trial block and MSN population ( $p<0.0001$ ,  $F_{3,48}=10.76$ ). These results suggest that discriminating MSNs underwent a significant divergence from the remaining MSN population in encoding correct hit and withholding trials over time (**Fig. 1.5C**).

#### *Discriminating MSNs form a temporally correlated ensemble*

The large scale of our silicon microprobe measurements provided a unique opportunity to analyze correlated dynamics among hundreds of simultaneously recorded striatal cell pairs in order to study how these populations interact at the network level. Numerous studies suggest that correlated spontaneous neural activity reflects the underlying connectivity of the network

**Figure 1.4**

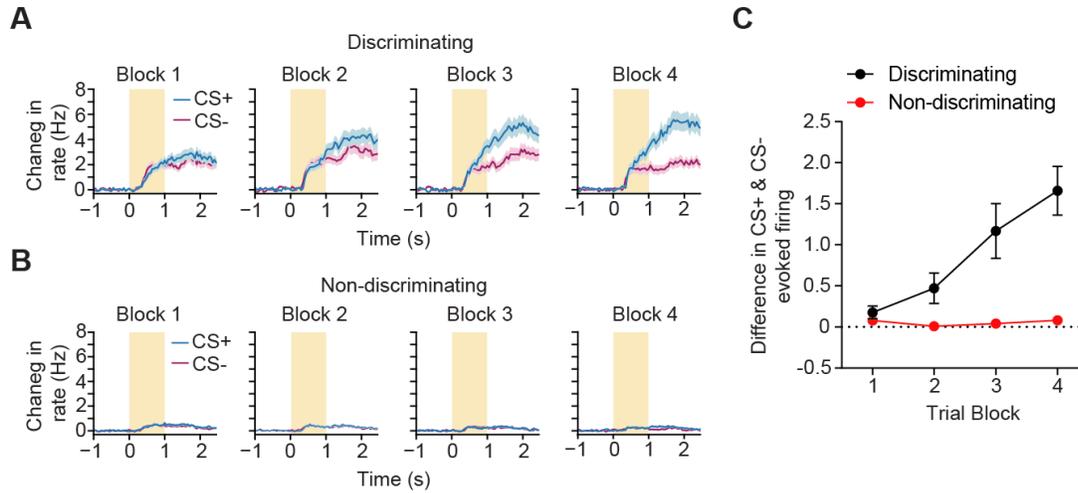


**Figure 1.4.** Mapping discriminatory activity across the striatal cross-section. *A*: Outline of the cross-section of the striatum spatially divided into a 4x3-compartment grid. Values represent the total number of recorded MSNs allocated into each of the grid's compartments based on the estimated recording position of each unit. *B*: Mean baseline subtracted firing rates for all MSNs positioned in each of the 12 boxes illustrated in *A*. Color conventions are identical to Fig. 3A. *C*: Combined map of the location of discriminating (blue) and non-discriminating (gray) MSNs recorded in all mice. Centers of all recordings were all aligned along the dotted line. *D, left panel*: Difference between the mean correct CS+ and CS- firing rates for all neurons binned by their mediolateral recording position. *Right panel*: Difference between the mean correct CS+ and CS- firing rates for all neurons binned by their dorsoventral recording position. Correlations were performed between position and difference in rate for MSNs pooled from all recordings (n=841).

(Ko et al., 2014; Ringach, 2009; Stern et al., 1998). As such, the resting state has become an important measurement modality in neuroscience (Cole et al., 2014; Raichle, 2010; Reimer et al., 2014). To focus on spontaneous activity we calculated pairwise spike count Pearson correlations of units recorded in the same session during intermittent periods when animals were at rest. This corresponded to times when mice were neither moving on the treadmill, nor licking or receiving explicit olfactory or reward stimuli (**Fig. 1.6A**).

We found that the strength of resting correlations decreased as a function of MSN pair separation ( $p < 0.0001$ ,  $F_{(9,24695)} = 63.75$ , one-way ANOVA; **Fig. 1.6B**). This spatial clustering suggests that neighboring cells are more likely to share information, possibly via overlapping input. We next tested whether correlated resting state dynamics show any resemblance to correlated activity during the behavioral task. Indeed, we identified a significant relationship between resting and task-evoked signal correlations during correct CS+ hit trials ( $p < 0.0001$ ,  $r = 0.104$ , permutation test for correlations; **Fig. 1.6C**). Since the small magnitudes of resting state correlation coefficients are difficult to interpret (Cohen and Kohn, 2011), we identified correlations that could not occur by chance ( $p < 0.01$ , permutation test for correlations, see Materials and Methods). We found that MSNs with significant resting correlations had a higher signal correlation coefficient than uncorrelated MSNs ( $p = 0.0008$ , paired t-test; **Fig. 1.6D**). Together, these results demonstrate that the resting state of the striatal network contains an intrinsic organization that is related to behaviorally modulated dynamics. This relationship is consistent with a model in which common input drives correlated spontaneous as well as task-evoked striatal activity. Using resting state correlation analysis, we asked if discriminating cells were distinct from the remaining population in terms of their likelihood of showing significant resting correlations. Indeed, we found that discriminating MSNs were more likely to exhibit significant correlations with one another than to non-discriminating MSNs ( $p = 0.0096$ , paired t-test; **Fig. 1.6E**). In addition to being more correlated in their spontaneous activity, discriminating MSNs were more excitable as revealed by their higher spontaneous firing rate ( $p = 0.0109$ ,

**Figure 1.5**



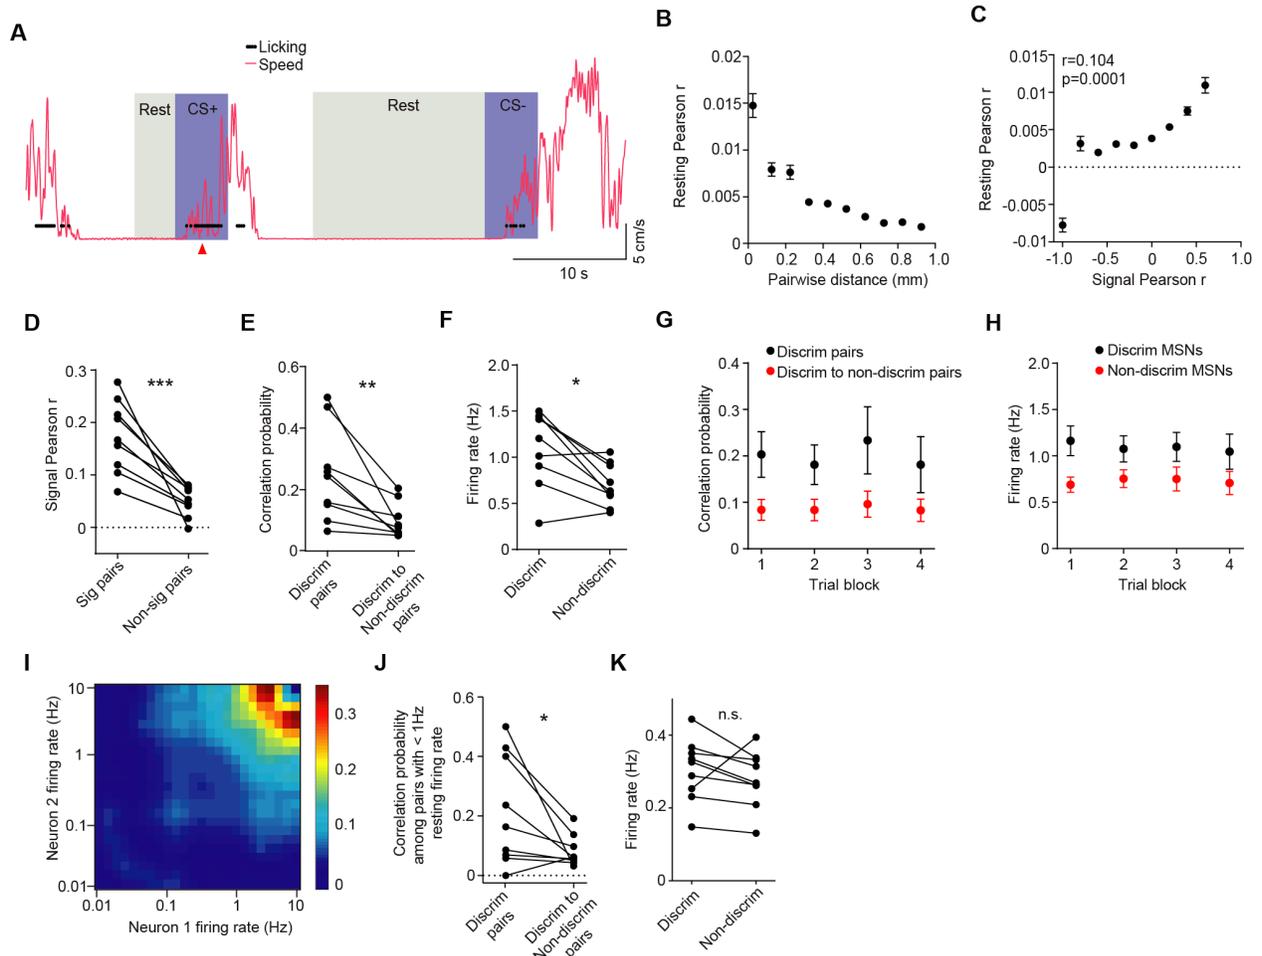
**Figure 1.5.** Evolution of activity in discriminatory MSNs during training. *A*: Mean baseline subtracted firing rate for all discriminating MSNs during the period starting 1s before cue onset until reward delivery for all CS+ trials (blue) and all CS- trials (magenta). Each panel depicts firing activity for each trial type in blocks of 25 trials. *B*: Same as *A* for non-discriminating MSNs. *C*: Mean firing rate difference between CS+ trials and CS- trials for discriminating and non-discriminating MSNs in blocks of 25 trials. A two-way, repeated measures ANOVA revealed a significant effect of trial block ( $p < 0.0001$ ) and population ( $p = 0.0003$ ) and a significant interaction between the two variables ( $p < 0.0001$ ). Averages were computed across individual animals,  $n = 9$ . Error bars represent SEM.

paired t-test; **Fig. 1.6F**). We next investigated whether discriminating MSNs changed their correlation strengths over time during the course of the recording session. The recording was divided into four blocks corresponding to 25 CS+ trials, and network connectivity was analyzed separately during resting periods occurring within each block of trials. We found that there was no effect of time in either the probability of significant correlation among discriminating MSNs or between discriminating and non-discriminating MSNs ( $p=0.143$ ,  $F_{(3,48)}=1.85$ , two-way, repeated measures ANOVA; **Fig. 1.6G**). Similarly, there was no significant effect of time in the resting state firing rates ( $p=0.8$ ,  $F_{(3,48)}=0.34$ , two-way, repeated measures ANOVA; **Fig. 1.6H**). These findings suggest that resting state correlations among discriminating MSNs remains relatively stable over the course of the first training session. However, we cannot rule out the possibility that resting correlations or firing change over more extended periods of training.

Firing rate is known to influence neural correlations between cells (Cohen and Kohn, 2011; la Rocha et al., 2007). On one hand, there may be a biological basis for this relationship: for example, experiments in cortical circuits show that highly active cells are more likely to be coupled (Yassin et al., 2010). But to examine whether temporal correlations are purely explained by higher firing, we looked at the relationship between resting firing rate and correlation probability. As expected, the probability of finding significant resting state correlations was greatest when both MSNs had relatively high ( $>1$  Hz) firing rate (**Fig. 1.6I**). Next, to account for the effects of firing rate we examined the correlation probability for pairs of cells whose resting rate did not exceed 1 Hz. Discriminating MSNs remained more likely to be correlated to one another than to non-discriminating MSNs ( $p=0.042$ , paired t-test; **Fig. 1.6J**) despite their no longer having significantly higher resting state firing ( $p=0.3$ , paired t-test; **Fig. 1.6K**). Together, these results show that with respect to non-discriminating MSNs, discriminating neurons represent a functionally specialized network in the striatum that is more tightly correlated during both spontaneous and task-evoked activity.

*FSIs contribute to striatal microcircuit synchrony*

**Figure 1.6**



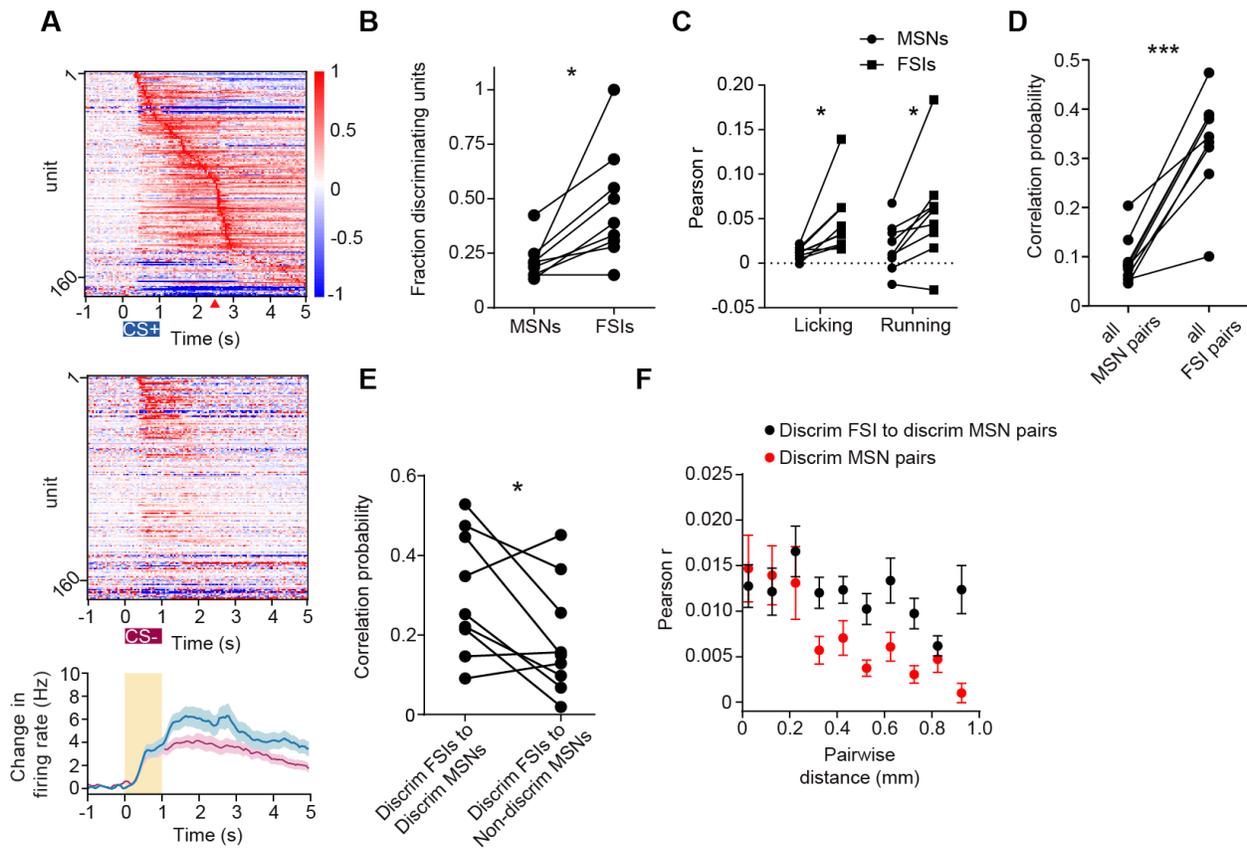
**Figure 1.6.** Correlated resting state activity in the striatum. **A:** Sample data depicting resting state identification. The black circles and magenta traces represent individual licks and running speed on the treadmill, respectively. Blue shaded regions label a 5 second window following cue onset. Gray shaded regions depict resting periods that would be concatenated with other resting periods for resting state analysis. **B:** Mean resting correlation coefficient for all MSN pairs plotted as a function of pairwise distance ( $p < 0.0001$ , one-way ANOVA). Data are binned in 0.1 mm increments. **C:** Mean resting correlation coefficient for all MSN pairs plotted as a function of the pair's signal correlation during correct CS+ trials. Binned data show a strong relationship between these parameters, and unbinned data are also correlated ( $n = 23758$  pairs, permutation test for correlations). Removing the outlier point in the left-most bin did not change the significance of the correlation ( $p < 0.0001$ ,  $r = 0.104$ , permutation test for correlations). **D:** Mean signal correlation coefficient for functionally connected (FC) and non-FC MSN pairs during spontaneous activity in the resting state. Points represent mean values of individual animals ( $p = 0.0008$ , paired t-test,  $n = 9$ ). **E:** Probability of finding significant resting correlations (*i.e.*, functional connections) among discriminating MSN pairs, and between discriminating to non-discriminating MSN pairs ( $p = 0.0096$ , paired t-test). Points represent the fraction of pairs spaced within 0.025 to 1 mm recorded from individual animals.

**Figure 1.6 (continued).** *F*: Mean resting state firing rate of discriminating MSNs and non-discriminating MSNs ( $p=0.011$ , paired t-test). Points represent the mean rate in individual animals. *G*: Resting state correlation probabilities among discriminating MSN pairs and between discriminating to non-discriminating MSN pairs calculated during resting times that occurred in different blocks of the recording. Each trial block represents resting periods detected within blocks of 25 CS+ trials. *H*: Resting firing rates for discriminating and non-discriminating MSNs calculated during blocked resting periods. *I*: Mean probability of finding significant pairwise resting correlations, as a function of the firing rate of each cell in the pair. Color scale represents significant correlation probability. *J*: Resting state correlation probabilities among pairs of discriminating and between pairs of discriminating to non-discriminating MSNs that had firing rates less than or equal to 1 Hz ( $p=0.042$ , paired t-test). *K*: Mean resting state firing rate of discriminating MSNs and non-discriminating MSNs having firing rates  $< 1$  Hz. ( $p=0.3$ , paired t-test). All error bars are SEM.

The striatal microcircuit contains a population of fast spiking interneurons which are thought to be involved in regulating striatal output signals (Berke, 2011; Koós and Tepper, 1999). As observed for MSNs, FSIs activity was modulated during the task, forming trajectories that qualitatively resembled the MSN population in their responses to CS+ and CS- trials (**Figs. 1.7A**). We found that FSIs were even more likely to be classified as discriminating than MSNs ( $p=0.011$ , paired t-test; **Fig. 1.7B**) and were more correlated to licking rate ( $p=0.016$ , paired t-test; **Fig. 1.7C**) and running velocity ( $p=0.012$ , paired t-test) than MSNs. Approximately 15% of the recorded units were putative FSIs, providing a sufficiently large population for resting state correlation analysis. Resting state firing between FSIs was also more likely to be correlated than between MSNs ( $p<0.0001$ , paired t-test; **Fig. 1.7D**). We found that discriminating FSIs were more likely to be correlated in their resting state activity to discriminating MSNs than to non-discriminating MSNs ( $p=0.04$ , paired t-test, **Fig. 1.7E**). We also compared the distributions of resting correlations among discriminating MSN pairs and pairs of discriminating MSNs and FSIs as a function of pairwise distance. We found that there were significant effects of both population ( $p<0.0001$ ,  $F_{(9,4446)}=31.66$ , two-way ANOVA; **Fig. 1.7F**) and pairwise distance ( $p<0.0001$ ,  $F_{(9,4446)}=4.96$ ), suggesting that discriminating FSIs are more likely to be coupled with discriminating MSNs at greater distances than discriminating MSNs are coupled to each other. Together, these results show that discriminating FSIs and MSNs form a temporally correlated ensemble.

We next searched for potential microcircuit mechanisms for how striatal cell types might interact locally. We analyzed spike time cross correlations to assess the temporal relationship between cells on a millisecond timescale. We used a spike time jitter test (Fujisawa et al., 2008) to identify pairs with significant low latency cross correlations, consistent with the occurrence of direct synaptic coupling between these cells. To characterize average significant cross correlation trends we performed an analysis on the entire recording session and pooled results from all animals in the study ( $n=9$  mice). Only 0.1% (38/28452 pairs) of MSN pairs showed

**Figure 1.7**



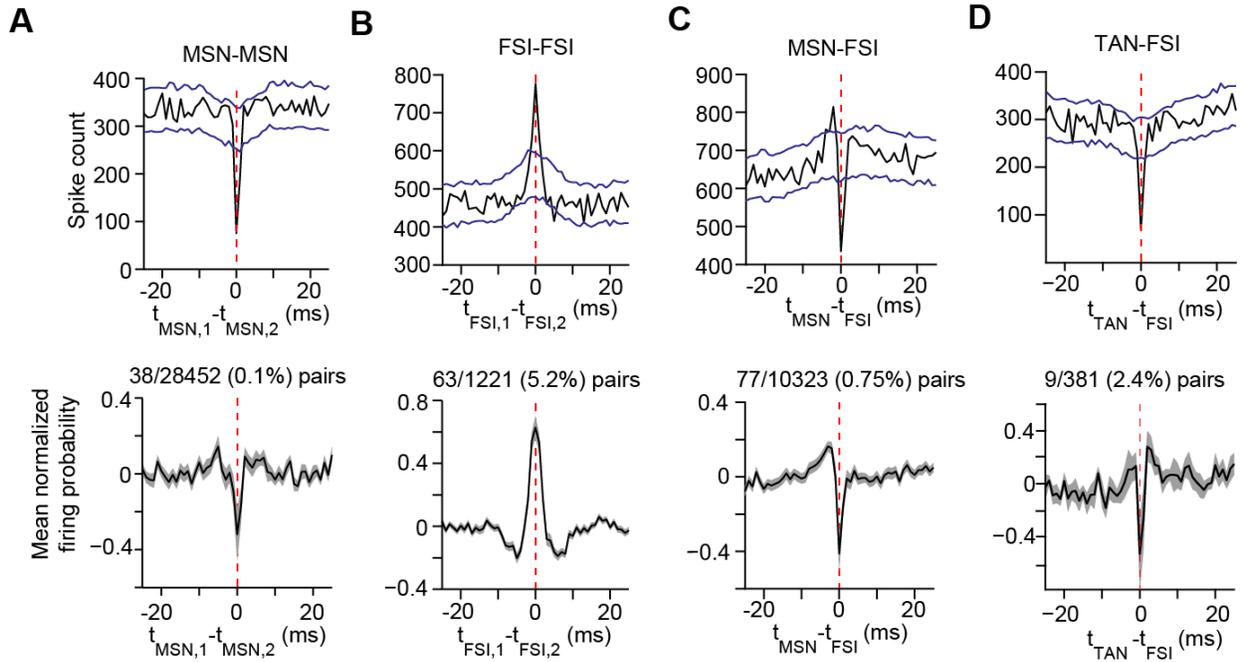
**Figure 1.7.** Discriminating FSIs and MSNs form functional ensembles. **A:** Mean baseline subtracted and normalized firing rates for all 178 FSIs recorded from 9 animals during correct CS+ trials (top) and correctly withheld CS- trials (center). Units in both plots are sorted by latency to peak firing in the top plot. Cues are presented between 0 and 1 s, indicated by colored rectangles. Reward delivery during CS+ trials is indicated with the red triangle. Bottom panel shows the mean baseline subtracted firing rate for all FSIs depicted in the heat plots. Orange rectangle represents odor delivery time. **B:** mean fraction of cue-discriminating MSNs and FSIs ( $p=0.011$ , paired t-test). **C:** Mean Pearson correlation coefficient between spiking activity of all MSNs or FSIs, and lick rate, ( $p=0.016$ , paired t-test) or treadmill velocity ( $p=0.012$ , paired t-test) behavior. **D:** Probability of finding significant resting correlations among pairs of MSNs and FSIs ( $p<0.0001$ , paired t-test). **E:** Probability of finding significant resting correlations among pairs of discriminating FSIs and MSNs and between pairs of discriminating FSIs and non-discriminating MSNs ( $p=0.0395$ , paired t-test). **F:** Mean resting correlation coefficient for all discriminating FSI and MSN pairs and all discriminating MSN pairs plotted as a function of pairwise distance. A two-way ANOVA revealed a significant effect of population ( $p<0.0001$ ) and pairwise distance ( $p<0.0001$ ). Data pooled from all animals are binned in 0.1 mm increments. Error bars are all SEM.

evidence of low latency cross correlations typical of monosynaptic coupling (**Fig. 1.8A**). FSIs may also contribute to MSN activity through local interactions. Because of gap junction coupling FSIs form a highly interconnected network (Hjorth et al., 2009; Koós and Tepper, 1999; Lau et al., 2010; Russo et al., 2013). In agreement with their predicted synchrony, we found that FSIs had a higher interaction rate than MSNs (5.2%, 63/1221 pairs; **Fig. 1.8B**), and their cross correlation showed positive and symmetric coupling. Furthermore, we found that a subset of FSIs are coupled with MSNs (0.75%, 77/10323 pairs; **Fig. 1.8C**) in a manner consistent with feedforward inhibition (Koós and Tepper, 1999). Our results show that the likelihood of finding significant FSI-MSN interactions with this analysis is relatively small, and another study failed to find these interactions altogether (Gage et al., 2010). We also detected a small population (~4%) of tonically active neurons in our recordings, and found that some FSIs were coupled to TANs (2.4%, 9/381 pairs; **Fig. 1.8D**). We detected very sparse TAN-MSN (0.05%, 1/2085 pairs) and zero TAN-TAN connections (0/88 pairs). Taken together the cross correlation analysis suggests that FSIs form a highly synchronized subnetwork of cells that can alter the timing of MSN activity, although the impact of these interactions on circuit activity during behavior is not yet fully understood.

### **Discussion for Chapter 1**

This study used large-scale neural recordings to characterize striatal network activity in mice learning to associate specific odor cues with rewards. We initially focused our analysis on single-unit activity during the task and identified a subpopulation of putative MSNs whose firing responses discriminated between correctly performed CS+ and CS- trials. Similar responses in individual striatal units have been previously reported during cue discrimination tasks (Nicola et al., 2004; Setlow et al., 2003; L. Tremblay et al., 1998). When compared to non-discriminating units, discriminating MSNs displayed higher firing rate change during CS+ trials. The entire population of discriminating MSNs also responded with higher activity to CS+ trials than to CS- trials. These discriminatory responses were not uniformly distributed in the striatum. By mapping

**Figure 1.8**



**Figure 1.8.** *A, top panel:* Spike time cross correlogram between one pair of MSNs exhibiting significant cross correlation. Blue lines represent upper and lower 99% confidence intervals of the time-jittered cross correlation. *Bottom panel:* Mean jitter subtracted and normalized cross correlogram for all MSN pairs exhibiting significant cross correlation. The fraction (0.1%) indicates the proportion of MSN pairs recorded within 0.025 to 1 mm that exhibited significant cross correlation according to the jitter test. Dotted red lines are aligned to a time lag of 0 s. *B:* Same as A but for FSI pairs. *C:* Same as A but for MSN-FSI pairs. *D:* Same as A but for TAN-FSI pairs. Error bars represent SEM.

the recorded neural activity, we showed that discriminating MSNs were more likely to be found in the lateral side of the striatum. This appears to be consistent with the known anatomical organization of this structure; specifically, that the dorsolateral striatum receives significant input from the sensorimotor areas of the overlying cortex (Alexander et al., 1986; McGeorge and Faull, 1989). Furthermore, single-unit measurements in the lateral striatum of rodents have been previously found to represent sensory and motor activity of different body parts, including those of the orofacial area and forelimbs (Cho and West, 1997). In our study, discriminating MSNs were significantly more correlated with lick rate than non-discriminating MSNs. Hence, this discriminating population could be involved in generating striatal output signals that mediate stimulus-specific anticipatory licking. We also showed that discriminating MSNs were modulated by running speed. However, running speed was equally correlated to both discriminating and non-discriminating neuron firing. This may also be evidence that some discriminatory neurons are modulated by both types of behaviors, supporting a report on the multiplexed coding properties of striatal neurons (Rueda-Orozco and Robbe, 2015). These results do not rule out that other aspects of the task also contribute to discriminating MSN coding.

Since there have been reports that the dorsomedial striatum is engaged during early stages of procedural learning (Thorn et al., 2010; Yin et al., 2009; 2005), it was slightly unexpected that we observed little change in activity in the medial striatal subregions (**Fig. 1.4B**). This might reflect differences in learning or behavioral strategies between our study and other work. The dorsolateral striatum has been implicated in mediating stimulus-response associations that appear in well-trained animals (Corbit and Janak, 2007). However, some studies have reported that lesions to the dorsolateral striatum also impair acquisition of stimulus-response associations as well (Featherstone and McDonald, 2005). Our results show that at least in the case of the Pavlovian reward association task employed here, the lateral striatum appears to be more active than the medial regions in the early stage of training. Further work needs to elucidate the role of specific microcircuits in the lateral striatum in acquiring and

expressing this behavior. In addition to evidence that the medial and lateral striatal subregions have different functional contributions to behavior, the ventral striatum (nucleus accumbens) is also known to have unique reward processing functions in comparison to the dorsal striatum. This is thought to be partially based on the limbic inputs that this area receives (Voorn et al., 2004). We did not detect a strong relationship between MSN position along the dorsoventral axis and the degree of differential coding between CS+ and CS- trials. This was evidence that the dorsolateral and ventrolateral striatum both contained discriminating units. Thus as an approximation, we assumed no difference in temporal correlation properties among pairs of discriminating neurons throughout the striatum. Although this is likely to be an oversimplification, our recordings lacked the single-unit throughput to reliably examine differences in correlations between the dorsal and ventral striatum. Furthermore, our assumption does not imply that these subregions encode identical information, and indeed, the ability to map activity across a section of the striatum revealed marked differences in cue-evoked neural dynamics.

We took our analysis of discriminating units in a new direction by investigating temporal correlations among this population. In order to avoid the potentially confounding effects of behavior and stimuli on neural activity and correlations, we examined spontaneous activity, which coincided with periods when animals were at rest. An extensive body of literature has shown that spontaneous neural activity is related to behaviorally evoked activity (Arieli et al., 1996; Raichle, 2010; Ringach, 2009). We extended this principle to striatal microcircuits by demonstrating that resting state MSN correlations are correlated to signal correlations. We speculate that experience-dependent plasticity in the striatum may establish neural ensembles whose synaptic connectivity predisposes them to fire together both during behavior (leading to high signal correlation), and thus also during rest (leading to high resting correlation) (Fregnac, 2003). A novel finding of this study was that discriminating units are more likely to exhibit significant resting state correlations to each other than to non-discriminating units. Correlated activity among neurons is attributed to common sources of input (Cohen and Kohn, 2011). We

therefore postulate that an important factor that mediates the observed pattern of correlations between discriminating MSNs is shared glutamatergic connections (Stern et al., 1998; C. J. Wilson, 2013). Glutamatergic axons that innervate the striatum are widely distributed (Kincaid and C. J. Wilson, 1996), which may enable multiple MSNs to receive the same excitatory signals. But since individual axons form only a few synapses with any individual MSN, eliciting action potentials may require coordinated glutamatergic activity from many axon terminals representing a diverse range of input (Kincaid and C. J. Wilson, 1996). This may explain why the activity of the recorded MSNs was correlated with more than one type of behavior (e.g., licking and running).

We did not find that discriminating neurons became more correlated in their spontaneous activity during learning, nor did they show a change in their excitability. This observation could be explained either because our analysis methods did not have the temporal resolution to observe the changes, or that the discriminating MSN ensemble was already established prior to the experiment. In the latter scenario, discriminating neurons may already have been established by an animal's past experiences in its home cage prior to training, which would have involved some form of licking and running behaviors. Future work will be required to better understand how striatal neurons are selected to participate in a specific ensemble and how resting state correlations are shaped by experience. Training animals on tasks in which they must perform arbitrary movements that they never encountered before may be a promising direction.

In addition to shared glutamatergic inputs, temporal correlations among MSNs are likely to be mediated by striatal interneurons such as FSIs, which are thought to control MSN firing on millisecond timescales (Damodaran et al., 2014). FSIs receive input from many of the same external sources as MSNs (Fino and Venance, 2011) and are thus able to encode behaviorally relevant information. Indeed, we found that FSI firing rate was correlated to licking and running to an even greater degree than MSNs, and FSIs were more likely to be classified as

discriminating units than MSNs. Discriminating FSIs were also more likely to be correlated with discriminating MSNs over non-discriminating MSNs. Many studies have shown that FSIs provide feedforward GABAergic inhibition of MSN spiking activity (Gittis et al., 2010; Koós and Tepper, 1999; Mallet et al., 2005; Taverna et al., 2007). We found 0.75% of FSI-MSN pairs exhibited significant cross correlations *in vivo*, providing evidence for direct interactions between these two subpopulations of striatal neurons. It is interesting to note that another study did not find any significant cross correlations between these cell types *in vivo* (Gage et al., 2010). This suggests that FSI-MSN interactions may also occur over long timescales relative to single action potentials because of polysynaptic network effects in the striatum. FSI-FSI connectivity is complex as it involves both chemical and electrical synapses (Berke, 2011; Fukuda, 2009; Kita et al., 1990; Russo et al., 2013), whose collective influence on network activity in the intact brain, or behavior is not well understood.

Another possible mechanism by which MSN activity can be correlated is through direct MSN-MSN connections. Despite a well-known effect of lateral inhibition among MSNs (Czubayko and Plenz, 2002; Taverna et al., 2004; Tunstall et al., 2002), the role of this inhibitory coupling in striatal computation is still unclear (Tepper et al., 2008). Our cross correlation analysis found very few significant low latency interactions among nearly 30,000 MSN pairs (only 0.1%), suggesting that individual MSN-MSN interactions are weak compared to other factors that influence MSN activity (D. Jaeger et al., 1994). Furthermore, inhibitory post-synaptic potentials (IPSPs) between MSNs have been found to be weak (Tunstall et al., 2002) relative to IPSPs evoked by FSIs on neighboring MSNs (Koós and Tepper, 1999). Interestingly, the ratio of these IPSPs is in the same order of magnitude as the ratio of significant FSI-MSN to MSN-MSN cross correlation pairs in our study, suggesting that cross correlations are sensitive to synaptic strength between two neurons. Thus, our work appears to show that MSN-MSN interactions are sparsely detected using cross correlation analysis, and that FSI-MSN coupling is 7.5 times more prevalent at the level of individual cell pairs. However, because of the abundance of MSNs in

the striatum compared to FSIs, which only represent ~1% of the total population (Berke, 2011), the cumulative effect of lateral inhibition may have a significant impact on striatal microcircuit dynamics, particularly during periods when large groups of MSNs are synchronized (Carrillo-Reid et al., 2008).

Finally, temporal correlations in striatal microcircuits are believed to strongly depend on neuromodulatory signals such dopamine and acetylcholine. Dopamine has a well-known role in modulating MSN activity and plasticity (Gerfen and Surmeier, 2011). Abnormally low levels of dopamine found in Parkinson's disease have been linked to excessive synchrony of striatal ensembles (Jaidar et al., 2010), as well as altered FSI-MSN connectivity (Gittis et al., 2011), both of which could also significantly impact temporal correlations. Cholinergic interneurons are sparsely distributed in the striatum, but have been shown to significantly impact striatal activity (English et al., 2012). In addition to modulating the release of dopamine (Cachope et al., 2012; Threlfell et al., 2012) and GABA (Nelson et al., 2014) from midbrain dopaminergic terminals, cholinergic signaling has also been shown to regulate the efficacy of corticostriatal input in eliciting action potentials in MSNs (Pérez-Ramírez et al., 2015; Shen et al., 2005). Cholinergic interneurons are thought to correspond to putative TAN units (Aosaki et al., 1994). We found only 1 significant MSN-TAN cross correlation out of over 2000 pairs (<0.05%), which appears consistent with TANs influencing MSN activity on long timescales relative to single action potentials. On the other hand, significant FSI-TAN cross correlation events were relatively common (9/381, 2.4%). Coupling between these cells (Koos and Tepper, 2002) may indirectly influence MSN correlations as well. However, the role of interactions between TANs and other interneurons in coordinating MSN dynamics remains unclear.

In conclusion, large-scale neural recordings enabled an examination of temporal correlations among hundreds of electrophysiologically defined striatal neurons in animals undergoing reward conditioning. We identified a population of cue discriminating striatal neurons that were more highly correlated to each other than to non-discriminating units. These results

suggest that discriminating units represent a functionally specialized ensemble with a higher occurrence of shared connections from both external as well local sources. Thus, temporal correlations among specialized neurons may help to pattern a strong output signal that is sent to downstream basal ganglia nuclei in order to facilitate behavior.

## **Chapter 2: Differential encoding of time by prefrontal and striatal network dynamics**

### **Introduction:**

Anticipating events that will happen in the future is among the most important functions the brain performs. Indeed, it has been increasingly stressed that learning and memory are prospective brain functions—that is, they are only adaptive to the extent that they help animals anticipate and prepare for the future (Dudai and Carruthers, 2005; Schacter and Addis, 2007). In order to anticipate *when* events will happen, the brain has evolved mechanisms to tell time across a wide range of temporal scales (Buhusi and Meck, 2005; Buonomano, 2007).

Timing on the scale of hundreds of milliseconds to a few seconds is of particular importance in that it allows animals to predict and prepare for events unfolding within the immediate future. Within this range animals discriminate the temporal features of sensory stimuli—such as those used for communication—and generate timed motor responses to prepare for external events—such as expected rewards. The neural mechanisms underlying the brain’s ability to tell time on the scale of seconds remains unknown (Mauk and Buonomano, 2004; Merchant et al., 2013a); but a rapidly growing literature has reported that dynamically changing patterns of neural activity encode information about the amount of time elapsed since a given stimulus. These patterns of activity—which have been referred to as population clocks (Buonomano and Karmarkar, 2002; Buonomano and Laje, 2010; Buonomano and Maass, 2009)—have now been observed in a wide range of different brain areas, including the striatum (Bakhurin et al., 2016; Chiba et al., 2008; Gouvêa et al., 2015; D. Z. Jin et al., 2009; Matell et al., 2003; Mello et al., 2015), prefrontal cortex (Brody et al., 2003; Carnevale et al., 2015; Genovesio et al., 2009; D. Z. Jin et al., 2009; Kim et al., 2013; Merchant et al., 2011; Oshio et al., 2008), parietal cortex (Crowe et al., 2010; Janssen and Shadlen, 2005), hippocampus (Kraus et al., 2013; Pastalkova et al., 2008), as well as in the bird song system (Hahnloser et al., 2002; Long et al., 2010). Additionally, pharmacological, lesion, and neuroimaging work

suggests a role of the basal ganglia (Coull et al., 2011; Meck, 1996) and prefrontal cortex (Dietrich and Allen, 1998; Kim et al., 2009; M. Xu et al., 2014) in timing.

The diversity of areas implicated in timing likely reflects the range of tasks and temporal scales examined. But additionally, it is possible that even within the same task, different areas track time in parallel (D. Z. Jin et al., 2009; Matell et al., 2003). To date, however, no single study has directly quantified the degree to which two different circuits encode time through simultaneous multiple-region recordings. Here we directly contrast the ability of two circuits, the striatum and orbitofrontal region (OFC) of the prefrontal cortex to encode time.

We examined the neural representation of time during a Pavlovian conditioning task in which a food reward is presented at a specific interval after a conditioned stimulus (CS). Mice exhibited anticipatory licking during the fixed cue-reward delay period. Silicon microprobe recordings of dozens of units from either the striatum or OFC, or both simultaneously, revealed that population activity in both circuits encoded an internal representation of elapsed time. This code was quantified by feeding the trial-by-trial spike pattern into a pattern classifier, and training it to read out elapsed time. The quality of the striatal population code for time was significantly better than that of the OFC. Our results support the hypothesis that many different brain areas simultaneously encode time, but that the striatum may play a privileged role in timing relative to the OFC as it holds a more accurate clock. We hypothesize that by continuously sampling the changing patterns of activity unfolding throughout the cortex and other inputs, the striatum implements a robust code for elapsed time via a temporal winners-take-all mechanism.

## **Materials and Methods:**

### *Animals and surgical procedures*

All procedures were approved by the University of California, Los Angeles Chancellor's Animal Research Committee. Singly housed male C57Bl/6J mice (n = 11, 15-22 weeks old at the time of recording, The Jackson Laboratory) were used in the experiments. Animals underwent an

initial head bar implantation surgery under isoflurane anesthesia in a stereotaxic apparatus to bilaterally fix stainless steel head restraint bars on the skull with dental cement. After training, animals underwent a second surgery under isoflurane anesthesia on the recording day to make craniotomies for acute microprobe recordings. An additional craniotomy was made over the posterior cerebellum for placement of an electrical reference wire. All behavioral training and recording sessions were carried out in fully awake head-restrained animals.

### *Behavioral task*

After a one week recovery period following the initial head bar implantation surgery, animals were food restricted and fed daily after each training session to maintain ~90% of their baseline weight. Water access was *ad libitum*. During daily training sessions, animals were mounted on the head bar restraint bracket on the recording rig and stood on a polystyrene treadmill ball (200 mm diameter, Graham Sweet Studios) that rotated along a single-axis during forward/backward ambulation. Animals were initially habituated to the head-fixed recording rig and trained to consume a liquid reward (5  $\mu$ L, 10% sweetened condensed milk). The reward was delivered from a tube positioned between an infrared lick meter (Island Motion) by actuation of an audible solenoid valve (Neptune Research). During daily reward-only training sessions, animals consumed 100 rewards and were exposed to a constant stream of pure air through a tube positioned next to the nose (100 rewards per session, 13-21 s inter-trial interval (ITI), sampled from a normal distribution, 1.5 L/min air flow). Once animals could consume  $\geq 90\%$  of the rewards for two consecutive days, they began conditioning with olfactory cues using an olfactometer. Odorants were introduced by bubbling air (0.15 L/min) through aromatic odorants diluted 1:10 in mineral oil (Sigma-Aldrich), and merging this product with the 1.5 L/min stream of pure air. The constant flow of pure air into which odors are introduced decreased the possibility that animals used decaying concentrations of odorant as a temporal cue. During daily training sessions, animals received pseudorandom presentations of each odor stimulus (1 s duration,

17-29 s ITI, sampled from a normal distribution). Isoamyl acetate served as the CS+ odor, as its offset was followed by a 1.5 s delay and a reward delivery. Citral served as the CS- odor as it was not followed by any explicit outcome. Animals received 100 presentations of each trial type in random order during each training session. The solenoid valves controlling the odors were sound isolated and thus inaudible to the mouse. Typically during the first or second day, animals began predicting the delivery of the reward by licking in anticipation during the interval between the odor and the reward. Correct CS+ trials were defined as those trials during which licking was initiated prior to reward delivery (between 0.7 and 2.5 s following stimulus onset). Correct CS- trials were defined as those containing no licking activity for 5 s following stimulus onset. False alarm CS- trials were defined as those trials during which licking was initiated between 0.7 and 2.5 s following stimulus onset. Once animals demonstrated correct responding on  $\geq 90\%$  of trials, they underwent surgery for recording. During the recording session, animals received 100 CS+ trials with 85% reward probability and 100 CS- trials. Animals performed between 54 and 99 correct CS+ trials and between 1 and 56 false alarm CS- trials.

### *Electrophysiological recordings*

Procedures for developing and recording with silicon microprobes are described elsewhere (Shobe et al., 2015). One recording was performed per animal. Each area was targeted with a silicon microprobe containing a total of 256 electrodes that were divided across 4 or 5 prongs. The electrodes spanned between 0.825 to 1.05 mm of the distal tip of the prongs. Data in this study were aggregated from two groups of animals. In the first group ( $n = 5$ ), recordings took place in the anterior striatum only (silicon prong tip positions: 1.2 mm anterior, 0.8 to 2.2 mm lateral, -3.4 to -5.7 mm ventral relative to bregma). In the second group ( $n = 6$ ), we simultaneously recorded from the orbitofrontal region of the prefrontal cortex (2.2 mm anterior, 0.26 to 2.05 mm lateral, -3.6 mm ventral relative to bregma), and both the anterior and posterior regions of the striatum (anterior striatum: 1.2 mm anterior, 0.78 to 2.1 mm lateral, -5 mm ventral;

posterior striatum: -0.5 mm anterior, 2.4 to 3.2 mm lateral, -4.3 mm ventral tip position relative to bregma) using multiple 256 electrode probes attached together (Shobe et al., 2015). Thus the striatal dataset analyzed in this study was composed of the anterior striatal recordings performed in the first group, combined with anterior and posterior striatal recordings performed in the second group. The OFC dataset was composed of orbitofrontal recordings performed in the second group. Because of the wide spatial distribution of recording sites above the prong tips, the anterior striatal dataset contained units sampled from both dorsal and ventral striatal areas. Positions of units included in analysis are illustrated in **Figure 2.5A**. Spike sorting was performed on the data using custom, semi-automated software written in MATLAB (Mathworks, Cambridge MA). The placement of silicon probes was confirmed histologically at the end of each experiment by coating the prongs with a fluorescent dye (Di-D, Thermo Fisher) prior to implantation.

#### *Delineation of anatomical subregions*

In each animal, the recordings in the anterior striatum consisted of predominantly ventral or dorsally positioned units, with one recording containing units evenly distributed in each area. We used the mean electrode position of -4.2 mm DV to divide the anterior striatal recordings into dorsal or ventral regions. To divide the OFC into medial and lateral subregions, we used the mean electrode position of 1.19 mm in the OFC.

#### *Unit classification*

Analysis was performed on putative principal neuron populations, *i.e.*, pyramidal cells in the OFC and medium spiny neurons (MSNs) in the striatum. We used spike waveform trough-to-peak duration to distinguish putative MSNs and pyramidal neurons from non-principal neurons. Putative fast spiking interneurons (FSIs) were separated from principal cells in both the OFC and the striatum by their narrow waveform (maximum FSI  $t_{tr-pk} = 0.475$  ms, minimum principal

neuron  $t_{tr-pk} = 0.55$  ms, and maximum MSN  $t_{tr-pk} = 1.25$  ms (Bakhurin et al., 2016)). We also used a measure of firing rate regularity (coefficient of variation, CV) to exclude putative tonically active neurons from the striatal recordings (maximum CV = 1.5, (Bennett and C. J. Wilson, 1999)). We recorded a total of 690 putative MSNs out of a total of 1115 striatal units and 505 putative pyramidal cells out of a total of 654 cortical units.

#### *Identification of lick-modulated units*

We determined licking modulated units by correlating estimated firing rates with licking rate around lick episodes that occurred throughout the recording, including within and outside of trial periods. Licking episodes were defined as containing two licks that were separated by at most 250 ms (4 Hz). Licking episodes could not occur within 5 s of each other. To calculate the correlations, we binned individual licks occurring within a 2 s window around each lick episode into 50 ms time bins. For each unit, we binned spikes occurring around each licking episode within a 2 s window into 50 ms bins. The resulting episode vectors reflecting licking and spiking counts for each episode were concatenated into 2 vectors and convolved using a Gaussian function (SD = 100 ms) to obtain licking and spiking rate estimates across all lick episodes in the recording. A Pearson correlation was performed between the lick-rate vector and each spiking rate vector for each unit. We considered a unit to be lick-rate-modulated if it demonstrated a positive correlation coefficient with a p-value below 0.01.

#### *Elapsed time prediction analysis*

All analyses were performed independently on data collected from each animal and each brain region, using correctly performed CS+ trials or CS- trials with false alarm licking. All decoding models were generated using only simultaneously recorded cells from individual animals. For each trial, neural population activity was analyzed over the 2.5 s interval between cue onset and reward delivery. Over this interval, we transformed the activity of each neuron in the

simultaneously recorded population into an analog rate code estimate by: (1) convolving its spike train with a decaying exponential function ( $\tau = 100$  ms); (2) calculating its firing rate estimate as a binned average (100 ms time bins) of its convolved spike train. This procedure resulted in 25 population firing rate vectors, one per 100 ms time bin, in the trial.

Elapsed time was decoded from the population firing rates in each trial by requiring a classifier to label each rate vector in the trial as coming from one of the 25 time bins. The classification task was performed with a multi-class support vector machine (SVM) with a radial-basis function (RBF) kernel, as implemented in the LIBSVM library (v. 3.20, (Chang and Lin, 2011)). This SVM uses a one-against-one multi-class approach to distinguish the population firing rates encoding a given time bin from those encoding each of the 24 other time bins (Hsu and Lin, 2002; Knerr et al., 1990; Kreßel, 1999). In the one-against-one multi-class approach, binary classifiers are trained to distinguish between the population codes for each pair of distinct time bins  $(i, j)$ , for a total of 300 binary classifiers. SVM output is represented in 25 readout units, one per time bin. Given a test population rate vector, readout  $i$  generates a classification score indicating how closely this vector resembles the population code encoding bin  $i$ . It is calculated as an aggregate of the outputs of the 24 binary classifiers  $(i, 1), (i, 2), \dots, (i, i-1), (i, i+1), \dots, (i, 25)$ . The SVM predicts that the test vector encodes time bin  $k$ , whenever readout  $k$  produces the highest score of all 25 readout units (**Fig. 2.2**).

Individual animals showed varying numbers of correctly performed trials. To ensure that the decoding performance across animals was compared under equivalent conditions, the predicted time bins in all figures were generated with a Monte-Carlo cross-validation strategy. The rate vectors from each trial were tested on 30 independently trained SVMs, where each SVM was trained on the rate vectors from  $M$  randomly sampled trials excluding the test trial. Because the minimum number of correct CS+ trials for an individual animal was 54, we chose  $M$  to be 53.

We controlled for the number of simultaneously recorded units used to train and test the models. The number of simultaneously recorded cells,  $N$ , used to generate each decoding model, and the number of animals used for averaging is always indicated on the figure, or in the figure caption.  $N$  varied from 29 to 55 because of subregion-specific limitations in the number of simultaneously recorded units. Furthermore, to test the effect of the population size on model performance (**Figs. 2.3, 2.4, 2.9**), we compared random samples sizes of 5, 10, 15, 20 and 40 units taken from the entire striatum or entire OFC. During each of the 30 repetitions of the Monte-Carlo cross-validation, we also randomly sampled  $N$  distinct units from the population for training and testing. To maximize decoder performance, the RBF SVM regularization parameters were optimized for each brain region of each animal. Specifically, the misclassification cost parameter,  $C$ , and the data complexity parameter,  $\gamma$ , were optimized via a grid search with five-fold cross-validation. Across all datasets, the predominant value of  $C$  was 4 (range: 1-16), and of  $\gamma$  was 0.25 (range: 0.0156-0.25).

#### *Comparing population coding between correct CS+ and false alarm CS- trials*

To determine the extent to which the CS+ code for time generalized to CS- trials, we trained the classifier in the same way as described above, using 55 cells per area and using  $M = 53$  trials per Monte-Carlo cross-validation repetition. We then tested the models on the 25 rate-vectors generated for each false alarm CS- trial available for each animal (identical binning and rate estimation procedure as done for CS+ trials). This procedure was repeated 30 times, whereby random combinations of 55 units and 53 trials were employed in training the model.

#### *Lick onset prediction analysis*

For each trial, neural population activity was transformed into estimated population firing rate vectors using 100 ms bins, as in the elapsed time prediction analysis. This sequence, or *trajectory*, of neural population activity started 1 s before cue onset and ended 200 ms after the

latest lick onset time of all correctly performed CS+ trials. As a result, the number of time bins (and population rate vectors) analyzed per trial varied between 31 and 37 across animals.

Lick onset bins were predicted from the population firing rates in each trial with an RBF SVM binary classifier. The SVM's output is represented by a single readout that scores how closely each population vector in the test population trajectory predicts lick onset. The predicted lick onset bin was the one in which the readout was at its highest value. Testing was performed with a Monte-Carlo approach similar to the elapsed time prediction whereby each trial was tested on 30 SVMs independently trained on  $M = 53$  randomly sampled trials. The dataset contains disproportionately fewer lick onset bins than non-lick onset bins, because only a single bin out of the 31 to 37 bins per trial can be a lick onset bin. To avoid the resulting bias in the SVM model, the training set for each SVM was altered by (i) randomly down-sampling the subset of non-lick onset bins by 75%, and (ii) expanding the set of SVM target bins to include one bin immediately preceding and one bin immediately following the actual lick onset bin in each trial, for a total of 3 target bins per trial. The misclassification cost and data complexity regularization parameters for the RBF SVMs were optimized for each brain region of each animal similarly to the elapsed time prediction analysis. Across all datasets, the predominant value of  $C$  was either 2 or 8 (range: 2 to 128), and of  $\gamma$  was 0.125 (range:  $10^{-7}$  to 0.5).

The binary SVMs were retrained for each pre-lick time to determine how far in advance the neural trajectory could predict lick onset (**Fig. 2.10D**). At each pre-lick time, the SVMs were retrained to predict a new set of target bins that were appropriately shifted backward in time from the actual lick onset bin. During training, the data down-sampling procedure was accordingly altered to down-sample the subset of non-target bins

### *Trial shuffling*

Trial shuffling was used as a control for elapsed time prediction. This procedure disrupts correlations in simultaneously recorded population activity, but preserves the correct bin order

for each unit. To create trial shuffled activity, each unit's firing rate estimate in each time bin of each trial was replaced with the same unit's firing rate estimate in the same time bin of a randomly selected trial. This control could not be performed with the lick onset prediction analysis because of the resulting dissociation between lick onset times and unit activity.

### *Bin shuffling*

Bin shuffling was used to generate population responses that were dissociated from their correct temporal order. To create bin shuffled activity, each unit's firing rate estimate in each time bin was replaced with the same unit's firing rate estimate in a randomly selected bin of the same trial. Prediction analyses from trial and bin shuffled data involved training and then testing on the respectively transformed datasets. To allow for direct comparisons between observed, trial shuffled, and bin shuffled controls, care was taken to make sure that we subsampled the same units and trials for analysis.

### *Temporal warping of internal time representation*

Given our hypothesis that the population code for elapsed time and lick-onset time share a common internal timing representation, the population's encoding of the animal's internal representation of time should co-vary with the lick onset time. To measure this effect, correct CS+ trials were divided into three approximately equally sized sets corresponding to each of the terciles of the animal's lick onset distribution, and SVMs were trained to classify elapsed time in the 1<sup>st</sup> (3<sup>rd</sup>) tercile trials, and then tested on the 2<sup>nd</sup> and 3<sup>rd</sup> (1<sup>st</sup>) tercile trials. Biases in the resulting error distribution would then reveal an underlying comodulation. A more direct measurement of this effect was performed by comparing the temporal relationship between trial-averaged trajectories of 1<sup>st</sup> and 3<sup>rd</sup> tercile trials—if the 3<sup>rd</sup> tercile trajectory was consistently slower than the 1<sup>st</sup> tercile trajectory, this would indicate that the two timing variables comodulate one another. Population spike trains were convolved with a Gaussian function (mean

= 0, SD = 100 ms) and then trial averaged separately over the 1<sup>st</sup> and 3<sup>rd</sup> tercile trials, to produce characteristic 1<sup>st</sup> ( $T_1$ ) and 3<sup>rd</sup> ( $T_3$ ) tercile trajectories. These trajectories were compared by temporally aligning them as follows: (i) a Euclidean distance matrix was constructed by comparing the population code at each moment along  $T_1$  to the population code at each moment along  $T_3$  resulting in an  $N_T \times N_T$  distance matrix (**Fig. 2.7C**), where  $N_T = 2500$ , given the time resolution of the spike trains (1ms); (ii)  $T_3$  and  $T_1$  were temporally aligned with a dynamic time warping procedure that calculated the deviation of  $T_3$  from  $T_1$ , over the course of time, as the path along the distance matrix between the beginning and the end of  $T_1$  with minimum cumulative distance (black trace in **Fig. 2.7C**). The relative speed (temporal warping) of  $T_3$  with respect to  $T_1$  was indicated by the difference between the respective times at which the two trajectories were temporally aligned (**Fig. 2.7D**)—When  $T_3$  ran slower than  $T_1$ , this difference would be positive and monotonically increase, and when it ran faster the difference would be negative and monotonically decrease.

### *Effective dimensionality*

The effective dimensionality of each recorded population was calculated from trial-averaged population firing rate estimates. To control for the difference in the number of units measured across different recordings, the effective dimensionality for each recording was calculated as a mean over 30 randomly sampled subpopulations of size 55. Performing principal component analysis (PCA) on the dynamics of a single such sample produced a list of 55 principal components (PCs) ordered by the percentage of variance in the population dynamics explained by each PC. The effective dimensionality was calculated as the minimum number of PCs required to explain 95% of the variance in the dynamics (Rajan et al., 2016).

### *Statistical analysis*

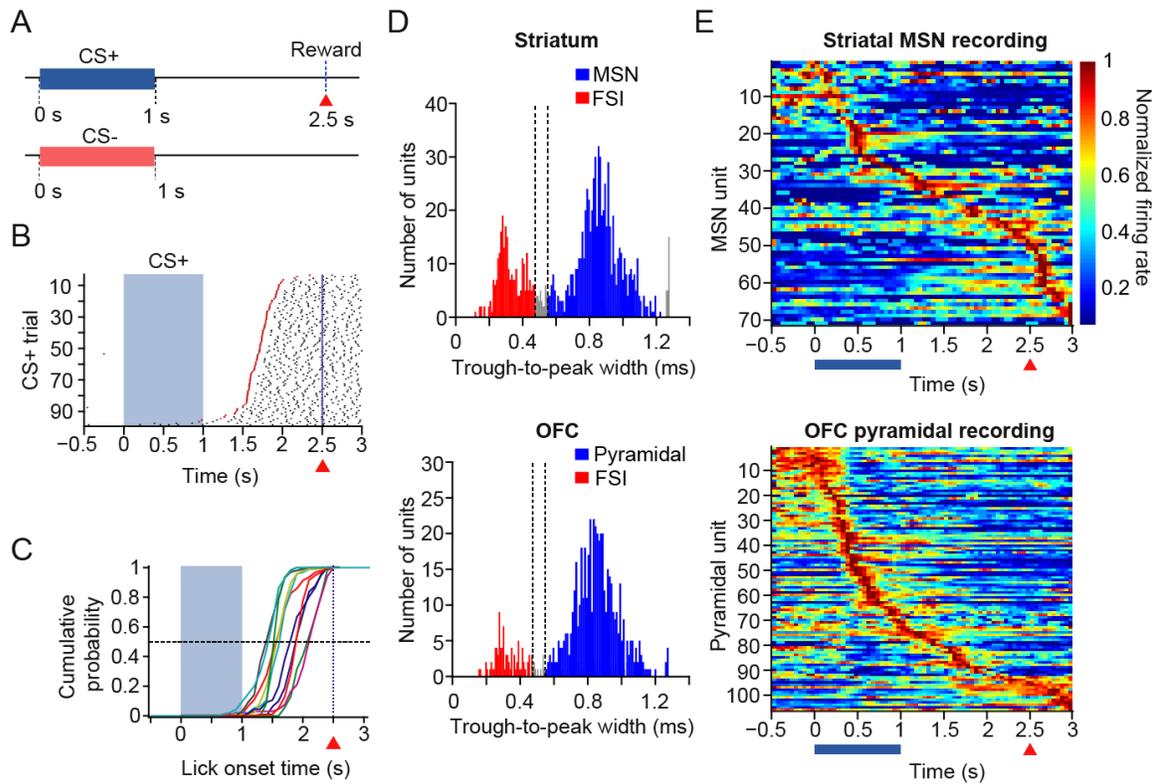
To determine the efficacy of the SVM models in elapsed time prediction, we calculated the correlation between the correct bin number and the predicted bin number. A single correlation coefficient was calculated from all the test data (*i.e.*, 25 time bins per correct CS+ trial in the dataset, 30 repetitions each) for a given brain region of a given animal. For clarity, graphs display correlation coefficients but statistical analysis was performed using Fisher's z-statistic for correlation coefficients (Fisher transformation). For the lick onset time prediction analysis, model accuracy was measured by the root mean squared error (RMSE) of the predicted lick onset bins. A single RMSE value was calculated from all the test data (*i.e.*, 30 repetitions for the lick onset bin in each of the correct CS+ trials) for a given brain region of a given animal. During hypothesis testing, we assumed that the population size used in the analysis represented a repeated measure because units were sampled from the same population of units. Brain region (*i.e.*, striatal versus OFC networks) was considered a repeated measure only when recorded in the same animal (**Figs. 2.5D, 2.6B, 2.6E, 2.9, 2.10**). Two-way repeated-measures and mixed-model ANOVA analysis was performed using GraphPad Prism (version 6.0). Two-sided paired and unpaired t-tests were performed using standard functions in MATLAB.

## **Results:**

### *Behavior*

We obtained large-scale recordings from the striatum and OFC in head-fixed mice ( $n = 11$  mice) previously trained to perform an odor discrimination task (Bakhurin et al., 2016; Shobe et al., 2015). In this task, mice were presented for 1 second with one of two olfactory stimuli. One of the odors (CS+) was followed by a reward delivered 2.5 s from cue onset. The delivery of the reward was not contingent on any instrumental actions of the animal. The second odor (CS-) was followed by no specific outcome (**Fig. 2.1A**). Following repeated presentations of the CS+ trials, animals learn to generate anticipatory licking behavior that preceded the reward delivery (**Fig. 2.1B**). Previous experiments from our group have demonstrated that animals time their

**Figure 2.1**



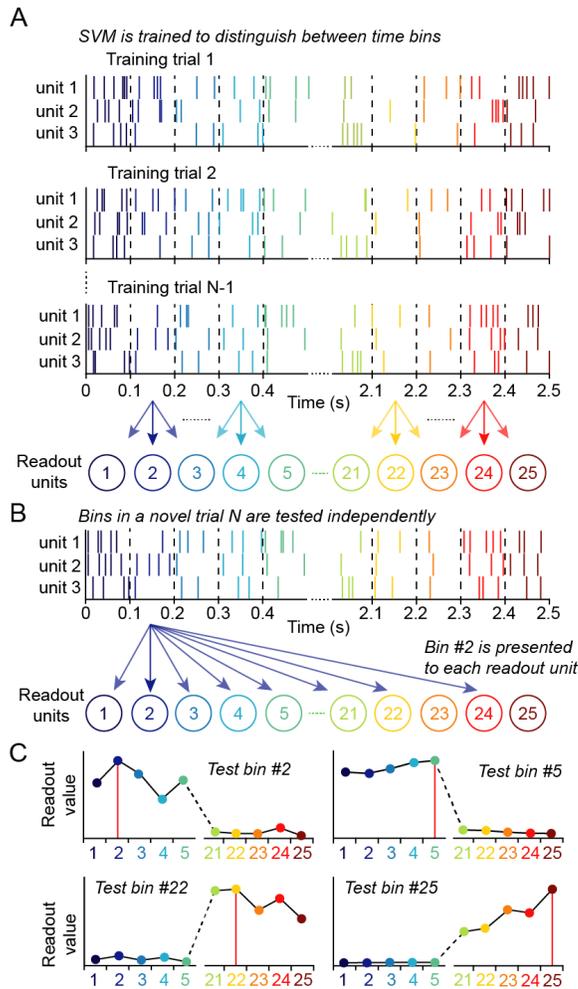
**Figure 2.1.** Large-scale recording of OFC and striatal networks during reward-predictive behavior. *A*: Task schema. Mice received pseudo-randomly ordered presentations of a CS+ odor that predicted reward delivery 2.5 s after odor onset and an unrewarded CS- odor. Rectangles represent odor-on time. Red triangle and vertical blue dashed line indicate reward delivery. *B*: Example of anticipatory licking behavior of one mouse during CS+ trials. Shaded blue rectangle represents odor presentation time. Black tick marks indicate individual licks, and red ticks denote lick onset times that are used for subsequent analysis. Trials are sorted by descending latency to first lick. *C*: Cumulative distributions of lick onset times during CS+ trials for all mice included in the study ( $n = 11$  mice). *D*: Distribution of the trough-to-peak width (ms) recorded from striatal units (top), and OFC units (bottom). Vertical dotted lines depict the threshold margin (0.475 to 0.55 ms) for segregating putative FSIs (red histograms) from putative principal cells (striatal MSNs and OFC pyramidal cells, blue histograms). Gray bars reflect unclassified cells. *E*: Individual population-level recordings from the striatum (top) and the prefrontal cortex (bottom) during correctly performed CS+ trials. Each row in a matrix represents the mean normalized firing rate of one recorded putative projection neuron in the corresponding brain area. Units are sorted by their latency to maximum firing rate. Blue rectangles indicate CS+ odor presentation time and red triangles mark the time of reward delivery.

anticipatory licking response depending on the cue-reward delay duration (data not shown), consistent with timed reward-guided behavior found in many other studies (Bermudez and Schultz, 2014). Our recordings were performed in animals that had experienced 5-10 training sessions and were performing above a criterion of at least 90% correctly performed trials (see Methods) prior to the recording day. The onset of anticipatory CS+ licking responses was concentrated during the cue-reward delay period for all animals studied (mean lick onset time = 1.8 s, SD = 0.25, **Fig. 2.1C**). We focused our analysis on correct CS+ trials, as these displayed discrete behavioral evidence that animals timed their behavior to anticipate the reward.

#### *Large-scale striatal and orbitofrontal recordings*

After animals reached criterion performance on the task, we used silicon microprobes (Shobe et al., 2015) to record population activity from either the striatum, OFC, or simultaneously from both of these areas, as the mice performed the task. We focused our analysis on putative principal cells in these brain regions: striatal MSNs and cortical pyramidal cells. If these brain areas contain a code for time, principal cells would be the most likely to transmit that signal to downstream brain regions (Buonomano and Merzenich, 1995). To identify these populations, we measured the action potential duration of each unit and used a threshold margin to segregate putative principal cells from fast spiking interneurons. In both the striatum and OFC, the distribution of spike widths across all cells was bimodal (**Fig. 2.1D**). Based on the separation of these distributions we only included putative principal cells in our analysis. We analyzed data from animals containing at least 55 principal units per region (n = 9 striatal recordings, and 6 OFC recordings). Our datasets contained between 55 and 120 simultaneously recorded principal neurons. We found that on average, the population of striatal and prefrontal neurons exhibited highly heterogeneous firing activity during the cue-reward interval (**Fig. 2.1E**). This observation is qualitatively similar to the sequential firing patterns reported from other cortical (Crowe et al., 2010; Harvey et al., 2012; Stokes et al., 2013) and striatal recordings (Bakhurin et

**Figure 2.2**



**Figure 2.2.** Schematic of the support vector machine decoding of elapsed time. **A:** Training the SVM. Single-trial spiking activity of each unit in a simultaneously recorded population (only 3 units represented) is transformed into a firing rate estimate for the unit during the 2.5 s interval following odor presentation onset (not shown here). The rate estimates are binned (100 ms time bins) to construct 25 population activity patterns per trial. Using a one-against-one multiclass strategy, the SVM trains a set of binary classifiers to distinguish the population activity pattern in each time bin from every other time bin. SVM output is conceptualized as 25 readout units, one per target time bin, that learn to distinguish activity patterns in their respective target time bin from those in all other bins. **B:** The model is tested using a Monte-Carlo cross-validation approach in which each activity pattern from novel trials (i.e. those excluded from the training set) is tested on trained SVM models. Illustrated is the testing of bin #2 of the test trial. **C:** Readout units score each test activity pattern for how closely it corresponds to their respective target bins. The target time bin of the readout with the maximal value is chosen as the predicted time in a winner-take-all manner (marked with a red vertical line). Actual readout values are depicted here.

al., 2016; Gage et al., 2010; Mello et al., 2015; Rueda-Orozco and Robbe, 2015; Thorn and Graybiel, 2014). However, from the average firing rate representation it was not evident whether the dynamics were robust at the single-trial level. We thus examined whether it was possible to decode elapsed time and lick onset time on a trial-by-trial basis during correctly performed CS+ trials.

### *Decoding time from network dynamics*

To investigate the possibility that neural network activity could provide a mechanism for the stable representation of time, we used a support vector machine (SVM) decoder to detect and measure the reoccurrence of dynamic population activity in striatal and OFC networks on a trial-by-trial basis. An SVM was trained to identify population activity in each of the 25 time bins (100 ms bin duration) between stimulus onset and reward delivery.

Each unit's firing rate for the 25 time bins of a given trial was estimated from its spike train over that trial (see Methods). Next, population firing rate dynamics across multiple trials were used to train the SVM classifier (**Fig. 2.2A**). During testing, population activity from time bins of novel trials were presented to the trained SVM. SVM output for the population activity in a given time bin was represented by a vector of values generated by 25 readout units, where each readout value  $i$  represented a prediction score that the input pattern was from time bin  $i$ . This resulted in a vector of 25 readout values per test time bin (**Fig. 2.2B**). For each test time bin, the SVM predicted its bin label as the index of the maximal readout (**Fig. 2.2C**). Testing was performed with a Monte-Carlo cross-validation approach that controlled for the variance in the number of trials and size of the simultaneously recorded population across brain regions and animals (see Methods).

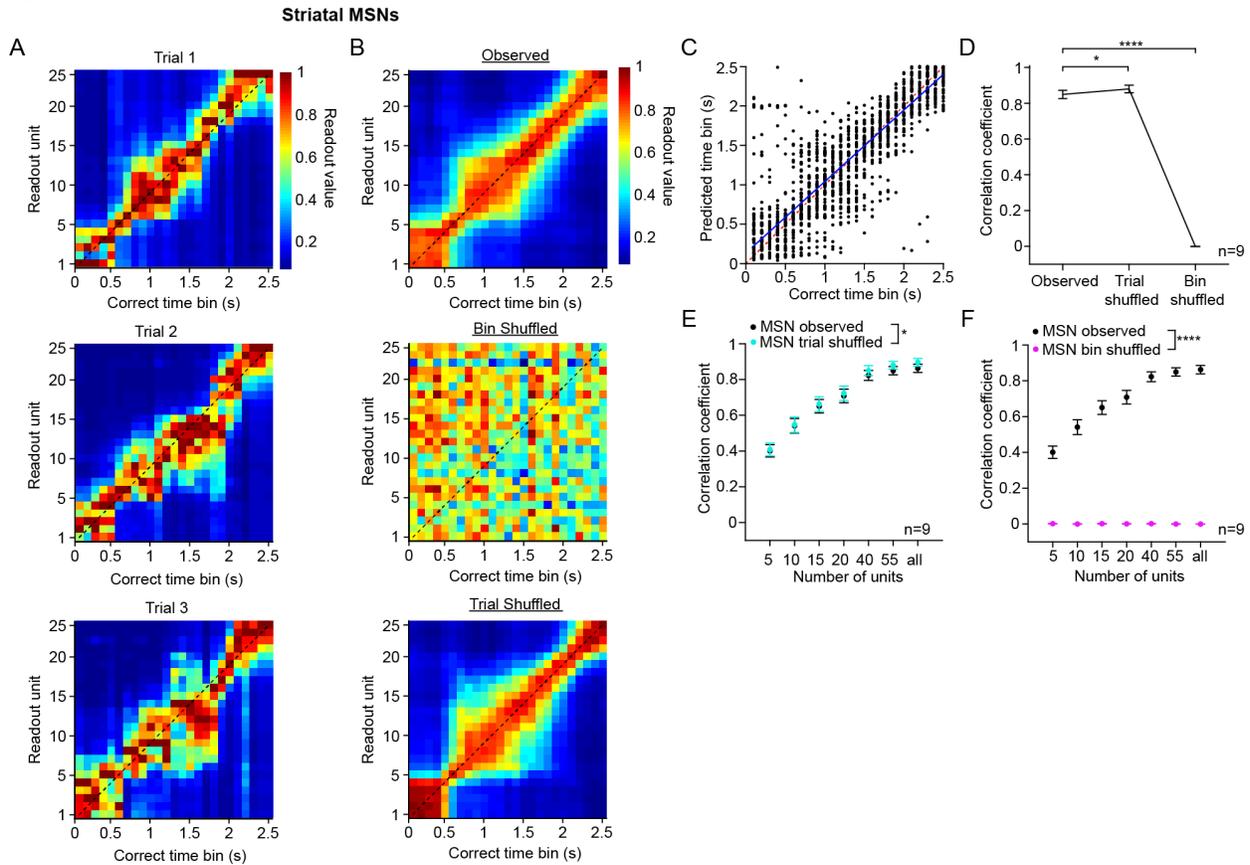
### *Elapsed time encoding by striatal and cortical networks*

We first investigated the ability of striatal MSN dynamics during single CS+ trials to be sorted into the correct temporal order by an SVM. Strong SVM performance would suggest that striatal neuron populations stably encode an internal representation of time elapsed from stimulus onset, and may provide a mechanism by which downstream regions could readout temporal information from striatal activity. We found that the highest SVM readout values during testing generally fell along the diagonal line in single trial cross-temporal classification matrices (**Fig. 2.3A**).

**Figure 2.3B**, *top* illustrates the average classification matrix over all trials in a single striatal recording and reveals the presence of a time code in the recorded dynamics. For each recording, we repeated the analysis on two different control patterns. First, to evaluate the temporal encoding efficacy of striatal population dynamics, we trained and tested an SVM on the dataset after scrambling its temporal dynamics by bin shuffling, wherein the sequence of firing activity for each unit within each trial was independently shuffled (see Methods). The control confirmed that bin shuffling completely eliminated the ability of the SVM to identify a code for elapsed time in the population activity (**Fig. 2.3B**, *center*). Since the above analysis was based on a set of 55 simultaneously recorded cells, we are able to determine the effect of noise correlations on the time code (Averbeck et al., 2006; Averbeck and D. Lee, 2006; Nirenberg et al., 2001; Schneidman et al., 2003)—in other words, does decoding based on simultaneously recorded cells hamper or improve performance. To do this, we measured decoding performance after independently shuffling the firing activity of each unit across trials (see Methods). While bin shuffling population activity rendered time bin predictions entirely random, trial shuffled controls performed very similarly to models trained on observed data (**Fig. 2.3B**, *bottom*).

In order to quantify performance, and the effects of bin and trial shuffling on the quality of the time code, we calculated the Pearson correlation coefficient between the correct and predicted time bin values in each recording (**Fig. 2.3C**). Across all striatal datasets, population

**Figure 2.3**



**Figure 2.3.** Striatal networks encode elapsed time. *A:* Cross-temporal classification matrices visualizing SVM model performance on striatal network data recorded during individual correctly performed CS+ trials. Each column represents the normalized readout values normalized across SVM readout units for the activity pattern from the corresponding correct time bin (X-axis). Peaks in each column reflect the predicted time chosen by the model. The black dotted line lies along the diagonal. *B: Top,* average of classification matrices generated across all correct CS+ trials for one striatal recording. *Center,* average classification matrix across all correct CS+ trials after bin shuffling each unit’s activity in the same recording. *Bottom,* average classification matrix across all correct CS+ trials after trial shuffling each unit’s activity in the same recording. *C:* Scatter plot of predicted versus correct time bins across 80 correctly performed CS+ trials for one striatal recording. Predicted bin numbers (Y-axis) were jittered (Gaussian noise, mean = 0, SD = 0.2) to separate overlapping points. The blue solid line represents the regressed line describing the correlation between actual and predicted time. The red dotted line lies along the identity line. *D:* Mean correlation coefficients between predicted and correct time bins across all striatal recordings (55 units per animal, n = 9) for observed, bin shuffled and trial shuffled data types. SVM classification of population activity was repeated 30 times (see Methods). SVM models trained on trial shuffled activity performed better than when trained on observed (non-shuffled) activity patterns (p = 0.023, paired t-test). Bin shuffled models performed at chance level, significantly worse than non-shuffled models (p < 0.0001, paired t-test).

**Figure 2.3 (continued).** *E*: Comparison of SVM performance using non-shuffled and trial shuffled network activity as a function of the number of units used for training and testing. There was a significant effect of data type ( $F_{1,8} = 7.9$ ,  $p = 0.023$ ), and number of units ( $F_{6,48} = 109.7$ ,  $p < 0.0001$ , two-way repeated-measures ANOVA). *F*: Bin shuffled models performed worse than non-shuffled models for each population size used in the model ( $F_{1,8} = 178.0$ ,  $p < 0.0001$ , two-way repeated-measures ANOVA). All error bars are SEM.

dynamics were highly predictive of elapsed time during the task (mean Pearson correlation coefficient = 0.85, SD = 0.069, n = 9). However, while bin shuffling reduced time prediction to chance levels, elapsed time decoding performance on the trial shuffled control was slightly, but significantly, better than on the observed data (mean Pearson correlation coefficient = 0.88, SD = 0.064,  $p = 0.024$ , paired t-test on Fisher transformed coefficients, **Fig. 2.3D**). This shows that the neurons are not noise independent and that the noise correlations—the within trial correlations between neurons—impair decoding.

Next, we examined if performance was dependent on the size of the striatal population used in decoding time. A two-way, repeated-measures ANOVA detected a significant effect of population size on classification performance ( $F_{6,48} = 109.7$ ,  $p < 0.0001$ ). The analysis again revealed a significant effect of trial shuffling ( $F_{1,8} = 7.9$ ,  $p = 0.023$ , **Fig. 2.3E**). These results show that under physiological conditions, striatal noise correlations are detrimental for neural coding of elapsed time, in agreement with the detrimental role of correlations found in other studies (Averbeck et al., 2006; Averbeck and D. Lee, 2006; Cohen and Maunsell, 2009; Mitchell et al., 2009; S. Tremblay et al., 2015). A separate two-way, repeated-measures ANOVA comparing observed and bin shuffled data at different population sizes also revealed a significant effect of bin-shuffling ( $F_{1,8} = 178.0$ ,  $p < 0.0001$ , **Fig. 2.3F**).

We next applied these same analyses to OFC pyramidal cell dynamics using the same procedures and numbers of cells. We found that OFC network dynamics also encoded elapsed time during the task (**Fig. 2.4A**). Interestingly, in contrast to the striatal code, we found no significant difference in the encoding efficacy between observed (mean Pearson correlation coefficient = 0.7, SD = 0.104, n = 6) and trial shuffled OFC network activity (mean Pearson correlation coefficient = 0.72, SD = 0.13,  $p = 0.21$ , paired t-test, **Fig. 2.4B**). While a two-way, repeated-measures ANOVA demonstrated that decoding performance using OFC population dynamics also depended on population size ( $F_{6,30} = 49.6$ ,  $p < 0.0001$ ), the analysis did not show a significant difference between the encoding efficacy of observed and trial shuffled data ( $F_{1,5} =$

2.4,  $p = 0.18$ , **Fig. 2.4C**). These findings suggest that temporal encoding in OFC dynamics is potentially less sensitive to noise correlations than in the striatum. Finally, as expected, bin shuffling the OFC dynamics eliminated the temporal code and resulted in chance level decoding performance (two-way, repeated-measures ANOVA,  $F_{1,5} = 109.5$ ,  $p < 0.0001$ , **Fig. 2.4D**).

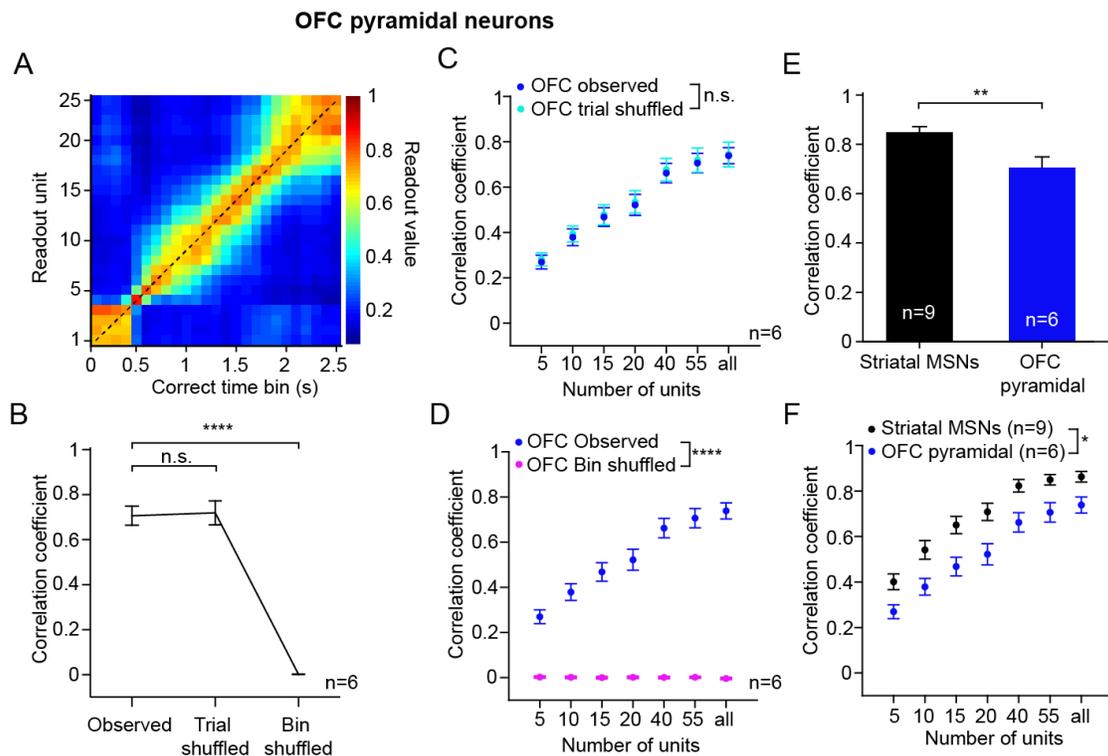
#### *Striatal networks outperform prefrontal networks in encoding elapsed time*

Consistent with the striatal results above, other experimental studies have reported the presence of a time code in the striatum (Gouvêa et al., 2015; Mello et al., 2015). We found that OFC networks also encode time, suggesting that this information is distributed throughout multiple brain areas. An important and unaddressed question pertains to the relative quality of this neural code in the striatum and OFC. We thus compared the performance of OFC and striatal network dynamics in encoding elapsed time. SVM classification performance was significantly better when trained and tested on striatal activity than on OFC activity ( $p = 0.0092$ , unpaired t-test, **Fig. 2.4E**). A two-way mixed-model ANOVA between brain region and population size revealed that this effect was consistent across a broad range of population sizes ( $F_{1,13} = 9.5$ ,  $p = 0.01$ , **Fig. 2.4F**). These results suggest that striatal networks show a significantly more robust representation of time compared to the OFC.

#### *Dorsal and ventral striatum equally encode elapsed time*

In the above analysis we adopted an unbiased approach for quantifying temporal coding in the striatum, in that we incorporated units from both anterior and posterior areas of this structure (**Fig. 5A**, *left and center*). Most of our recorded units were from the anterior striatum, but it is unclear to what extent this subregion by itself contained a better neural code than the OFC. We therefore repeated our comparative analysis after excluding posterior striatal MSNs (this reduced the minimum number of simultaneously recorded cells from 55 to 48). We found that the anterior striatum alone continued to have an improved code for time over the OFC (Pearson

**Figure 2.4**



**Figure 2.4.** Striatal networks encode elapsed time better than OFC networks. *A*: Average cross-temporal classification matrix across all correct CS+ trials for one OFC recording. Color scale is the same as in *Fig. 3B*. *B*: Mean correlation coefficients across all OFC recordings (55 units per animal,  $n = 6$ ) for observed, bin shuffled and trial shuffled data types. SVM classification of population activity was repeated 30 times (see Methods). SVM models trained on trial shuffled activity were not significantly different from those trained on non-shuffled activity patterns ( $p = 0.21$ , paired t-test). Bin shuffled models performed at chance level and significantly worse than the non-shuffled models ( $p < 0.0001$ , paired t-test). *C*: Comparison of SVM performance using non-shuffled and trial shuffled network activity as a function of the number of units. There was no significant effect of data type ( $F_{1,5} = 2.4$ ,  $p = 0.18$ ), but we observed a significant effect of the number of units ( $F_{6,30} = 49.6$ ,  $p < 0.0001$ , two-way repeated-measures ANOVA). *D*: Bin shuffled models performed worse than non-shuffled models for each population size used in the model ( $F_{1,5} = 109.5$ ,  $p = 0.0001$ , two-way repeated-measures ANOVA). *E*: Comparison of SVM model performance between all striatal and OFC recordings (55 units per region,  $n = 9$  striatal recordings and 6 OFC recordings) showed that the classification performance of models trained on striatal network data was significant better ( $p = 0.0092$ , unpaired t-test). *F*: Mean performance of SVM classification as a function of number of units used in training and testing for each brain region. A mixed-model ANOVA revealed a significant effect of number of units ( $F_{5,65} = 191.9$ ,  $p < 0.0001$ ) and a significant effect of brain region ( $F_{1,13} = 9.0$ ,  $p = 0.01$ ). The ANOVA excluded the ‘all units’ column as it contained inconsistent numbers of cells between regions. All error bars are SEM.

correlation coefficients: mean anterior striatum = 0.83, SD = 0.083, n = 9; mean OFC = 0.69, SD = 0.083, n = 6,  $p = 0.0083$ , unpaired t-test, **Fig. 2.5B**). Next we focused on differences between dorsal and ventral areas of the anterior striatum. Few studies have investigated whether the ventral striatum encodes time; however, the role of this area in reward prediction suggests that it may have a code for time. We took advantage of our widely distributed recording positions to compare the decoding performance of dorsal and ventral striatum MSNs (we used datasets with at least 35 simultaneously recorded MSNs). We found that SVM models trained on dorsal or ventral units performed as well as models trained with units taken randomly from either dorsal or ventral areas ( $F_{2, 16} = 0.02$ ,  $p = 0.98$ , one-way ANOVA, **Fig. 2.5C**). Together, these results suggest that the quality of temporal coding appears to be evenly distributed across the striatum, and that this area consistently outperforms the OFC.

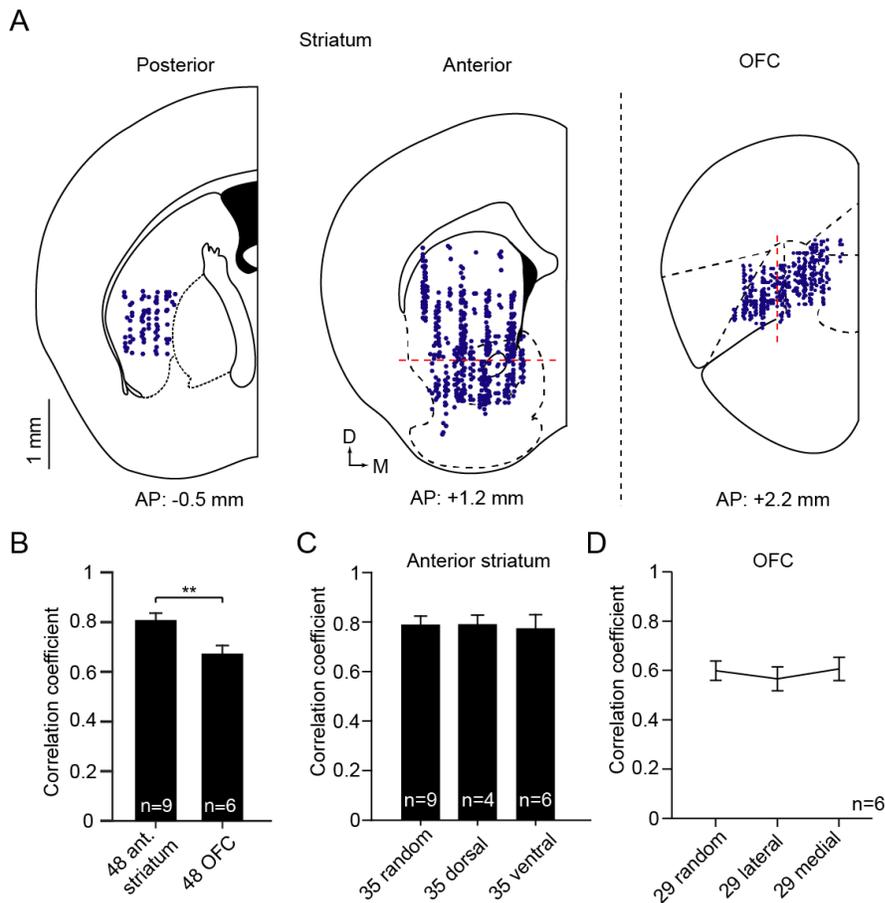
#### *Medial and lateral OFC equally encode elapsed time*

Our cortical recordings were mostly positioned within the OFC (**Fig 2.5A, right**). However, this area is comprised of several different anatomical subdivisions, raising the possibility that certain subregions encode time better than others. We therefore examined whether medial or lateral fields within our OFC recordings had a differential neural representation of time (we used datasets with at least 29 simultaneously recorded pyramidal cells). We found that models trained on medially or laterally positioned OFC units were just as effective at representing time as models using units taken randomly from either medial or lateral areas ( $F_{2,10} = 0.48$ ,  $p = 0.64$ , one-way, repeated-measures ANOVA, **Fig. 2.5D**). These findings suggest that the encoding of time via population dynamics is not localized to specific regions of the OFC.

#### *Lick-related movement does not explain the striatum's improved encoding of time*

Timing and movement are intimately related. Indeed in the current task licking should be driven in part by an internal representation of time, but on the other hand it is possible that some of the

**Figure 2.5**

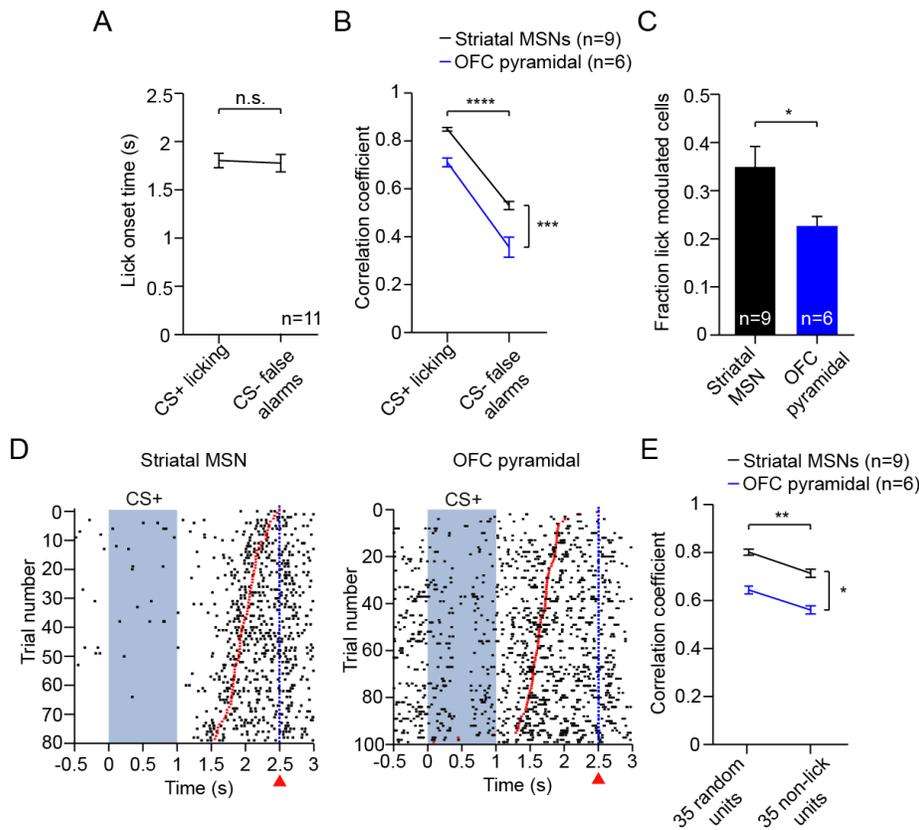


**Figure 2.5.** Population encoding of elapsed time is distributed throughout striatum and OFC. **A:** Illustrations of recording positions of all principal units included in analysis from posterior striatum (left), anterior striatum (center), and OFC (right). Dotted red lines indicate boundaries used to separate units recorded in dorsal and ventral striatum (center) or those recorded in lateral and medial OFC (right). Scale bar represents 1 mm. AP positions are distance from bregma. Section diagrams were adapted from Franklin and Paxinos (2008). **B:** Comparison of elapsed time decoding performance between models trained on recordings from OFC and anterior striatal neurons showed that anterior striatum performs better than OFC ( $p = 0.0083$ , unpaired t-test). **C:** Recordings in the anterior striatum were grouped based on whether they included predominantly dorsal or ventrally recorded neurons ( $N \geq 35$  cells), with one recording being distributed into both subregions. Dorsal and ventral populations performed as well as populations containing 35 cells drawn uniformly at random from both areas ( $F_{2,16} = 0.02$ ,  $p = 0.98$ , one-way ANOVA). **D:** All recordings in the OFC were bisected into lateral and medial populations. Lateral and medial populations performed as well as populations containing 29 cells drawn uniformly at random from both areas ( $F_{2,10} = 0.48$ ,  $p = 0.64$ , one-way repeated-measures ANOVA). All error bars are SEM.

code for time we observed might directly reflect neurons encoding motor behaviors. If the encoding or planning of motor activity were the primary basis for the observed code for time during reward-anticipatory licking after CS+ cues, then we would predict any licking episode would also encode time. We therefore examined whether population coding for time transferred to false alarm CS- trials, in which animals errantly licked following CS- odor presentations. Licking onset time was conserved between CS+ trials and false alarm CS- trials (mean CS+ lick onset time = 1.8 s, SD = 0.25 s; mean CS- lick onset time = 1.8 s, SD = 0.30 s;  $p = 0.80$ , paired t-test, **Fig. 2.6A**). To quantify the extent to which time-related coding could be detected during false alarm trials, we trained the SVM decoder on correct CS+ trials, and tested the model on the 0 to 2.5 s interval during false alarm trials. The performance of these models tested on false alarm trials was severely attenuated as compared to their performance when tested on correct CS+ trials (**Fig. 2.6B**). A two-way, mixed-model ANOVA revealed a significant effect of trial type used for testing ( $F_{1,13} = 33.0$ ,  $p < 0.0001$ ). The ANOVA did not show a significant interaction between brain region and trial type, demonstrating that both striatum and OFC saw an equal decrement in model performance when tested on CS- trials ( $F_{1,13} = 0.5$ ,  $p = 0.48$ ). These results suggest that temporal encoding is not only sensitive to licking, but also to other task variables such as the cue context.

To further examine the contribution of lick-related movement to our data, we identified principal cells that were positively correlated with lick rate. While we identified lick-rate-modulated cells in both areas, the striatum contained a significantly greater proportion of these cells than the OFC (mean striatal lick-modulated fraction = 0.35, SD = 0.127; mean OFC lick-modulated fraction = 0.226, SD = 0.049,  $p = 0.044$ , unpaired t-test, **Fig. 2.6C**). **Figure 2.6D** depicts two examples of lick-rate-modulated neurons from the striatum (*left*) and the OFC (*right*). We re-trained and tested the decoder after excluding these cells from the population. We found that removing lick-rate-modulated cells reduced decoder performance below what would be expected after removing the same number of randomly selected cells ( $F_{1,13} = 17.2$ ,  $p =$

**Figure 2.6**



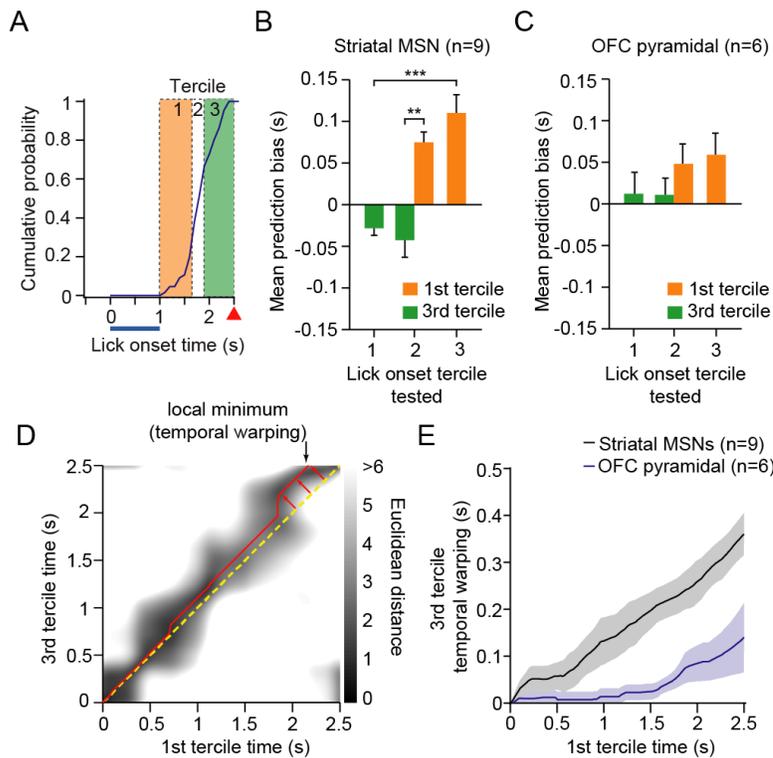
**Figure 2.6.** Population coding of elapsed time is specific to CS+ trials and not fully explained by licking behavior. **A:** Mice showed similar licking onset times during CS+ trials and CS- false alarm trials ( $p = 0.80$ , paired t-test). **B:** Comparison of performance in decoding elapsed time for SVM models trained on correct CS+ trials and tested on either correct CS+ trials or on CS- false alarm trials (55 units per region,  $n = 9$  striatal recordings and 6 OFC recordings). There was a significant effect of trial type ( $F_{1,13} = 33.0$ ,  $p < 0.0001$ , two-way, mixed-model ANOVA), and a significant effect of brain region ( $F_{1,13} = 18.3$ ,  $p = 0.00091$ ), with no significant interaction ( $F_{1,13} = 0.5$ ,  $p = 0.48$ ). **C:** Mean fraction of recorded principal cell populations showing significant activity modulation by licking in each brain region ( $p = 0.044$ , unpaired t-test). **D:** Example licking-modulated principal cells recorded in each region (left, Striatal MSN; right, OFC pyramidal). Shaded blue rectangle represents odor presentation time. Black tick marks indicate individual spikes, red ticks denote lick onset times, and blue dotted line shows reward delivery time. Trials are sorted by descending latency to first lick. **E:** Comparison of elapsed time decoding performance between models generated using all cells or all non-lick-modulated cells. Performance showed a significant decrease with the exclusion of lick-modulated cells ( $F_{1,13} = 17.2$ ,  $p = 0.0011$ , two-way, mixed-model ANOVA). The striatum maintained an improved code for time over the OFC after excluding lick-modulated cells ( $F_{1,13} = 7.4$ ,  $p = 0.017$ ). We did not observe a significant interaction between region and population ( $F_{1,13} = 0.9$ ,  $p = 0.35$ ). All error bars are SEM.

0.0011, two-way, mixed-model ANOVA, **Fig. 2.6E**). But crucially, the decoder still performed significantly above chance levels, demonstrating that a code for time was still present without lick-rate-modulated cells. In addition, we found that the striatum still performed better at representing time over the OFC despite the exclusion of lick-modulated cells ( $F_{1,13} = 7.4$ ,  $p = 0.017$ , two-way, mixed-model ANOVA). The ANOVA did not reveal a significant interaction between brain region and the type of population used in analysis ( $F_{1,13} = 0.936$ ,  $p = 0.35$ ). Together, the results in **Fig. 2.6** show that although movement does indeed contribute to the observed code for elapsed time in both the striatum and OFC, it is not sufficient to fully explain the neural representation of time in these areas. Furthermore, we demonstrated that our main finding that striatal ensembles outperform OFC ensembles in terms of temporal coding is robust even after controlling for lick-rate-modulated cells.

*The striatal population code for elapsed time co-varies with lick onset time*

Until now our decoding analysis was performed on all correct CS+ trials irrespective of the animal's actual lick onset time. However, since we found that lick-related movement partially contributed to the neural code for time, this implies that the neural code may vary on a trial-to-trial basis depending on the precise timing of lick onset. If the population dynamics are sensitive to lick onset time, then a prediction is that the encoding trajectories are respectively traversed faster (slower) when an animal licks earlier (later) than the mean. To test this prediction, we took advantage of the trial-to-trial variability in the time at which animals initiated licking during CS+ trials (**Fig. 2.1C**). We determined if population dynamics in the striatum and OFC reflected this variable lick onset time. For each animal, we divided trials into three evenly sized groups representing early (1<sup>st</sup> tercile), intermediate (2<sup>nd</sup> tercile), and late (3<sup>rd</sup> tercile) lick onset time trials (**Fig. 2.7A**). We then trained SVM models on trials in the 1<sup>st</sup> or 3<sup>rd</sup> terciles, and tested each separately on trials in the remaining terciles. In the striatal population, we found that when testing the 1<sup>st</sup> tercile's model versus the 3<sup>rd</sup> tercile's model on trials from the 2<sup>nd</sup> tercile, the

**Figure 2.7**



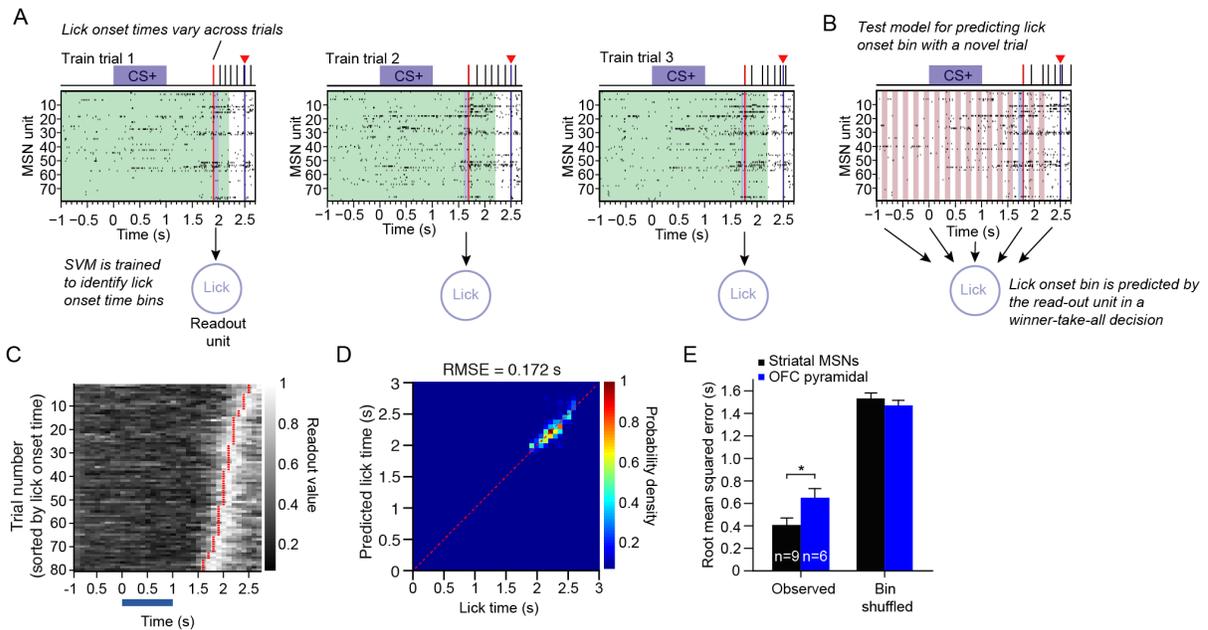
**Figure 2.7.** Striatal population coding of elapsed time shows higher sensitivity to lick onset variability than OFC. **A:** Schematic illustrating the division of correct CS+ trials into three sets based on tertiles of the lick onset distribution. **B:** Mean prediction biases of SVM decoders trained to predict elapsed time from striatal population data recorded in first tertile trials (orange), and tested on second and third tertile trials. Bias is defined as mean of predicted bin time minus correct bin time. Green bars show decoder biases when trained on third tertile trials and tested on data from first and second tertile trials. Training on first and third tertile trials and testing on second tertile trials produces opposing biases ( $p = 0.00034$ , paired t-test), as does training on first tertile trials and testing on third tertile trials when compared to training on third tertile trials and testing on first tertile trials ( $p = 0.002$ , paired t-test). **C:** Mean prediction biases of SVM decoders trained to predict elapsed time from OFC data under similar conditions as **B**. No significant difference in biases were observed when training on first and third tertile trials and testing on second tertile trials ( $p = 0.22$ , paired t-test), or when training on first tertile trials and testing on third tertile trials in comparison to training on third tertile trials and testing on first tertile trials ( $p = 0.06$ , paired t-test). **D:** Illustration of temporal alignment procedure on one striatal recording (88 cells). Distance matrix represents the Euclidean distance between all pairs of population activity patterns in the trial-averaged trajectories for the first and third tertile trials. Red line traces the minimum distance path along the distance matrix, between the beginning and the end of the mean first tertile trajectory. A deviation (red arrows) of this path from the diagonal (dashed yellow line) measures the temporal warping of the mean third tertile trajectory relative to the mean first tertile trajectory. The upward shift observed here indicates that the mean third tertile trajectory is consistently slower. **E:** Mean temporal warping of striatal (black) and orbitofrontal (blue) third tertile trajectories relative to their respective first tertile trajectories. All error bars are SEM.

evaluations showed opposing classification error biases ( $p = 0.00034$ , paired t-test, **Fig. 2.7B**). In other words, the model trained on the 1<sup>st</sup> tercile consistently classified time bins in the 2<sup>nd</sup> tercile as having occurred earlier than they had. Conversely, the model trained on the 3<sup>rd</sup> tercile consistently classified time bins in the 2<sup>nd</sup> tercile as having occurred later than they actually had. Furthermore, when testing the 1<sup>st</sup> tercile's model on the 3<sup>rd</sup> tercile's trials or testing the 3<sup>rd</sup> tercile's model on the 1<sup>st</sup> tercile's trials, these evaluations also showed opposing classification error biases ( $p = 0.002$ , paired t-test). Altogether, these results show that internal representation of time in the striatum appears to co-vary with the timing of lick onset, consistent with earlier work suggesting that the latency of the motor response was driven by the neural code for time (Gouvêa et al., 2015).

In contrast to the striatum, in the OFC, we did not find any significant effects of training classifiers on the 1<sup>st</sup> or 3<sup>rd</sup> terciles and testing those models on the 2<sup>nd</sup> tercile's trials ( $p = 0.22$ , paired t-test, **Fig. 2.7C**). Testing 1<sup>st</sup> or 3<sup>rd</sup> tercile classifiers on the 3<sup>rd</sup> or 1<sup>st</sup> terciles' trials, respectively, also did not result in biased classification error deviations ( $p = 0.06$ , paired t-test), although there was a trend. Thus in contrast to the striatum, the temporal code in the OFC may not co-vary as effectively with movement onset time.

It was possible that the decoded biases quantified above did not fully establish the extent of the underlying relationship between the internal representation of time and lick time, due to potential artifacts imposed by binning and smaller training datasets after grouping by terciles. To better determine the extent of temporal co-variation between licking and neural dynamics, we compared population trajectories averaged over the trials in the 1<sup>st</sup> tercile with population trajectories averaged over trials in the 3<sup>rd</sup> tercile. A temporal alignment procedure applied to the two trial-averaged trajectories (see Methods) revealed that while the two trajectories remained close to each other over the course of the trial interval, they were not uniformly aligned in time (**Fig. 2.7D**). Instead, the 3<sup>rd</sup> tercile trajectory consistently lagged behind the 1<sup>st</sup> tercile trajectory, illustrated in **Figure 2.7D** as an upward shift of the minimum distance

**Figure 2.8**



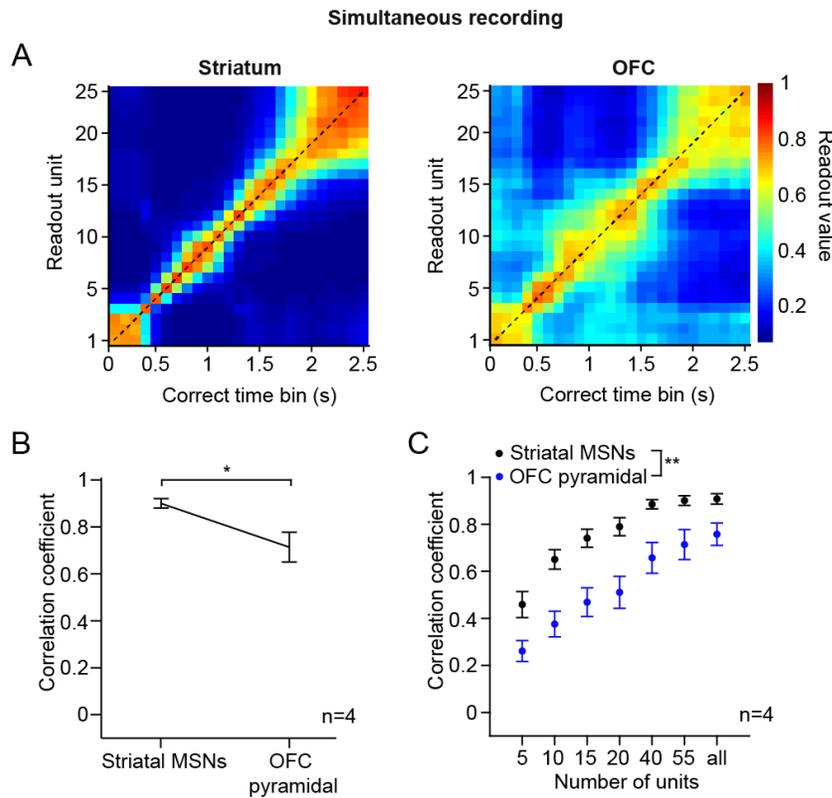
**Figure 2.8.** Striatal networks outperform OFC networks at predicting lick onset time. *A*: Illustration of lick onset time prediction analysis. Raster plots show the same MSN population’s activity during different correct CS+ trials. Upper schematic shows odor on time (blue rectangle), reward delivery (red triangle), and actual lick times (red/black lines) that correspond to the recorded raster plots. Each correctly performed CS+ trial has a lick onset time indicated by a red line. As in the elapsed time prediction analysis, in each trial, spiking activity of each unit was transformed into corresponding firing rate estimates (not shown), and the firing rates of simultaneously recorded units were binned (100 ms time bins) to construct population firing patterns for the trial. In each trial, the bin during which the first lick occurred is labeled as its lick onset bin (violet shading). A binary SVM classifier, represented here by a readout unit, was trained to distinguish between lick onset bins and non-lick onset bins (green shading). *B*: The model is tested using a Monte-Carlo cross-validation approach. Population activity patterns for all time bins in a trial are presented to the classifier, which predicts the lick onset bin for the trial as the time bin with the maximal readout value. *C*: Heatplot showing normalized trial-averaged readout values generated by the SVM trained and tested on striatal network activity of one mouse. Trials are sorted by decreasing latency to lick onset time, indicated by a red tick mark. *D*: 2D density plot showing the joint distribution of actual lick onset times and those predicted by the SVM from striatal network activity, for one mouse. Prediction performance is measured as the root mean squared error (RMSE). Lick onset bin classification was repeated 30 times for each trial (see Methods). Actual and predicted lick onset bins were jittered (Gaussian noise with 0 mean, 0.3 SD) to separate overlapping points. *E*: Comparison of mean predicted lick onset bin RMSEs across all striatal and OFC recordings (55 units per region,  $n = 9$  striatal recordings and 6 OFC recordings) showed that models trained on striatal network data performed significantly better ( $p = 0.032$ , unpaired t-test). Bin shuffled models based on striatal recordings performed significantly worse than corresponding non-shuffled models ( $p < 0.0001$ , paired t-test). Bin shuffled models based on OFC recordings also performed worse than corresponding non-shuffled models ( $p = 0.0002$ , paired t-test). All error bars are SEM.

curve between the two trajectories, away from the diagonal line. The magnitude of this shift is a measure of the temporal warping, or speed of progression, of the 3<sup>rd</sup> tercile trajectory with respect to the 1<sup>st</sup> tercile trajectory. In the striatum, temporal warping emerges very early on in the trial relative to lick onset timing (**Fig. 2.7E**), which suggests that the striatal activity encoding an internal representation of time undergoes “subjective” fluctuations that may drive trial-to-trial variability in lick onset. In contrast, warping was less prevalent, particularly near the beginning of the trial, in the OFC. Together these results suggest that the internal representation of time as encoded in the striatal dynamics, are co-modulated by the elapsed time and the lick onset time, and that these effects are less evident in the OFC.

#### *Striatal ensembles predict movement onset time*

Since striatal ensemble dynamics possessed a better code for time we hypothesized that the lick onset time could also be predicted with better accuracy from patterns of striatal activity than OFC activity. Using the ensemble firing rate pattern in each 100 ms time bin of a trial, a binary SVM classifier was trained to discriminate the population activity in the first time bin when an animal licked (*i.e.*, the lick onset bin) from the activity in all other time bins (**Fig. 2.8A**). SVM output for the population activity in a given bin is represented by a single readout unit whose value captures the propensity of lick onset occurring in that bin. To establish how well network activity predicted lick onset times, we used a Monte-Carlo cross-validation method to test trained SVM classifiers on population activity patterns in novel trials (**Fig. 2.8B**). The classifier generates one readout value for the activity pattern from each bin in a trial, and the predicted lick onset bin for the trial is chosen as the one with the maximal readout value. **Figure 2.8C** illustrates the readout value distributions decoded from the striatal dataset of an animal and its observed lick onset bins (red ticks) for all correct CS+ trials. To quantify the classification performance, we measured the root mean square error (RMSE) of the predicted lick onset times across all correct CS+ trials, as generated by the Monte-Carlo cross-validation approach (**Fig.**

**Figure 2.9**



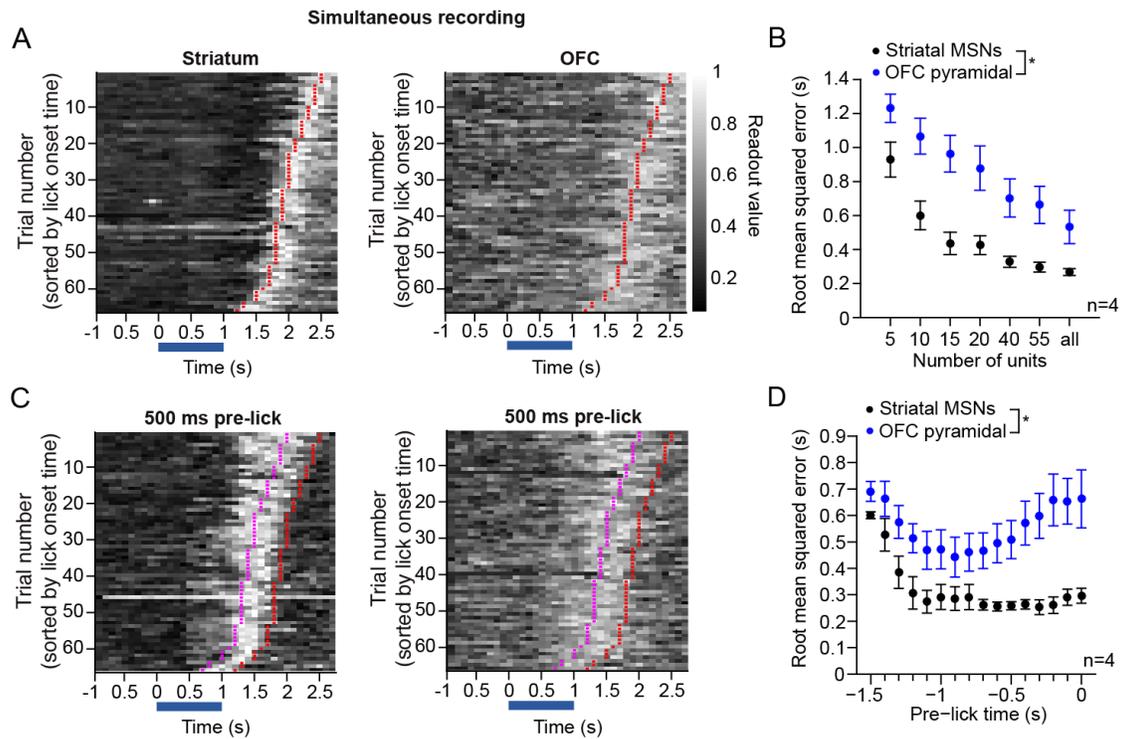
**Figure 2.9.** Simultaneous multi-region recordings indicate that striatum encodes elapsed time better than OFC. *A: Left*, average cross-temporal classification matrix showing mean performance of the elapsed time classifier across all correct CS+ trials for one striatal recording that occurred in parallel with a OFC recording in the same mouse. The classification matrix for the corresponding OFC recording is shown at *right*. *B*: Mean correlation coefficient across simultaneous striatal and OFC recordings (55 units per region,  $n = 4$ ) for each brain region. SVM classification of population activity was repeated 30 times (see Methods). SVM models trained on striatal activity performed better than when trained on OFC activity patterns ( $p = 0.013$ , paired t-test). *C*: Performance comparison of SVM models trained and tested on striatal and OFC network activity from simultaneous recordings, as a function of number of units. There was a significant effect of brain region ( $F_{1,3} = 58.1$ ,  $p = 0.0047$ ), and number of units ( $F_{5,15} = 73.4$ ,  $p < 0.0001$ , two-way repeated-measures ANOVA). All error bars are SEM.

**2.8D**). SVM models trained on the striatal network datasets (observed mean RMSE = 4.07, SD = 1.90; bin shuffled mean RMSE = 15.32, SD = 1.48,  $n = 9$ ,  $p < 0.0001$ , paired t-test, **Fig. 2.8E**) and the OFC network datasets (observed mean RMSE = 6.50, SD = 1.96; bin shuffled mean RMSE = 14.71, SD = 1.12,  $p = 0.0002$ ,  $n = 6$ , paired t-test, **Fig. 2.8E**) performed well above chance levels in predicting lick onset time. However, consistent with our hypothesis the SVM models trained on striatal activity outperformed those trained on OFC activity in predicting lick onset times during the task ( $p = 0.032$ , unpaired t-test, **Fig. 2.8E**).

*Simultaneous OFC and striatal recordings exhibit a superior code for elapsed time in the striatum*

The above analyses suggest that the network dynamics of the striatum constitute a better “clock” than the dynamics of the OFC. However, it is possible that these observations are partially influenced by differences in neural coding performance across animals. In a subset of our recordings ( $n = 4$ ) we were able to simultaneously measure at least 55 OFC pyramidal cells *and* 55 striatal MSNs within the same animal and session (Shobe et al., 2015). Thus, we examined if the observation that striatal dynamics contain a better code for elapsed time than the OFC was supported in these simultaneous dual region recordings. One particular advantage of this within-animal comparison is that the SVM models are trained and tested with network data in two brain regions that were recorded using identical behavioral conditions and trials. Thus, the networks share the same stimulus inputs, interval durations, and lick onset times. A cursory comparison of the cross-temporal classification matrices for simultaneously recorded brain regions in a single animal indicated that its striatal population encoded elapsed time more robustly than its OFC population (**Fig. 2.9A**). An accuracy comparison of the elapsed time decoded from population activity in the two brain regions, across all simultaneously recorded network activity datasets, reasserted that the striatal networks’ temporal encoding efficacy was consistently better (mean striatal correlation coefficient = 0.90, SD = 0.041; mean OFC

**Figure 2.10**



**Figure 2.10.** Simultaneous multi-region recordings show distinct pre-lick dynamics across striatal and OFC networks. **A:** Heatplots showing normalized readout values generated by SVM models trained to detect lick onset times. Heatplots reflect trial-averaged readout values of SVM models trained and tested on striatal (*left*) and OFC (*right*) network activity from simultaneous recordings from the same mouse (55 units per region,  $n = 4$ ). Trials are sorted by decreasing latency to lick onset time, indicated by a red tick mark. **B:** Mean performance of lick onset bin prediction as a function of number of units included in training and testing the SVM models, for each simultaneously recorded brain regions (55 units per region,  $n = 4$ ). A two-way, repeated-measures ANOVA revealed a significant effect of number of units ( $F_{5,15} = 178.4$ ,  $p < 0.0001$ ) and a significant effect of brain region ( $F_{1,3} = 18.9$ ,  $p = 0.022$ ). The ANOVA excluded the ‘all units’ column as it contained inconsistent numbers of cells between simultaneously recorded regions. **C:** Heatplots showing normalized readout values generated by SVM models trained to detect time bins occurring 500 ms prior to actual lick onset times. Heatplots reflect trial-averaged readout values of SVM models trained and tested on striatal (*left*) and OFC (*right*) network activity from simultaneous recordings from the same mouse. Trials are sorted by decreasing latency to actual lick onset time, indicated by a red tick mark. Magenta tick marks indicate 500 ms prior to lick onset (55 units per region,  $n = 4$ ). **D:** Mean RMSE values across all simultaneous striatal and OFC recordings (55 units per region,  $n = 4$ ) quantifying performance of SVM models trained and tested to predict time bins that occurred in advance of actual lick onset times. A two-way, repeated-measures ANOVA revealed a significant effect of time bin ( $F_{15,45} = 8.8$ ,  $p < 0.0001$ ) and brain region ( $F_{1,3} = 16.0$ ,  $p = 0.028$ ). All error bars are SEM.

correlation coefficient = 0.71, SD = 0.13, n = 4, p = 0.013, paired t-test, **Fig. 2.9B**). We also measured whether this effect was consistent at different sizes of the decoded neural population. A two-way, repeated-measures ANOVA between brain region and the decoded population size showed a significant effect of brain region ( $F_{1,3} = 58.1$ , p = 0.0047) and population size ( $F_{5,15} = 73.4$ , p < 0.0001, **Fig 2.9C**).

### *Simultaneous OFC and striatal recordings exhibit a superior lick onset time prediction in the striatum*

We also explored whether lick onset time prediction was significantly better using striatal population activity within the simultaneously recorded striatal and OFC data sets. We again observed that SVM models trained to identify population activity encoding lick onset time appeared to be more precise when decoding from striatal population activity than from OFC population activity (**Fig. 2.10A**). When comparing lick onset prediction performance between SVM models trained on simultaneously recorded networks as a function of the brain region and the size of the decoded population, we found that the striatal networks encode lick onset time with a significantly higher efficacy than OFC networks ( $F_{1,3} = 18.9$ , p = 0.022, two-way, repeated-measures ANOVA). There was also a highly significant effect of population size on lick onset prediction performance across the two brain regions ( $F_{5,15} = 178.4$ , p < 0.0001, two-way, repeated-measures ANOVA, **Fig. 2.10B**).

Lastly, we examined how far in advance could the neural activity predict lick onset time. This allowed us to further examine if lick onset prediction was not simply a product of neural activity directly driving motor responses. For this analysis we trained separate SVM classifiers on increasingly earlier target bins, moving the target bin backwards in time with respect to the actual lick onset bin (see Methods). In each brain region, a comparison of the readout value distributions when the classifier is trained on the actual lick onset bin (**Fig. 2.10A**) versus on the bin occurring 500 ms ahead of the actual lick onset bin (**Fig. 2.10C**) indicates that either time bin

can be decoded with similar reliability (**Fig. 2.10D**). A two-way, repeated-measures ANOVA revealed a significant effect of brain region ( $F_{1,3} = 16.0$ ,  $p = 0.028$ ). SVM models trained to decode pre-lick onset time bins from OFC population activity were consistently worse than those based on striatal population activity. There was also a significant effect of time ( $F_{15,45} = 8.8$ ,  $p < 0.0001$ ), whereby increasing the “look ahead” time resulted in a decay in classification performance at about 1 second. This effect was more pronounced in the striatum, which maintained a relatively constant RMSE until approximately 1 second prior to actual lick onset. The OFC on the other hand, showed higher error at actual lick onset as compared to the striatum. Interestingly, the OFC showed a slight improvement of classification further away in time from the actual lick onset. These results suggest that striatal and OFC dynamics are causally related to behavior, and that the activity patterns preceding the lick by up to 1 second encode when the animal will lick. Again the results show that striatal dynamics provide more information about lick onset time.

An important question related to both the nature of the time code and the mechanisms underlying the neural dynamics in the striatum and OFC relates to the “complexity” of the dynamics. One way to measure the complexity of patterns of neural activity is through its effective dimensionality—a PCA-based measure calculated as the number principal components required to explain 95% of the variance in the trial-averaged population dynamics (Rajan et al., 2016) (see Methods). The calculations show that the OFC dynamics (mean = 15.75, SD = 2.06) are significantly higher-dimensional than MSN dynamics (mean = 11.25, SD = 1.5,  $p < 0.0001$ , paired t-test), and suggest that the OFC may encode other variables that are not immediately relevant to the task.

## **Discussion for Chapter 2**

This study took advantage of the high single-unit recording throughput of silicon microprobes (Shobe et al., 2015) to examine the neural coding properties of large prefrontal and striatal

ensembles during a conditioning paradigm in which mice learned to anticipate the timing of reward. We used a machine-learning algorithm to quantify the ability of dynamically changing patterns of network activity to encode time at the single trial level within individual animals. Using activity patterns in these two brain regions, we decoded time elapsed from the onset of a reward-predictive cue. Relatedly we also could predict the time of anticipatory licking onset. Our results show that the striatum consistently outperformed the OFC in terms of the ability to encode time. This superior time representation of striatal ensembles was confirmed via simultaneous recordings in the OFC and striatum.

The large scale of the recordings (at least 55 simultaneously measured units per animal) enabled quantitative comparisons of decoding performance without the need for pooling units across sessions or subjects. While a few studies have recorded simultaneously in the cortex and striatum (D. Z. Jin et al., 2009; Matell et al., 2003), to our knowledge this is the first study to quantitatively compare the quality of the time code between the striatum and one area of the cortex, the OFC, while recording simultaneously from large populations of neurons. Therefore, we anticipate that the approaches outlined in this work offer new opportunities for understanding the coding properties of neural ensembles across multiple brain areas during behavior (E. N. Brown et al., 2004; Buzsáki, 2004).

Computational models (Buonomano and Laje, 2010; Laje and Buonomano, 2013; Medina et al., 2000) and recent experimental work (Carnevale et al., 2015; Crowe et al., 2014; Stokes et al., 2013) have suggested that motor timing may be encoded in dynamic patterns of neural activity: a “population clock”. *In vivo*, population clocks in the form of either simple sequential patterns of activity or complex high-dimensional patterns have now been observed in many different brain areas including the striatum (Bakhurin et al., 2016; Chiba et al., 2008; Gouvêa et al., 2015; D. Z. Jin et al., 2009; Matell et al., 2003; Mello et al., 2015) and multiple areas of the prefrontal cortex (Brody et al., 2003; Carnevale et al., 2015; Dietrich and Allen, 1998; Fuster, 2001; Genovesio et al., 2009; D. Z. Jin et al., 2009; Kim et al., 2013; Merchant et

al., 2011; Oshio et al., 2008; M. Xu et al., 2014). Our data are consistent with these previous reports of dynamic time-varying coding properties of neural ensembles, and support population clock models for the coding of time. This study provides evidence that different brain regions, here the striatum and the OFC, may be part of a distributed but regionally specialized network for encoding time.

The brain's code for time and the timing of movements are highly interdependent – animals use a timing mechanism to determine when to generate actions, but those actions in turn cause changes in brain activity, potentially influencing any observable code for time. We have addressed such interdependences in our data and show that while licking-related activity can contribute to the population clock, there is a robust code in a number of control analyses aimed at removing potential motor influences. Because we trained mice to learn to time a single interval, we exploited the fact that licking behavior demonstrated variable onset timing, similar to the timing variability of lever pressing during fixed-interval tasks (Matell et al., 2003; Mello et al., 2015). Our study benefited from this natural variability in that we could demonstrate that population codes reflected early or late onset times within the single interval. Interestingly, we found evidence that licking behavior that occurred particularly late or early strongly co-varies with the speed at which temporal codes evolved along the entire duration of the trial. This was particularly apparent in the striatum, which supported results in earlier work on population coding in the dorsal striatum (Gouvêa et al., 2015). Overall these analyses support the notion of high-dimensional multiplexed representations within the striatum and OFC (Fusi et al., 2016; Rigotti et al., 2013).

Our study focused on the OFC region of the prefrontal cortex, an area that has not received extensive attention with respect to encoding of time. Parts of the prefrontal cortex, including the dorsolateral, medial, and premotor areas have been previously shown to encode time (Crowe et al., 2014; Kim et al., 2009; Merchant et al., 2013b; 2015; Onoe et al., 2001). It is therefore possible that these other areas exhibit a better code for time than both the regions that

we investigated in the OFC and the striatum. Despite not having been extensively studied in the context of coding time per se, neurons in the OFC has been previously shown to be sensitive to time during reward expectation and other related behaviors such as temporal discounting, in which rewards received earlier in time are preferred to those associated with a greater delay period (Moorman and Aston-Jones, 2014; Roesch et al., 2006), but see (Jo et al., 2013). In this study, we now show that the OFC is capable of encoding time using population clocks. We also add to a growing literature that the OFC contains movement-related coding (Feierstein et al., 2006; Furuyashiki et al., 2008; Simon et al., 2015). A further implication of our data is that the OFC may be better suited for representing higher-dimensional information about the behavioral task. This was reflected in the effective dimensionality analysis and is consistent with the role of the OFC in complex cognitive processing (Fuster, 2001), which has been hypothesized to involve recurrent circuit activity (Mante et al., 2013; Rigotti et al., 2013).

A major finding of this study is that ensembles in the striatum outperform the OFC in terms of time encoding, even after controlling for motor effects. A fundamental question thus pertains to how these differences in coding arise. Is time encoding generated within the striatum, through a cortico-striatal loop (Merchant et al., 2013a), or is the striatum instantiating a readout of the dynamics generated within neocortical areas? These questions cannot yet be answered, but we suggest that the most consistent interpretation given the superior performance of the striatum observed here is that the striatum is serving as a read-out of the dynamics generated in the cortex. The striatum is innervated by a plethora of areas including prefrontal regions and motor areas (Hintiryan et al., 2016; McGeorge and Faull, 1989; Voorn et al., 2004). As a result of its functionally diverse inputs, the striatum as a whole integrates information that spans the cue-reward delay period, including sensory stimuli, reward prediction, and action initiation (X. Jin and Costa, 2010; Jog et al., 1999; Nicola et al., 2004; Roitman et al., 2005; Rueda-Orozco and Robbe, 2015; L. Tremblay et al., 1998). Our results therefore suggest an important role for sensorimotor integration in the striatum (Reig and Silberberg, 2014), which

may lead to a more refined representation of time in this structure than what was found in an upstream frontal cortical area.

We subdivided our dataset into dorsal and ventral striatal ensembles, and showed that each subregion separately performed as well in encoding time as when both subregions were pooled together. Most studies of timing in the striatum have focused on the dorsal striatum (Bartolo et al., 2014; Chiba et al., 2008; D. Z. Jin et al., 2009; Mello et al., 2015), with the exception of a study that investigated how dopaminergic signaling in dorsal and ventral areas contributed to timing behavior (Meck, 2006). Our results show that the ventral striatum encodes time as effectively as the dorsal striatum. We do not suggest that the ventral and dorsal areas are encoding the same kinds of information (Bakhurin et al., 2016), nor would they need to in order to represent time in their distinct patterns of dynamic activity. It is also possible that the synchronization of a temporal code could be attributed to local striatal microcircuitry (Bakhurin et al., 2016; Barbera et al., 2016) or due to basal ganglia feedback loops (Haber et al., 2000).

Several studies using fMRI approaches in humans have shown that dorsal areas of the striatum, including both the caudate and the putamen, are involved in interval timing tasks (Coull et al., 2011; Harrington et al., 2004; Wiener et al., 2010). It is important to point out that ventral striatal subregions were not often significantly modulated in these studies, which leads to a discrepancy between our findings and those performed in humans. This inconsistency could be explained by differences between the timing tasks performed by subjects across these studies. For example, successful performance in our task relies on animals anticipating the delivery of a reward, thus explaining why the ventral striatum—a region commonly implicated in reward processing—may be recruited. Tasks developed for humans may rely less on the simple Pavlovian associations that we used to train mice. In addition, fundamental differences between single-unit spike and BOLD signal measurements could make direct comparison of our data with fMRI experiments challenging.

The neuronal architecture of striatum differs dramatically from that of the neocortex: striatal circuits are characterized by recurrent inhibition (Tepper and Bolam, 2004), while neocortical circuits contain recurrent excitation. Theoretical studies have established that in contrast to inhibitory circuits, excitatory recurrent circuits are ideally suited to generate self-sustaining time-varying patterns of activity (H. Jaeger and Haas, 2004; Laje and Buonomano, 2013; Sussillo and Abbott, 2009)—although it is possible for such patterns to emerge from circuits that exhibit recurrent (feedback) inhibition (Mauk and Donegan, 1997; Medina et al., 2000). Because of the lateral inhibition interactions among MSNs (Taverna et al., 2008) and influence of local interneurons (Tepper et al., 2010), the striatal microcircuit may be well suited to refine those signals into an improved time code through a temporal “winners-take-all” mechanism by ensuring that the time-varying patterns of activity within cortical areas only activate a subpopulation of MSN cells at a time (Carrillo-Reid et al., 2011; Humphries et al., 2009; Ponzi and Wickens, 2012).

## **Chapter 3: External circuit influences on encoding of movement preparation in the striatum**

### *Introduction*

The basal ganglia are a set of subcortical nuclei that are believed to be involved in selecting and generating actions (Graybiel, 2008; Mink, 1996; Redgrave et al., 1999). The striatum is the major input nucleus of the basal ganglia (Bolam et al., 2000), which receives excitatory input from nearly every area of the neocortex (Finch, 1996; McGeorge and Faull, 1989). The striatal microcircuitry receives these glutamatergic inputs and transforms their signals into output signals that coordinate the activity of the basal ganglia output nuclei (Kreitzer and Malenka, 2008). Striatal medium spiny projection neurons (MSNs) are GABAergic and are also the predominant recipient of glutamatergic inputs on their dendritic spines. These neurons are typically highly hyperpolarized at rest and thus depend on coordinated glutamatergic input to fire (Calabresi et al., 1987; Plotkin et al., 2011; C. J. Wilson et al., 1990).

MSNs receive glutamatergic input from a diverse number of individual axons, and are thought to integrate across many cortical signals in order to produce their spiking activity (Carter et al., 2007; Kincaid et al., 1998; Kocsis et al., 1977; C. J. Wilson, 2013). In addition to receiving diverse glutamatergic signals, MSNs also produce widespread axon collaterals that form unidirectional inhibitory synapses onto other MSNs (Taverna et al., 2004; Tunstall et al., 2002; C. J. Wilson and Groves, 1980), and receive strong feedforward inhibition from local interneurons (Koós and Tepper, 1999; Taverna et al., 2007). While the physiological importance of these local inhibitory mechanisms are not well understood (Tepper et al., 2008), striatal output is believed to be the result of an interaction between glutamatergic input and local microcircuit interactions (Bolam et al., 2000; Kreitzer and Malenka, 2008). In behaving animals, striatal output signals can show highly complex relationships to external stimuli, context, reward history, and consequences of behavior (Lansink et al., 2012; Seo et al., 2012; Tai et al., 2012; L. Tremblay et al., 1998; A. Y. Wang et al., 2013). Furthermore, large networks of striatal MSNs

display dynamic trajectories that course through these populations during behavior (Bakhurin et al., 2016; Gage et al., 2010; Rueda-Orozco and Robbe, 2015; Thorn and Graybiel, 2014).

Several computational models have sought to understand the origin of these dynamic patterns and how they arise from known interactions in the striatal microcircuit (Connolly and Burns, 1993; Fukai and Tanaka, 1997; Humphries et al., 2006; Murray, 2017; Wickens et al., 1995). Several of these models demonstrate that complex network activity can be generated in the striatal microcircuit in the presence of tonic glutamatergic signals (Humphries et al., 2009; Ponzi and Wickens, 2010). These observations have been supported in slice experiments after bath application of NMDA (Carrillo-Reid et al., 2008). However, the extent to which striatal dynamic activity arises from local microcircuit interaction or is entirely dependent on the dynamic patterns of cortical inputs is not well understood.

Here we causally test the contribution that external glutamatergic signaling makes to the generation of striatal dynamic activity. We expressed inhibitory opsins in projection neurons of the premotor cortex in mice, a region that has been previously demonstrated to project to the sensorimotor striatum (Gremel and Costa, 2013; Li et al., 2016) and to be important for guiding licking behavior (Komiyama et al., 2010). We used large-scale electrophysiology in combination with optogenetic inhibition of projection terminals to study how eliminating M2 input would impact striatal dynamics related to cued licking behavior. Our results reveal that striatal populations integrate cortical inputs with other excitation sources in a complex manner to generate dynamic activity during behavior.

## **Materials and Methods**

### *Animals and surgical procedures*

All procedures were approved by the University of California, Los Angeles Chancellor's Animal Research Committee. Singly housed male C57Bl/6J mice (n= 31, 12-16 weeks old, The Jackson Laboratory) were used in the experiments. Animals underwent an initial surgery under

isoflurane anesthesia in a stereotaxic apparatus to bilaterally fix stainless steel head restraint bars (10 mm x 7.5 mm, 0.6 g) on the skull. We generated 3 different groups for this study, each of which required different surgery procedures. For behavioral testing, 9 animals received bilateral injections of AAV5-CAMKII-eNpHR3.0 into M2 (2.5 mm AP,  $\pm$ 1.5 mm ML, -1.3 mm and -1.1 mm DV, 450 nl total volume) and received bilateral fiber optic implants (200  $\mu$ m diameter fibers) over the injection site. An additional 8 mice received bilateral injections of AAV5-CAMKII-YFP to the same site along with bilateral fiber optic implants. For simultaneous recording in M2 and DLS, 5 mice only received headbars. 7 mice undergoing electrophysiological recordings with optogenetic manipulations received bilateral injections of AAV5-CAMKII-eNpHR3.0 (n = 3) or AAV5-CAMKII-ARCH3.0 (n = 4) into M2 but did not receive chronic optical fiber implants. For electrophysiological recordings, animals were anesthetized with isoflurane for a second surgery on the recording session day to make a craniotomy for acute microprobe insertion. Prior to the simultaneous M2 and DLS recordings, rectangular craniotomies (1.5 mm AP x 2 mm ML) were made in order to allow access to the following coordinates relative to bregma: For M2, we targeted AP, 2.5 mm, ML, 2.0 mm, and for DLS, we targeted AP, 1.0 mm, ML, 2.25 mm. For DLS-only recordings using the optomicroprobe, a smaller craniotomy was made (1.0 mm AP x 1.5 mm ML) centered over the following coordinates over the DLS: AP, 1.0 mm, ML 2.25 mm. In both cohorts, an additional craniotomy was made over the posterior cerebellum for placement of an electrical reference wire.

### *Behavioral task*

After recovery from the first surgery, animals were food restricted and fed daily after each training session to maintain ~90% of their baseline weight. They received water *ad libitum*. During daily training sessions, animals were mounted on the head bar bracket on the recording rig and stood while headfixed on a platform. Delivery of the reward solution (5  $\mu$ L, 10% sweetened condensed milk) was from a tube positioned between an infrared lick meter (Island

Motion), and was controlled by an audible solenoid valve actuation (Neptune Research). Animals were initially habituated to head fixation by receiving rewards alone (maximum 100 rewards per daily session, 13-21 s inter-trial interval, ITI, sampled from a normal distribution), and exposed to a constant flow of odorless air (1.5 L/min) through a tube. After animals successfully consumed 90% of delivered rewards for two consecutive days, they progressed to the Pavlovian conditioning using an odor delivered via an olfactometer. The odor cue was introduced by bubbling air (0.15 L/min) through an aromatic liquid (isoamyl acetate) diluted 1:10 in mineral oil (Sigma-Aldrich), and mixing this product with the 1.5 L/min stream of air. The odor cue was presented for 1 second, was followed by a 2 s delay period, and the subsequent delivery of the reward 3 s after odor onset. After repeated presentation of the odor-reward pairing, animals generate anticipatory licking behavior. Performance on the task was quantified by counting the fraction of trials containing at least one lick prior to reward delivery, called a “hit” trial. After mice demonstrated greater than 90% successful hit trials, they underwent craniotomies for recording or began optogenetic manipulation testing.

#### *In vivo optical stimulation*

All stimulation experiments commenced 4 weeks following viral injection to allow for sufficient expression of inhibitory opsins in M2 projections. For bilateral M2 inhibition cohort, mice first received 40 trials without laser (“Pre”), then 40 trials with laser presented on each trial (“Laser”), followed by another 40 trials without laser to test for behavioral recovery (“Post”). During the laser on block, the laser was activated continuously beginning 2 seconds prior to cue onset and terminated alongside reward delivery. For recording combined with laser presentation, we presented laser for short or long durations within the trial. Individual animals received different combinations of short and long trials. 5 mice received long laser trials randomly on 50% of the trials. The laser was activated continuously beginning 2 seconds prior to the cue onset and terminated 3 seconds after reward delivery for a total period of 8 seconds. 2 mice received short

and long laser trials randomly interleaved with trials without laser presentation. These mice received laser presentations at cue onset that had a duration of 0.2 s (20% of trials), presentations at -0.8 s prior to cue onset that had a duration of 1 s (20% of trials), presentations at -2 s prior to cue onset that had a duration of 8 s (20% of trials) and 40% of trials that contained no laser. Only trials with 8 s laser were used for analysis here. All mice received 20 trials of laser presentation at the end of the recording to test for laser effects on spontaneous activity.

### *Electrophysiological recordings*

All recordings were performed using silicon microprobes (Shobe et al., 2015). Each device contained 256 recording channels distributed over 4 prongs (64 channels per prong in a honeycomb array pattern spanning 1.05 mm). For simultaneous M2-DLS recordings, two separate probes were combined to form a 512-channel device. The M2 layer had prongs spaced by 400  $\mu\text{m}$  and the DLS layer had prongs spaced by 200  $\mu\text{m}$ . For the cohort that underwent electrophysiological recording and M2 terminal suppression, we used a device with prongs spaced by 200  $\mu\text{m}$  that was integrated with two optical fibers (200  $\mu\text{m}$  diameter) with their centers spaced 400  $\mu\text{m}$  apart. The optomicroprobe device was cleaned after each recording session in a trypsin solution and deionized water, and could be reused. Laser power coming from the integrated optical fibers was calibrated prior to each experiment and was tested after completing the recording to ensure stability. On the recording day, animals were anesthetized and underwent a brief surgery to make a craniotomy over the recording sites (described above). During the 6-hour recover period, craniotomies were sealed with a silicone compound (Kwik-Cast, World Precision Instruments). After recovery, animals were installed head-fixed on the platform, the silicone compound was removed, and a reference wire was placed on the surface of the cerebellum. The microprobes were inserted into the brain under the control of a motorized micromanipulator. For simultaneous M2 and DLS recordings, the position

of the most lateral shaft in M2 was AP +2.5 mm, ML +2.5, and DV -2.0 relative to bregma. The recording position of the most lateral shaft position in DLS was AP +1.0, ML +2.5, and DV - 4.0. For DLS-only recordings, the position of the most lateral shaft position in DLS was AP +1.0, ML +2.5, and DV - 4.0. Mineral oil was applied to the craniotomy to prevent drying. Data collection began after a 45 minute settling period after reaching the final recording depth, using custom built acquisition software at a sampling rate of 25 kHz per electrode (Shobe et al., 2015).

### *Spike sorting and unit classification*

Spike sorting and all analysis of neural data were performed with custom MATLAB scripts (Mathworks, Natick MA), (Shobe et al., 2015). Striatal and cortical units were classified using previously described criteria (Bakhurin et al., 2017) to identify putative medium spiny neurons (MSNs) and putative cortical pyramidal cells. These principal cell classes were distinguished from putative interneuron populations on the basis of their spike waveform peak-to-trough width and coefficient of variation of baseline firing rate. Putative MSNs and pyramidal cells had wider spike waveforms (minimum width = 0.55 ms, maximum width = 1.25 ms) than interneuron populations. Putative tonically active neurons (TANs) were segregated from MSNs in striatal recordings by the regularity of their baseline firing (maximum coefficient of variation = 1.5). No analysis was performed on putative interneuron populations.

### *Histology*

We confirmed expression levels of eNpHR-YFP or YFP bilaterally in M2 and in DLS projections using immunolabelling of GFP. After recording, animals were sacrificed by overdose of Isoflurane, were transcardially perfused with 10% formalin, and brains were harvested and immersed in 10% formalin solution overnight. Brains were sectioned coronally at 100  $\mu$ m thickness on a vibratome. Sections were incubated overnight with chicken primary antibodies to GFP (1:1000, Abcam) at 4°C. The next day, sections were washed in PBS and were incubated

in Alexafluor 488-conjugated donkey antibodies to chicken (1:200, Jackson Immunoresearch). Sections were imaged under an epifluorescence microscope to confirm viral expression.

### *Firing rate analysis*

Firing rates were calculated using 50 ms bins and were aligned to either cue onset times, lick onset times, or to laser onset times. To compare firing rates for each unit during laser-off ( $R_{\text{OFF}}$ ) and laser-on ( $R_{\text{ON}}$ ) trials, we aligned the firing rate to the onset of the first lick, and averaged over all hit trials. Firing rate differences ( $R_{\text{OFF}} - R_{\text{ON}}$ ) were calculated for each unit individually before averaging across all recorded units. To identify cells that were significantly modulated by the task, we aligned their firing rates to the onset of the odor cue and compared the firing rate distribution across all trials for each time bin (from 0 to 5 s after cue onset) with the distribution of baseline firing rates estimated for the 5 second period prior to cue onset for each trial. A permutation test was used to establish the probability that a difference between mean baseline mean and the mean rate for each time bin could be observed by chance. The bin was determined to be significant if that probability was less than 1%. At least two consecutive time bins had to pass the significance test over in order for the cell to be considered significantly modulated by the task.

### *Spike-LFP coherence*

LFP signals were extracted from the raw data recorded at every channel and downsampling to 1000 Hz. To calculate spike-LFP coherence between simultaneously recorded M2 and DLS activity, we aligned spiking activity from all M2 units and LFP signals sampled from electrodes in the DLS to the onset of each anticipatory licking bout for all hit trials, analyzing a window of -3 s to + 3 s from lick onset. Spike trains for each M2 unit were convolved using a Gaussian kernel (kernel width = 50 ms, SD = 5 ms). To calculate the coherence, we used the Matlab function *mscohere*. Coherence between 0 and 500 Hz was calculated for pairs of M2 units and 4 DLS

electrodes for each trial individually, and was then averaged to create the final coherogram. Only frequencies between 0.25 Hz and 60 Hz were considered for final analysis. Up to 200 pairs of units and electrodes were analyzed per animal.

#### *Time of maximum rate difference*

We identified the specific time of maximum firing rate effect we calculated the absolute value of  $R_{OFF} - R_{ON}$  for each DLS MSN and identified the time bin where this curve reached its maximum. To estimate the null uniform distribution, we sampled the same number of random times relative to lick onset as there were recorded DLS MSNs. To identify the fraction of MSNs either excited or inhibited by the laser, we performed a paired t-test between the distribution of firing rates across all  $R_{OFF}$  and  $R_{ON}$  trials for each time bin relative to licking onset. If the t-test showed that there was a significant difference between these distributions, the unit was classified as being excited if  $R_{OFF}(i) - R_{ON}(i) < 0$  and inhibited if  $R_{OFF}(i) - R_{ON}(i) > 0$ , where  $i$  is the time bin index. All cells were tested for each time bin. To determine the magnitude of the effect of laser as a function of peak firing rate time, we calculated  $R_{OFF}(m) - R_{ON}(m)$  for each unit, where  $m$  is the index of that unit's maximum firing rate during  $R_{OFF}$ . The response to laser was classified as excited if  $R_{OFF}(m) - R_{ON}(m) < 0$  and inhibited if  $R_{OFF}(m) - R_{ON}(m) > 0$ . The mean value for each response class was determined across all units achieving maximum firing in a given bin.

#### *Clustering*

K-means clustering of MSN responses was performed on lick-aligned mean firing rates that were calculated across all hit trials, base-line subtracted, and normalized to the peak firing rate magnitude. The Matlab function *kmeans* used Euclidean distance to iteratively assign individual MSN rates to randomly selected centroids. K was set to 3.

### *Regression analysis*

Univariate linear regression was used to measure linear relationships between  $R_{ON}$  and  $R_{OFF}$  for each neuron. The MATLAB function *regress* was used to identify the  $R^2$  value, using the least-squares approach, between the distribution of firing rates observed in  $R_{OFF}$  and the firing rates in the corresponding time bins during  $R_{ON}$ .

### *Lick onset prediction*

Lick onset time prediction using population codes was based on the approach taken in (Bakurin et al., 2017). Analysis was performed individually on each simultaneously recorded population. The decoder was trained on either laser-off trials or laser-on trials, and was subsequently tested on the same kind of trial (within-trial), or the other kind of trial (cross-trial). For each hit trial, firing rates were estimated for each MSN in the population in 100 ms time bins by (1) convolving the spike train during each trial with a decaying exponential function ( $\tau = 100$  ms) and (2) taking the mean of the convolved trace for each time bin. This resulted in a multidimensional trajectory of the neural population beginning 1 second prior to the onset of the cue and terminating 2 time bins after the maximum anticipatory lick onset time for that animal. This resulted in a variable number of time bins analyzed for each animal, varying between 28 and 42 time bins.

We trained a radial basis function support vector machine (SVM) learning algorithm to identify the specific population activity that corresponded to the time bin of lick onset. The SVM contained a single read-out that was trained to distinguish between time bins that contained the initial lick onset time and time bins that did not contain lick onsets. This can be conceptualized as the algorithm having a single readout unit that learns to detect population activity that corresponds to movement initiation. The predicted lick onset time bin is the one with the highest read-out score after the read-out unit is presented with every time bin in the test trial. As there is only one lick onset time bin per trial, the dataset contains a disproportionate number of non-lick

onset time bins. For this reason, we oversampled the lick onset time bin by including time bins immediately preceding and following the true lick onset bin. We additionally undersampled non-lick onset time bins by randomly selecting 33% of them as training examples for each training trial. In order to control for variable numbers of trials and cells recorded for each subject, we employed a Monte-Carlo subsampling procedure. Equal numbers of trials and cells were used for each subject. For each subject, this procedure performed analysis on 42 randomly selected trials out of all available trials and 20 randomly selected MSNs from the available population. Analysis was repeated 30 times in order to account for any variability resulting from sampling suboptimal trials and cells. Performance of the decoder was quantified as the root mean squared error (RMSE) between predicted and actual lick onset times calculated across all 30 runs of the algorithm.

For within-trial decoding, the SVM was implemented using a leave-one-out cross validation procedure in which the algorithm was trained on all trials except one that was left out for testing. This procedure iterated such that each trial was tested exactly one time. For cross-trial decoding, the SVM was trained using all trials of one trial type. Each trial of the other trial type was then tested on the model.

The misclassification cost ( $C$ ) and data complexity ( $\gamma$ ) regularization parameters for the RBF SVMs were optimized for each laser-off and laser-on dataset using a 5-fold cross validation procedure. Across all datasets, the values of  $C$  ranged from 0.0625 to 256 and the values of  $\gamma$  ranged from 0 to 0.25. The same  $C$  and  $\gamma$  parameters were used in cross-trial, within-trial, and bin shuffled analyses for each subject.

### *Bin shuffling*

Bin shuffling was used to generate population responses that were dissociated from their correct temporal order. To create bin shuffled activity, each unit's firing rate estimate in each time bin was replaced with the same unit's firing rate estimate in a randomly selected bin of the

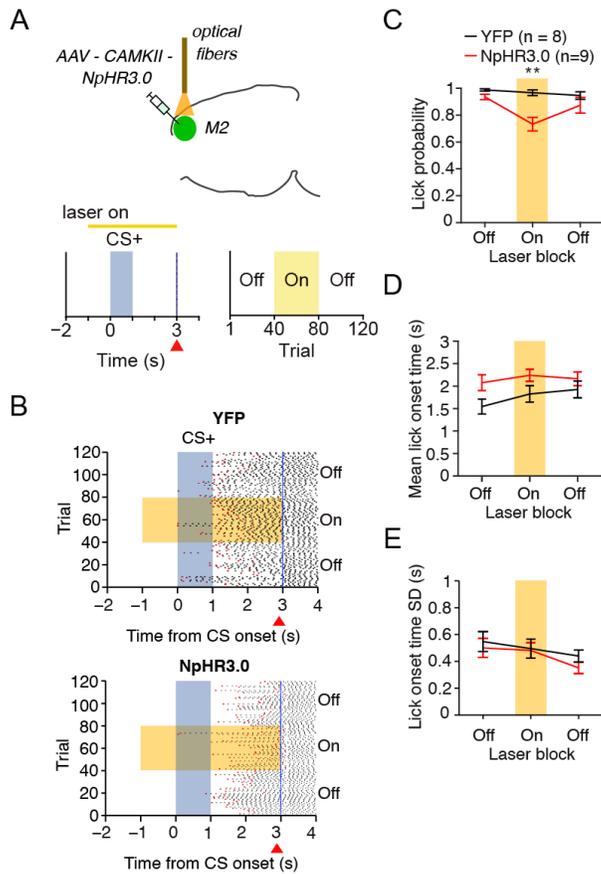
same trial. Prediction analyses from trial and bin shuffled data involved training and then testing on the respectively transformed datasets. To allow for direct comparisons between observed, trial shuffled, and bin shuffled controls, care was taken to make sure that we subsampled the same units and trials for analysis. Random performance was estimated by taking the mean of classifier performance trained and tested on bin shuffled laser-off and laser-on data across all subjects.

## Results

### *M2 inactivation reduces anticipatory licking probability*

We first sought to establish that the premotor cortex was involved in the expression of anticipatory licking behavior in the Pavlovian association task (Bakhurin et al., 2016; Shobe et al., 2015). Prior to training, we infected excitatory neurons in M2 with transgenes that promote expression of inhibitory halorhodopsin (eNpHR) under the control of the CAMKII promoter (**Fig. 3.1A**). Once animals generated anticipatory licking behavior on 90% of trials, we applied laser inhibition to the premotor cortex. Mice first received 40 trials without laser to establish their baseline performance. We then activated the laser for the duration of the pre-reward phase of the trial (-1 s to +3 s relative to odor cue onset) for 40 trials. This was followed by another 40 trials without laser to measure recovery from the manipulation. We found that mice showed a reduced probability to initiate licking behavior during the trial block containing laser presentations compared to the control group that received AAV-GFP virus injections (**Fig. 3.1B**). A 2-way, mixed model ANOVA revealed a significant effect of group on hit rate ( $F = 8.49$ ,  $p = 0.01$ , **Fig. 3.1C**). Interestingly, there was no significant effect of group in the 2-way, mixed model ANOVA performed on mean lick onset time ( $F = 2.97$ ,  $p = 0.11$ , **Fig. 3.1D**), or the standard deviation of lick onset time ( $F = 0.45$ ,  $p = 0.51$ , **Fig. 3.1E**). Furthermore, these manipulations did not affect the animals' ability to consume reward (data not shown). This suggests that M2 plays an important role in generating anticipatory licking behavior in our task, but is not involved in

**Figure 3.1**



**Figure 3.1.** M2 inhibition reduces anticipatory licking probability. **A:** Experiment schemas. AAV-mediated expression of eNpHR in excitatory neuron populations in M2 carried out bilaterally. 200  $\mu$ m optical fibers were implanted over injection sites. After reaching well-trained status, mice underwent optical stimulation. Laser (yellow bar) was delivered during the pre-reward phase of the trial beginning 1 s prior to cue onset. Laser delivery was performed in a blocked fashion where baseline performance was established, laser presented, and then laser was withheld to confirm behavioral recovery. Blue rectangle shows cue onset time and yellow rectangle the laser-on block. Red triangle and blue vertical line indicate reward delivery. **B:** Example licking raster plots for control (top) and experimental animals (bottom) showing reduction in anticipatory licking in the eNpHR group. Laser-off and laser-on blocks are indicated on the right. Coloring conventions are the same as in **A**. Red tick-marks indicate licking onset time. **C:** Probability of anticipatory licking for control and experimental groups as a function of laser-stimulation block. A two-way mixed-measures ANOVA revealed a significant effect of group ( $p = 0.01$ ,  $F_{1,15} = 8.494$ ). **D:** Lick onset time for control and experimental groups as a function of laser-stimulation block. There was no significant effect of group in a mixed-model two-way ANOVA ( $p = 0.105$ ). **E:** Lick onset time standard deviation for control and experimental groups as a function of laser-stimulation block. There was no significant effect of group in a mixed-model two-way ANOVA ( $p = 0.5$ ). All error bars are SEM.

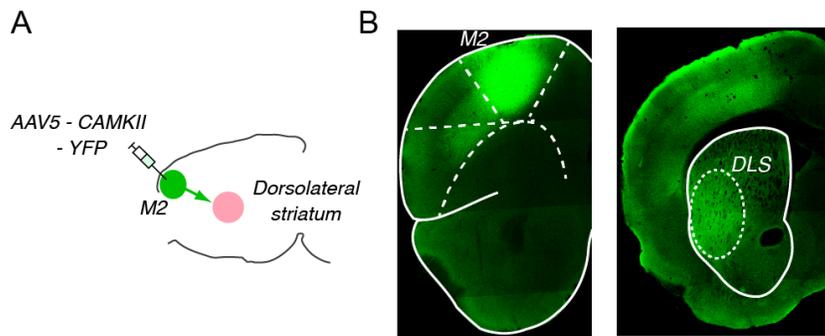
coordinating the specific parameters of these anticipatory movements or in actual reward consumption.

#### *Synchrony between M2 and DLS during cued action preparation*

We next determined whether M2 interacts with the DLS to drive anticipatory licking behavior. We first injected M2 with CAMKII-YFP virus to trace projection targets of the premotor cortex (**Fig. 3.2A**). We found robust labeling of terminals in the DLS, indicating that M2 sends projections to the striatum (**Fig. 3.2B**). This was consistent with previous reports of the existence of such connections (Gremel and Costa, 2013; Li et al., 2016). We wished to investigate the physiological activity of cell populations in both areas in order to further characterize the potential interaction between these two regions. After reaching high performance in the Pavlovian association task, wildtype mice underwent surgery to create craniotomies over M2 and DLS. After recovery, mice were implanted with large-scale, silicon based microelectrode arrays that targeted both areas simultaneously (Shobe et al., 2015), (**Fig. 3.3A**). We found that both M2 and the DLS generated dynamic population activity during the task, with 78% (352/451 units) of recorded M2 pyramidal neurons and 59% (191/322 units) of recorded DLS MSNs showing significant modulation during the task. In addition, both areas contained neurons that activated in the times prior to licking initiation (M2 pyramidal neurons: 53%, 238/451; DLS MSNs: 53%, 172/322, **Fig. 3.3B**). When we aligned the population activity to lick onset across all animals, we found that M2 and DLS showed sequential activity dynamics that began prior to movement generation (**Figs. 3.3C, 3.3D**). These results are consistent with both regions being involved in preparation of movement.

To investigate the possibility that these two brain regions are functionally coupled in the times prior to lick initiation, we estimated the spike-LFP coherence between these two brain regions. We found that spike-LFP coherence was elevated between these two brain areas in the times prior to movement initiation (**Fig. 3.4A**), specifically with the beta frequency band (15 Hz –

**Figure 3.2**



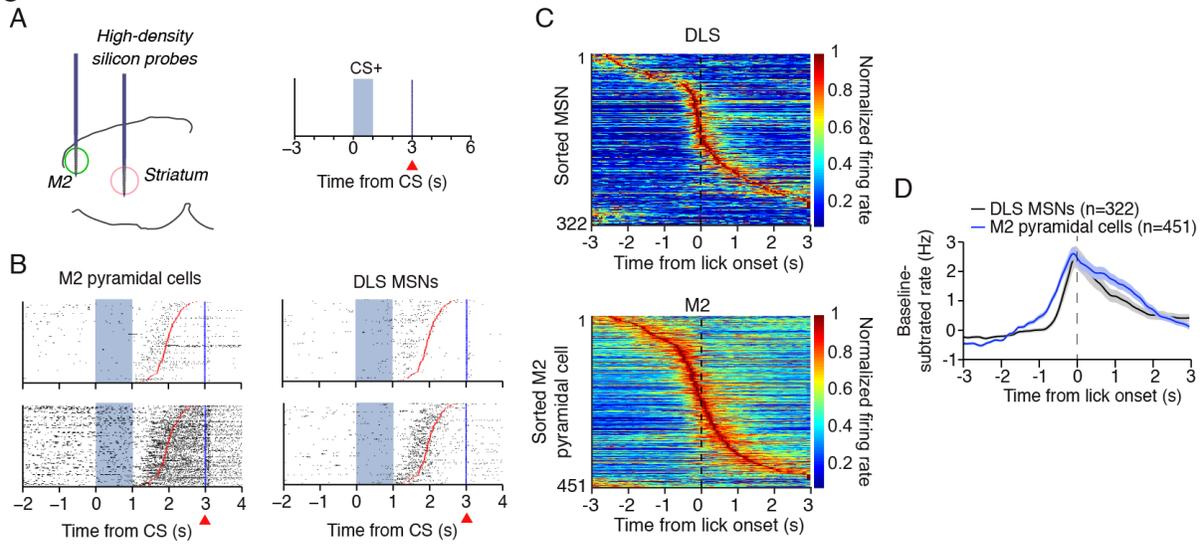
**Figure 3.2.** M2 projects to dorsolateral striatum. *A:* Experiment schema. AAV-mediated expression of YFP in excitatory neuron populations of M2. *B:* Fluorescent imaging revealed robust YFP expression in M2 and in DLS.

30 Hz). This suggests that these regions show synchronous activity during movement preparation periods, which is thought to reflect inter-area functional interactions. We performed the same analyses on the initiation of spontaneous licking bouts that occurred outside of the task. These events showed reduced firing rates in the two populations, and did not show changes in spike-LFP coherence in the times prior to their generation (**Fig. 3.4B**). When comparing spike-LFP coherence between cued and spontaneous licking, cued licking showed increased coherence within the beta-frequency band (**Fig. 3.4C**). This suggests that M2 and DLS interact prior to movement generation and that this interaction may be specific to movements that are linked to environmental stimuli through learning.

#### *M2 provides a dynamic drive onto DLS neurons*

We next used a causal approach to test the contribution that M2 makes onto DLS dynamic activity during the preparatory phases of movement. We injected viruses bilaterally into M2 to express the inhibitory opsins NpHR and Arch in excitatory neurons under the control of the CAMKII promoter (**Fig. 3.5A, left**). To observe the effects of M2 input suppression on DLS population activity, we used an integrated large-scale recording microprobe and optical light delivery device. This was implanted in the DLS in the region that receives M2 projections. We presented laser for the entire duration of the trial (-2 s to +6 s relative to cue onset) to maximize the effect of the laser on DLS activity (**Fig. 3.5A, middle**). We studied the effects of laser on MSN activity by calculating average firing rates during laser-off trials and laser-on trials (**Fig. 3.5A, right**). We first observed that DLS MSNs generated dynamic activity when aligning the population activity during laser-off trials to the start of anticipatory licking (**Fig. 3.5B, left**). Next, on laser-on trials we observed that suppressing M2 input dramatically reduced firing rates of many MSNs during the task (**Fig. 3.5B, right**). This resulted in a reduced mean firing rate across the recorded MSN population (**Fig. 3.5C**). We first calculated the difference between the mean firing rates during laser-off trials and laser-on trials for each recorded MSN (**Fig. 3.5D**).

**Figure 3.3**



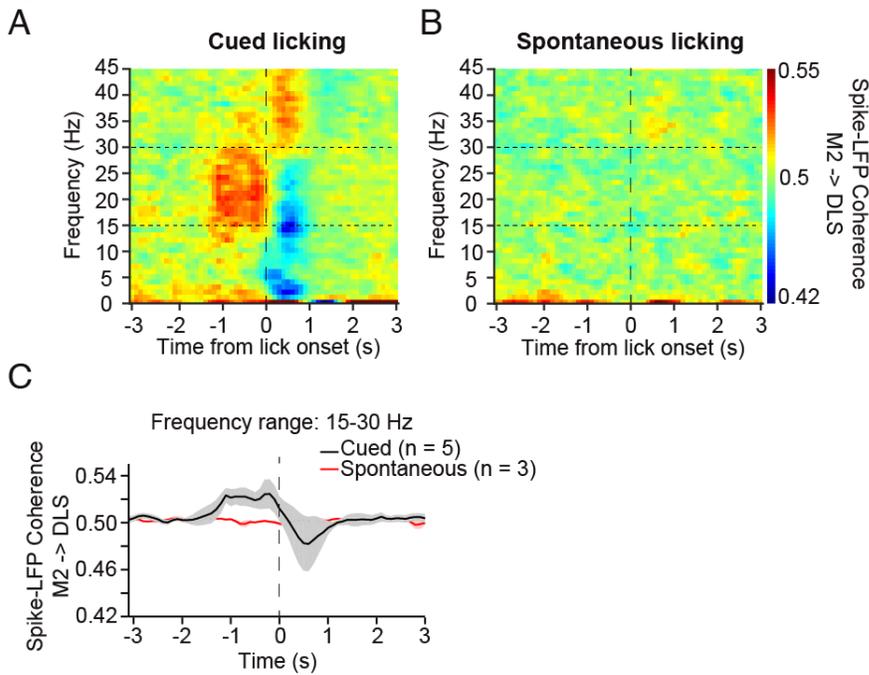
**Figure 3.3.** M2 and DLS show population activity related to movement preparation. *A:* Experiment schema. Dual-region recordings using large-scale silicon microprobes in M2 and DLS. *B:* Example raster plots for two pyramidal neurons recorded in M2 (left) and two MSNs recorded in DLS. Rows are trials sorted according to latency to lick onset time. Red ticks indicate lick onset time for each trial. Black ticks demarcate individual spikes. *C:* Response of all recorded M2 pyramidal neurons (top) and DLS MSNs (bottom) aligned to anticipatory lick onset time (dotted line). Each row in the matrix represents the mean normalized firing rate of one recorded neuron in the corresponding brain area. Units are sorted by their latency to peak firing rate. *D:* Population average baseline-subtracted firing rate for all M2 pyramidal cells and DLS MSNs relative to lick onset time. All error bars are SEM.

We found that on average, MSNs were maximally affected by the laser presentation in the times immediately preceding anticipatory lick onset (**Fig. 3.5E**). While this was consistent with the possibility that we reduced glutamatergic input to the striatum, we were interested in the effects of input suppression on microcircuit dynamics.

We identified the specific time of maximum firing rate differences for each unit (**Fig. 3.6A**). We found that most MSNs showed the greatest difference in rate in the times around movement initiation, and there was a trend toward the majority of MSNs having a maximum difference less than 0 s ( $p = 0.059$ , Wilcoxon sign rank test, **Fig. 3.6B**). The striatal network is known to be regulated by feedforward and, to a lesser extent, feedback inhibition mechanisms (Tepper et al., 2008). We asked if suppressing a specific cortical input resulted in largely inhibited cell activity or whether any neurons were excited as a result of the reduction of glutamatergic input. We thus quantified the probability of detecting significantly inhibited and excited neurons for each time bin relative to lick onset. We found that MSNs were largely inhibited by M2 input suppression and that the rate of detecting MSNs with significant rate increases during laser presentation did not change in the times relative to lick onset (**Fig. 3.6C**). Thus, M2 suppression resulted in a general and widespread reduction in MSN activity. Most neurons were active in the times immediately preceding lick onset in our MSN population. Thus peak firing times closely predicted the times of maximum laser effect (**Fig. 3.6D**). In order to control for the potential confound of most neurons being maximally active prior to licking, we calculated the mean magnitude of firing rate reduction and increases as a function of the peak firing time for each unit. This analysis revealed that when controlling for unit number per time bin, we can still observe a high degree of rate change for units that peak in their firing rate prior to movement initiation (**Fig. 3.6E**).

Lastly, we have previously shown that striatal dynamics show structured spontaneous activity during inter-trial intervals while animals are at rest (Bakhurin et al., 2016). We asked if M2 input suppression would affect DLS spontaneous activity. We found that activating laser

**Figure 3.4**



**Figure 3.4.** M2 and DLS are synchronized prior to cued licking behavior. *A*: Example spike-LFP coherogram showing coherence between M2 spiking activity and DLS LFP activity aligned to onset of anticipatory licking during CS-reward intervals. Vertical dotted line is aligned to lick onset and horizontal lines indicate beta-frequency range. *B*: Example spike-LFP coherogram showing coherence between M2 spiking activity and DLS LFP activity aligned to onset of spontaneous licking outside the task. Lines are the same as in *A*. *C*: Mean spike-LFP coherence in the beta-frequency band (15-30 Hz) averaged across individual animals. Only 3 mice had enough spontaneous licking events to perform analysis. Error bars are SEM.

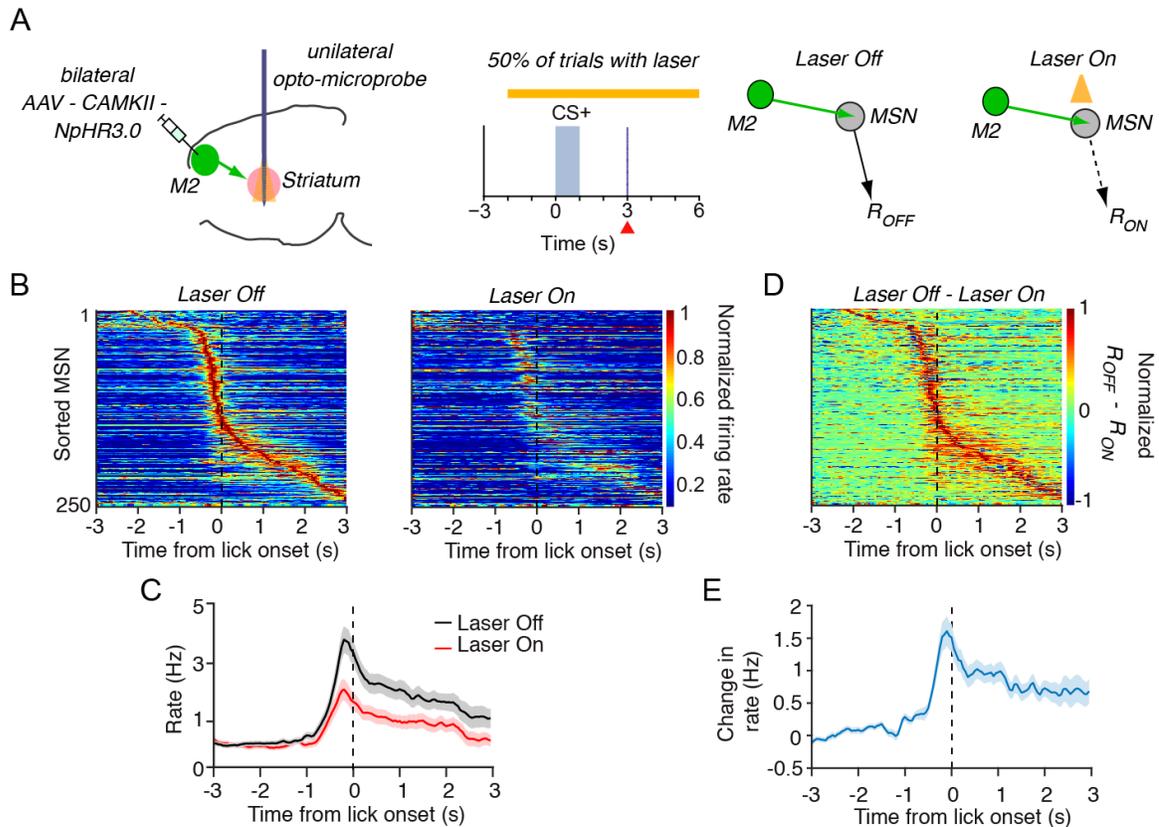
during times when animals were at rest outside of the task did not change MSN firing rates. This suggests that M2 does not continuously interact with the DLS, and that this circuit may be preferentially recruited during cued anticipatory movement generation. Together, these results suggest that M2 input to DLS MSNs is dynamic and that this cortical input does not simply provide a continuous tonic source of excitation to the network. The results also suggest that M2 input to driving DLS activity is largely restricted to cued licking events.

#### *Input suppression displaces DLS population codes for action initiation*

We have previously demonstrated that striatal dynamics contain stable population codes that track movement initiation (Bakhurin et al., 2017). We tested the effects of M2 input suppression on these population dynamics by training a decoder on population dynamics recorded during laser-off trials, and testing the model on other laser-off trials or on laser-on trials. In support of our earlier findings, we found that striatal dynamics in the DLS tracked lick onset time reliably, in that the decoder could identify time bins that contained lick onset times regardless of whether they occurred early or late within the delay during laser-off trials (**Fig. 3.7A**). Next, the decoder was trained on laser-off trials and then tested on laser-on trials. In these cases, we observed that the decoder made more errors in identifying the lick onset time using data during laser-on trials (**Fig. 3.7A**). This suggests that input suppression changes population dynamics related to licking behavior.

To determine the nature of the change in the population dynamics during laser-on trials, we trained the decoder on laser-on trials and applied these models to both laser-on trials. If laser presentation results in worse performance of the classifier, it would suggest that M2 input suppression caused an increase in variability of striatal activity in addition to reductions in firing rates. We found that models trained and tested on laser-on trials performed similarly to those trained and tested on laser-off trials, though there was a trend toward worse classifier performance on laser-on trials ( $p = 0.08$ , paired t-test, **Fig. 3.7B**). When applying laser-off or

**Figure 3.5**



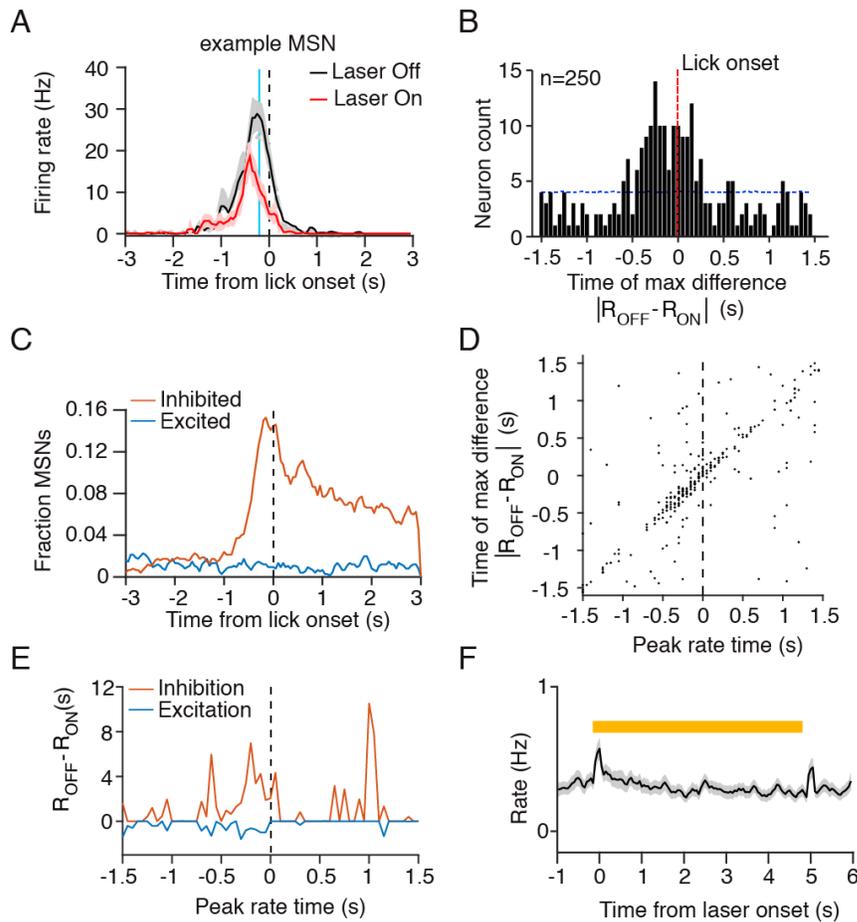
**Figure 3.5.** M2 provides dynamic drive onto DLS population activity. **A:** Experiment schemas. Left: Mice received bilateral virus injections to express eNpHR bilaterally in M2. After training, mice were implanted with the optomicroprobe unilaterally, allowing for large-scale recording from, and laser delivery to the DLS to inactivate M2 projections. Right: Laser was delivered for a total of 8 s beginning 2 s prior to cue onset and terminating 6 seconds after. Laser was pseudorandomly presented during 50% of the trials. **B:** Normalized mean firing rates for all DLS MSNs recorded in conjunction with laser light delivery. Left figure shows  $R_{OFF}$  and middle figure shows  $R_{ON}$  for each neuron. Cells in both matrices are sorted by their latency to peak firing relative to lick onset time in the  $R_{OFF}$  condition and both conditions are normalized to the max in the  $R_{OFF}$  condition. Right figure shows the  $R_{OFF} - R_{ON}$  curves for each neuron. **C:** Population averages of MSN firing rates in the  $R_{OFF}$  and  $R_{ON}$  conditions relative to lick onset time. **D:** Analysis is performed on mean firing rates recorded in the absence ( $R_{OFF}$ ) and in the presence ( $R_{ON}$ ) of laser. **E:** Average  $R_{OFF} - R_{ON}$  curve for the entire MSN population shows a maximum value immediately prior to lick onset time. Error bars are SEM.

laser-on models to their complimentary trial type (cross-trial analysis), both models performed significantly worse. This suggests that M2 input suppression does change the population coding of licking onset, but may simply do this via a significant, but consistent, displacement of DLS dynamic activity in state-space. Importantly, these results also show that in the absence of M2 input, DLS activity retains a dynamic organization of its activity patterns, albeit with potentially higher variability.

### *Input suppression reveals a dynamic corticostriatal transformation across MSNs*

The transformation that the striatal microcircuit performs on incoming excitatory input to generate the striatal output is not well understood. We took advantage of our ability to record firing activity in single neurons both with and without M2 input to attempt to understand how a glutamatergic input is incorporated into MSN spiking activity. Thus by eliminating the activity of glutamatergic input to the DLS, we assume that we can infer the function of corticostriatal synapses when they are intact (Phillips and Hasenstaub, 2016). For each MSN, we calculated the mean firing rate relative to lick onset during both laser-on and laser off trials (**Fig. 3.8A, left**). We next calculated the degree to which firing rates during laser-on trials were a linear function of the firing rates during laser-off trials by determining the  $R^2$  value of this relationship (**Fig. 3.8A, right**). We found that across the population of MSNs that we recorded, we observed many clearly linear relationships between laser-on and laser-off trials. However, we also observed a high number of cells in which simple linear approximations were not sufficient for explaining the relationships between these two trial types (**Fig. 3.8B**). This distribution was not strongly explained by the floor effects in cells that were not modulated by the task, as including only task-modulated cells in this distribution did not eliminate cells with very low  $R^2$  values (**Fig. 3.8C**). This suggests that the transformations that individual MSNs perform on cortical inputs are highly varied, with some cells simply increasing their firing rates in proportion to the amount

**Figure 3.6**



**Figure 3.6.** MSNs are maximally dependent on M2 input prior to lick initiation. **A:** Example MSN showing its  $R_{OFF}$  and  $R_{ON}$  curves relative to lick onset. Dashed line indicates lick onset time and blue vertical line the time of maximum difference between  $R_{OFF}$  and  $R_{ON}$ . **B:** Distribution of times of maximum difference between  $R_{OFF}$  and  $R_{ON}$ . Red dashed line shows lick onset time and blue dashed line the null uniform distribution. **C:** Fraction of all MSNs that show significant inhibition or excitation per time bin between laser-off and laser-on trials relative to lick onset. **D:** Time of maximum difference between  $R_{OFF}$  and  $R_{ON}$  as a function of peak firing rate time for each MSN. **E:** Magnitude of difference between  $R_{OFF}$  and  $R_{ON}$  during the time of peak firing rate across all recorded neurons. **F:** Mean firing rate of all MSNs during laser epochs presented outside of the task. No significantly modulated cells were detected for any time bin, including peri-laser-onset and -offset bins. Error bars are SEM.

of cortical input, and other cells performing highly nonlinear operations on M2 glutamatergic inputs.

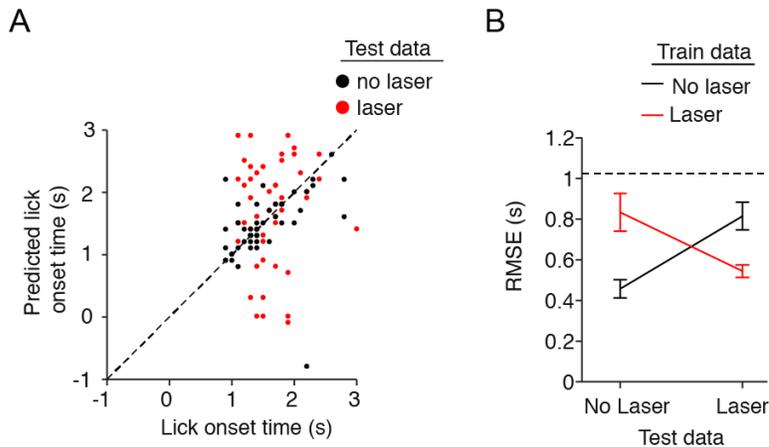
The MSNs in our DLS population are highly diverse in their mean responses relative to licking onset. We thus investigated whether the type of MSN response within the trial had any relationship to the kind of computational operation that MSN performs on M2 corticostriatal input. We first calculated the mean  $R^2$  response as a function of peak firing rate time relative to movement initiation. We found that MSNs with peak firing times closer to movement initiation showed higher  $R^2$  values, suggesting that MSNs have more linear relationships between laser-on and laser-off firing activity as the population activity transitions into movement initiation and movement maintenance periods (**Fig. 3.8D**). In addition, the distribution of  $R^2$  values became more diverse once animals initiated movements (**Fig. 3.8E**). This suggested that perhaps neurons that were differentially modulated by movement revealed different kinds of utilization of M2 corticostriatal inputs.

To investigate this possibility, we clustered the neurons based on their mean firing pattern, obtaining 3 classes of cell responses in the population. The clusters largely reflected neurons whose activity decreased during licking behavior (**Fig. 3.9A**), those cells whose activity was transiently active prior to licking onset (**Fig. 3.9B**), and those whose activity increased during licking (**Fig. 3.9C**). Interestingly, we did not observe strong trends in the distributions of  $R^2$  values across neurons in each of these clusters. This suggests that MSNs receiving M2 input utilize those inputs in a complex and diverse manner, and that similar firing patterns in individual MSNs may not arise from overlapping input activity.

### **Discussion for Chapter 3**

This study used optogenetic manipulations to study the effects of suppressing glutamatergic inputs on the generation of striatal dynamic activity. The striatum has recently been demonstrated to produce complex dynamic trajectories in behaving animals (Bakhurin et al.,

**Figure 3.7**

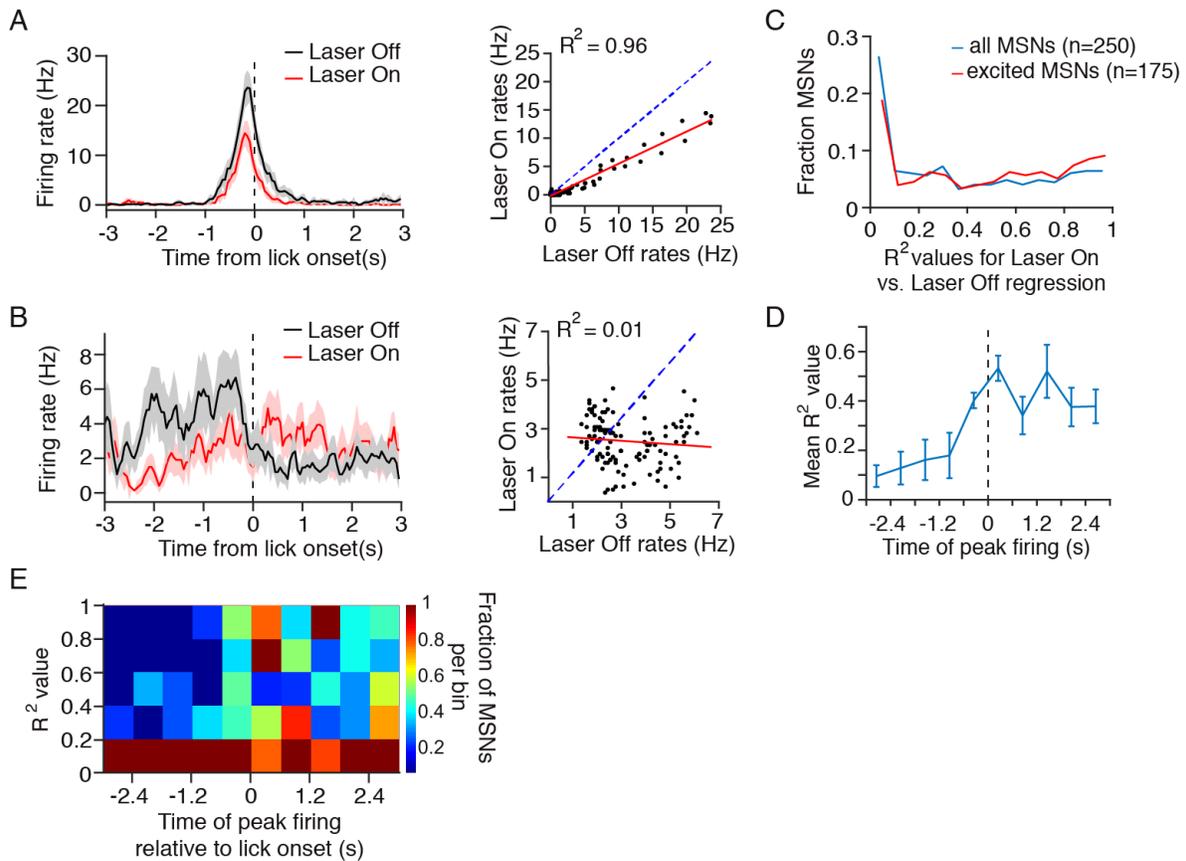


**Figure 3.7.** Dynamic codes for movement initiation remain after input suppression. *A*: Results of lick onset-time decoding in a representative animal. Scatter plots are the predicted lick onset times vs. actual observed lick onset times for each individual trial. Models were trained using non-laser trials and tested either on other non-laser trials (black, within-trial), or on laser-trials (red, cross-trial). *B*: Summary of within-trial and cross-trial decoding performance. Cross-trial testing of the decoder severely impaired lick-onset time prediction performance (a two way, repeated measures ANOVA revealed a significant interaction between training data groups). There was no significant difference between within-trial performance of the decoder on no-laser and laser trials ( $p = 0.133$ , paired t-test). There was no significant difference between cross-trial performance of the decoder on no-laser and laser trials ( $p = 0.234$ , paired t-test). Dotted line reflects random chance performance level. Error bars are SEM.

2017; Gage et al., 2010; Rueda-Orozco and Robbe, 2015; Thorn and Graybiel, 2014). Glutamatergic input is necessary for driving striatal activity. However, the extent to which dynamic population activity of striatal networks is mediated through internal synaptic mechanisms vs. being dependent on complex external signals is not clear. We showed that suppressing cortical inputs to a local striatal network revealed a complex dependency on external glutamatergic input for generating movement-related striatal dynamic activity. Our findings support a model in which cortical inputs to the striatum are not a simply tonic source of excitatory drive onto striatal microcircuitry. However, we also revealed that MSNs differentially depend on premotor input to shape their individual firing responses during movement initiation. This suggests that local microcircuit architecture may be fundamental for determining how a given MSN utilizes a specific glutamatergic signal in generating its spiking activity during behavior.

An important assumption underlies our study in that the MSNs that we recorded receive a myriad of distinct corticostriatal inputs, and not just those from M2. Striatal MSNs contain continuously active potassium channels that keep MSN membrane potentials at a hyperpolarized level (Calabresi et al., 1987; C. J. Wilson et al., 1990). In order to activate, MSNs require synchronized presynaptic excitatory activity in order to achieve depolarization that can lead to spiking activity (Plenz and Kitai, 1998; Plotkin et al., 2011; Stern et al., 1998). In addition to convergent activity, previous studies have also demonstrated that spatiotemporal activation patterns of presynaptic input may lead to distinct MSN responses (Carter et al., 2007). In our study, we found that continuous input suppression of M2-DLS synaptic activity during our task did not result in a uniform manipulation of MSN firing across time. Instead, we found that MSN populations showed dynamic dependence on M2 inputs, with a maximum interaction with those inputs prior to movement initiation. Thus it is possible that corticostriatal inputs may differentially contribute to local MSN activity in the DLS as a function of time relative to behavior.

**Figure 3.8**



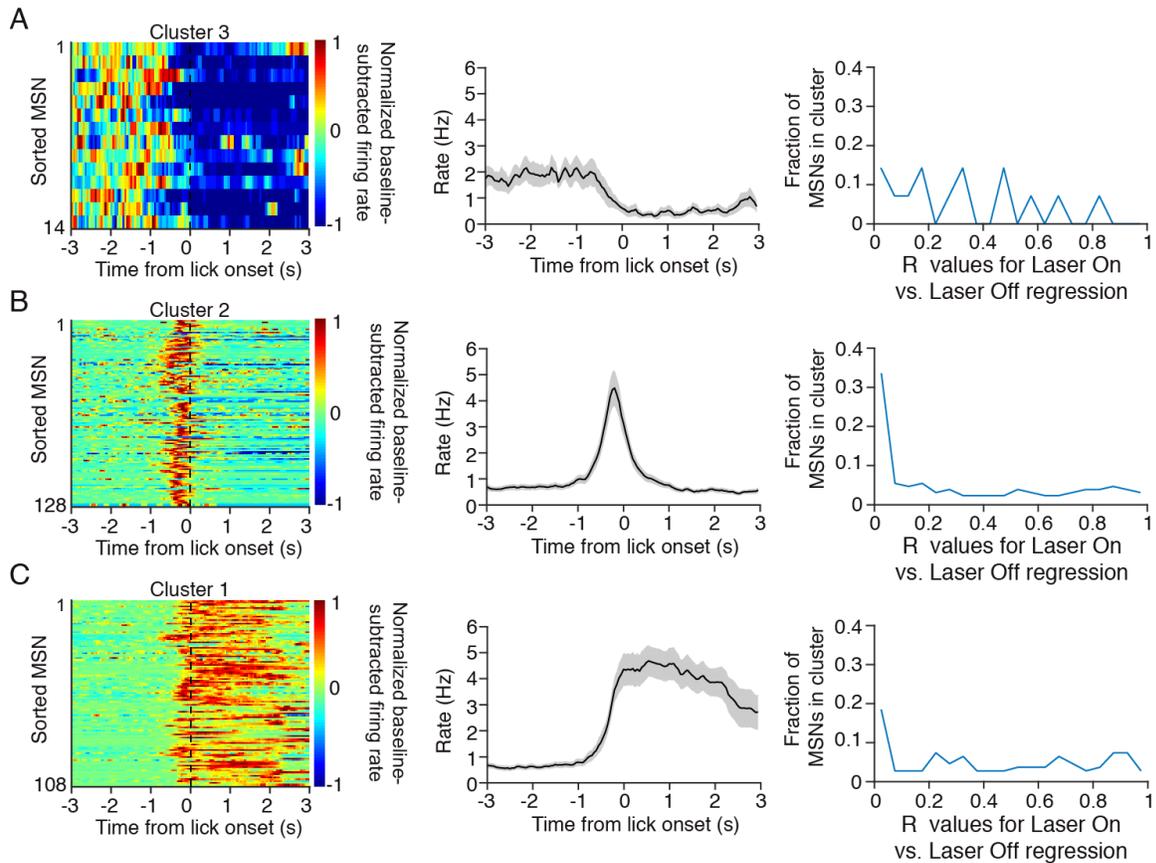
**Figure 3.8.** Input suppression reveals a diverse impact of M2 projections on MSN firing. *A:* Example regression analysis between  $R_{ON}$  and  $R_{OFF}$  firing rates. Left: Mean firing rates across all laser-off and laser-on trials showing firing rate reduction with M2 suppression. Right: firing rate during laser-on trials as a function of firing rates during laser-off condition per time bin. Points reflect the mean rate in each condition across all time bins. Dotted blue line is the unity line, and the red line reflects the line of best fit from least-squares regression. *B:* A second example MSN. Labeling convention the same as in *A*. *C:* Distribution of  $R^2$  values calculated between  $R_{OFF}$  and  $R_{ON}$  across all MSNs included in the study. *D:* Mean  $R^2$  values as a function of peak firing rate latency relative to lick onset across all MSNs. Error bars are SEM.

We also observed that laser presentation outside of the task resulted in little to no effect on MSN firing activity. This finding supports the possibility that M2-DLS interactions vary in their strength as a function of time, and furthermore suggests transient coupling between the two regions depending on task demands. These results also potentially exclude M2 from being the source of network organization in spontaneous striatal activity in the DLS recorded during resting state (see Chapter 1 and (Bakhurin et al., 2016)). However, this statement requires direct testing to determine that correlations among neurons are indeed intact during M2 input suppression. Alternative sources of spontaneous network activity in the striatum could be somatosensory regions or thalamic inputs (Reig and Silberberg, 2014; Smith et al., 2014).

In addition to dynamic influences of M2 on DLS firing rates, we also observed evidence that M2 and DLS showed increased spike-LFP coherence during preparatory phases of movement. Interestingly, these interactions were not present during spontaneous licking events that occurred outside of cue-reward intervals. However, because firing rates were also lower during such events, it is possible that M2-DLS coherence is an artifact of increased firing rates (Aoi et al., 2015; Vinck et al., 2010). In addition, though other groups have shown evidence for spike-field coherence between primary motor cortex and sensorimotor striatum, they showed increases between 6-14 Hz (Koralek et al., 2013), whereas we observed coherence in the  $\beta$ -frequency band, between 15-35 Hz. Though oscillations in the striatum at these frequencies have been linked to pathological conditions in humans (Boroud et al., 2006), they have also been linked to normal movement preparation (Courtemanche et al., 2003). Still, the discrepancy between our findings and those showing corticostriatal coherence at lower frequencies may be related to the different cortical areas being used.

Our results using bilateral M2 inhibition showed that these manipulations resulted in reduced likelihood of anticipatory licking behavior without changes in lick onset time or onset time variability. Though similar, these findings are inconsistent with previous work. First, it was shown in monkeys that brief premotor cortex stimulation would delay reaching behavior without

**Figure 3.9**



**Figure 3.9.**  $R^2$  values do not group according to MSN response classification. **A:** Left: Heat plot depicting mean responses of MSNs classified as inhibited. Each row in the matrix represents the mean firing rate of one recorded MSN normalized to its maximum rate aligned to the onset of cued anticipatory licking. Middle: Mean firing rates of all MSNs classified as inhibited. Right: Distribution of  $R^2$  values for all inhibited MSNs. **B:** Same as in A for MSNs classified as transiently responding. **C:** Same as in A for MSNs classified as sustained responding. Error bars are SEM.

affecting movement kinematics (Churchland and Shenoy, 2007). While we are unable to measure licking movement kinematics with much detail, we did not see a change in lick onset timing for licks that occurred prior to reward, suggesting that our manipulation did not change onset timing. One possibility for the difference between previous work and our work (besides species differences) is that our study did not require mice to undergo a variable waiting period. Perhaps if we were to introduce a waiting period that was terminated by a go signal, we could manipulate the preparatory function that premotor cortex may play in movement (Churchland et al., 2010). In another study with manipulations in M2 in mice, it was found that M2 suppression through PV interneuron stimulation using ChR2 would disrupt animals' ability to make correct discriminatory behaviors in a left/right lick decision-making task (Li et al., 2016). However, these results suggest that M2 inhibition resulted in increased directionality errors, but not reduced licking probability. This stands in contrast to our study, and may be attributed to the different inhibition approaches.

Our regression analysis was based upon the assumption that removing cortical input to MSNs allows us to infer the contribution that the input makes to normal MSN function (S.-H. Lee et al., 2012; Phillips and Hasenstaub, 2016). The general result of M2 projection inhibition is that MSNs show reduced firing rates after turning off these inputs. However, we also showed that individual neurons showed highly varied responses to these perturbations. In the simplest cases, M2 input provided a constant gain on spiking rate of MSNs. This suggests that for these MSNs with high  $R^2$  values, their temporal dynamics remained intact but with a reduced magnitude. For most other MSNs, however, M2 input suppression led to a dramatic reorganization of their firing activity. Interestingly, after we grouped MSNs based on their activity pattern with respect to licking, we found that neurons with similar activation profiles did not show uniform responses to M2 suppression. This therefore suggests that common firing activity in the striatum may not necessarily be the result of overlapping input. On the other hand, it is possible

that the effects of M2 suppression on individual MSN activity are largely determined by the MSN's local microcircuit architecture.

Our results suggest that cortical input to the striatum is not simply providing tonic drive to facilitate local dynamics. However, our results also show that in the absence of the premotor input, striatal networks continue to generate dynamic, albeit slightly more variable, activity, as evidenced by our decoding results. Whether this activity originates from within the striatal network or is dependent on other sources of excitation that we did not manipulate remains to be further determined. However, we propose that striatal dynamic activity with respect to movement planning and initiation results from a complex interaction of local microcircuit activity and cortical input.

## **General Discussion**

Here, I have described several studies on striatal network level activity in behaving mice. In chapter 1, I showed that plasticity mechanisms and topographic organization of striatal inputs result in long-lasting organization of networks of MSNs and FSIs. In chapter 2, I directly studied how striatal population dynamics may track time during predictive behavior, showing that striatal activity conforms to theorized mechanisms of time keeping in the brain. Finally, in chapter 3, I presented experiments that use optogenetic approaches to test the causal role of excitatory corticostriatal inputs in shaping striatal dynamic activity. This work represents a range of approaches that I have learned and applied to studying large-scale neural activity in behaving animals. In this last section, I would like to spend time providing future experiments for each study. Some proposed experiments are efforts to address some caveats or test alternative avenues of interpretation for the data I report above.

### *Memory formation and distribution in the striatum*

In chapter 1, I developed a Pavlovian odor discrimination task and used large-scale recording approaches to simultaneously capture spiking activity from dozens of striatal neurons as mice learned the association for the first time (Bakhurin et al., 2016). The study's primary finding involved neurons that discriminated between reward-predicting and non-predicting odor trials. Their activity patterns reflected the goal of the task that animals needed to learn, suggesting that they were involved in learning the behavioral discrimination between odors. One significant question regarding striatal-based learning is how subsets of cells like these are selected for during learning, as opposed to others. Striatal neurons receive a dense convergence of excitatory input from thousands of axons potentially all carrying distinct information (Kincaid et al., 1998). Thus it is conceivable that any MSN could become recruited into an ensemble that is involved in generating discriminatory behavior. A number of experiments could be proposed that would further characterize the linking of MSNs to specific cues, outcomes, and actions.

First, how would the discriminatory neurons in our study be recruited if animals must learn a new association, such as a reversal of odor reward relationships, such that the original CS- becomes the new CS+? At least one study has performed this kind of experiment while recording in rats running a T-maze, in which the rewarded arm is switched after rats reach high performance (Kubota et al., 2009). This study found very little alteration in the activity of striatal neuron populations after the reversal. My interpretation of these negative results is that as long as behavior remains similar, regardless of cues, striatal activity will remain fairly constant across the reversal. Thus, in our task, I would expect the same population of neurons to be recruited after a reversal. Our discriminating neurons were located in the lateral areas subregion of the striatum. Neurons in this subdomain has previously been associated with orofacial sensation and movement (Carelli and West, 1991; Cho and West, 1997), which links the neurons with our study with physical aspects of the behavior.

Second, an extension of the logic that behavior is the ultimate determinant of striatal activity would suggest that interesting interactions between cell populations could be elicited by having animals learn distinct actions. There are many questions that could be asked if it could be reliably demonstrated that divergent behaviors resulted in non-overlap in striatal activation. While licking may not provide enough diversity in repertoire, many studies have investigated limb-based skills in rodents that demonstrate the diversity of movements that could be studied (Kawai et al., 2015; Mathis et al., 2017; Yttri and Dudman, 2016). Restricting movements to a single limb would theoretically constrain activity to a subdomain of the dorsal striatum and thus allow testing of ensemble overlap in various skills. One simple prediction for ensemble interaction comes from extensive work in the field of memory allocation, where it is believed that memories are more likely to overlap in neurons that are more excitable (Ramirez et al., 2013; Silva et al., 2009). Testing of these frameworks within non-declarative motor-learning systems has not yet been undertaken.

This study involved recording from the striatum during the first day of learning. We were surprised to find that medial striatum was not activated in this task, and that all interesting activity was distributed on the lateral side of the structure. This contrasted with previous work that has shown that learning transitions from medial to lateral areas of the sensorimotor striatum (Thorn et al., 2010; Yin et al., 2005; 2004). These studies were performed using navigation-tasks or instrumental lever pressing, and thus, comparisons may be difficult to make. In addition, the finding that activity was restricted to lateral areas is more consistent with the known topographic inputs of the striatum. However, our study does not test for the necessity of the striatum for learning this association. Therefore, it is difficult to know whether activity that we did record was truly involved in learning, or whether striatal activity in our task is simply registering changes in motor behavior (Yin, 2016; 2014).

This study also used statistical characterization of simultaneously recorded networks to show that the discriminating neuron show higher rates of correlated activity with one another than with neurons that were not discriminating. Importantly, these correlations were measured outside of the task while animals were not moving in order to control for extraneous factors contributing to their statistical relationships. Our findings have parallels in other fields that show population activity outside of task performance that retain task-related organization (Jadhav et al., 2012; S. Xu et al., 2012). Another study has shown that correlated cortical neurons *in vivo* are more likely to be physically connected (Ko et al., 2011). Owing to the lack of consensus about the importance of lateral interactions in the striatum, we speculated that correlations in our recordings arose from common sources of input.

What could be the origins of these inputs? Presumably these would be extrastriatal glutamatergic axons, but whether they arise from one region or many would require some additional experiments. If these inputs come from one source, do correlations disappear after removing that input? If they come from multiple sources, are these input areas also correlated to each other and thus send synchronous inputs to the striatum?

We also performed analysis to ask if the probability of detecting of these correlations increased with learning. We did not see changes in correlations throughout the recording, which suggested that they were established prior to the experiment. One experiment that could be done to test this hypothesis is to record resting state activity prior to beginning the odor training. Spontaneous correlations among neurons should be present prior to any learning. This would reinforce the need to perform these experiments using more arbitrary movements that animals have not been exposed to prior to training, which is not the case with licking behavior as they have been performing this action their whole lives. One recent study has demonstrated interesting results with this kind of approach, showing that striatal neurons will increase their firing rates in response to task-specific excitatory inputs that were not previously encountered (Xiong et al., 2015). However, the correlations among these neurons were not explored.

Altogether, our work was some of the first to do these kinds of analyses in the striatum in behaving animals, and to bring network analysis measures to subcortical structures. While we could have made more of a direct link to behavior, our observations can be expanded into a number of follow-up experiments, a number of which I described here.

### *Expansion on time representation in striatal networks*

In chapter 2, I focused on characterizing striatal dynamic activity within the framework of time representation in neural networks (Bakhurin et al., 2017). In this study, we took advantage of our abilities to record large-scale network activity in multiple brain regions of the same animal. This work operated under the theory that the brain tells time by learning the patterns of activity that its neural networks generate as it performs its computational functions. Our laboratory routinely records dynamics that are very similar to those that can be generated in neural networks, and that can be detected in many other brain areas. We found that directly comparing neural network activity of two brain areas results in a difference in their ability to keep track of time, suggesting that while time encoding may be widely distributed in the brain, it is not

uniformly tracked. Many follow-up questions can be asked of these dynamics and their relationship to time.

The population clock framework makes two general assumptions that must hold in order for it to be truly implemented in the brain. First, population clocks must be reliable in that too much noise in network activity will render them poor clocks. This reliability is ensured by the brain through plasticity mechanisms and learning. Our study used machine learning approaches to demonstrate that both striatal and OFC activity was indeed reoccurring within individual trials. The second assumption is more difficult to demonstrate but is equally important: Is anything listening to these reliable population clocks? In neural network models of these clocks, a read-out population is trained to detect specific population configurations that emerge as a function of time.

How can these read-out populations be detected or identified *in vivo*? In one sense, disrupting the population activity should alter timing. Timing has been altered by inactivating striatal activity in another study (Mello et al., 2015), but not in a Pavlovian task like the one we have implemented. Furthermore, we do not gain understanding of how these timing sources are actually used by downstream areas, and do not exclude the possibility that read-out populations may be located within the same network that is generating the action. In addition, while many studies have explored the ways that distinct, spatially separated brain regions interact (Engel et al., 1991; Koralek et al., 2013; Popescu et al., 2009; Zandvakili and Kohn, 2015), these kinds of analyses have not been explicitly applied to tasks involving timing. Our multi-region, large-scale approaches are ideal for these kinds of experiments. For example, do specific activity patterns in upstream regions such as the striatum predict population activity in basal ganglia output nuclei, such as the substantia nigra?

Our study analyzed dynamics that were recorded within a small temporal interval. We argue that animals are using time in the task because, 1) their licking onset times during the interval are not random and uniformly distributed, but appear to center around a specific

setpoint, and 2) unpublished data from our group shows that mice can learn to time two different intervals that are presented in a randomized fashion. However, within this short interval, neural activity in response to the odor cue may overlap with neural activity that is driven by anticipatory licking occurring prior to the reward. The effect of a short interval could be that two distinct populations in the striatum have overlap in their activation patterns, and may therefore be interpreted as a continuously active dynamic trajectory. One way to determine if this is the case is to present the reward a much later time relative to the odor cue. Although we have previously shown that mice can time two different intervals, one being twice as long as the other, we have not performed striatal recordings using this data, and thus do not know how the same network handles short vs. long intervals. If our analysis was performed on overlapping cue and motor related activity, then these patterns may separate during the interval and would reveal a gap in striatal activity. Other studies using striatal dynamics have tested this by using very long intervals, and argue that striatal dynamics are continuously active and scale their transition speed in response to varying timing demands (Gouvêa et al., 2015; Mello et al., 2015). Do our recorded activity patterns expand to accommodate longer intervals as well? If they were not continuous during the longer interval, it would imply that the timing originated in an area outside of the striatum.

Finally, it would be useful to determine how other cortical areas are involved in representing the temporal interval. We chose the OFC as our cortical area to compare with the striatum, but often received criticism for this choice, as the OFC may not necessarily be involved in the behavior we were using. Several studies have shown that prefrontal cortex and premotor cortex are capable of, and sometimes required for, encoding time (Crowe et al., 2014; Kim et al., 2009; Merchant et al., 2013b). Our model is that the striatum integrates across many brain areas to generate its population code. Thus we proposed that cortical regions are perhaps less well tuned to time and the striatum is recruited to refine this signal. This statement requires further testing. Once it is determined that striatal output is indeed necessary for timing in our

task, it would be important to find the input areas that drive this activity. The work that I describe in chapter 3 shows how our lab is beginning to investigate these kinds of questions.

*The importance of local microcircuitry for generating striatal output*

In chapter 3, I investigated the effects of optogenetically inhibiting inputs to the lateral striatum from a premotor cortical area known to be important for guiding licking behavior. I showed evidence that these two brain regions interact specifically during cued licking behavior, and that M2 inputs appear to shape striatal activity predominantly during movement preparatory periods. Finally, I used the ability to transiently silence inputs onto MSNs to characterize how these inputs normally contribute to MSN firing activity. While it is clear that the MSN populations showed reduced overall activity after laser activation, these effects were not uniform across the population, leaving room for other excitatory inputs or local microcircuit interactions continue supporting striatal output.

Several challenges exist in the interpretation of these results. Although we eliminated the contribution of one source of excitatory input to the striatum, it is very likely that the neurons that we recorded from still received excitation from other sources that we did not manipulate. One candidate is the region of somatosensory cortex relating to orofacial sensation, which should project to a similar place as motor inputs do (Carmichael and PRICE, 1995; Hoffer and Alloway, 2001). In considering future experiments with this circuit, I would be interested in using retrograde tracing to determine what cortical regions could theoretically contribute most to the activity of the region in the DLS that we targeted. Then inhibitory opsins could be injected into these areas and the same kinds of recordings could be made in the striatum with optical inhibition of inputs from two regions. First, what would  $R^2$  distributions look like with this manipulation? It is difficult to speculate on the shape of this distribution, but if firing rates are reduced even more than with a single suppressing a single input, as we did in this study, perhaps the result would be a floor effect of the  $R_{ON}$  curve. The resulting distribution might be

pushed further towards lower  $R^2$  values. A second question involves the state of temporal dynamics in the absence of multiple glutamatergic inputs. In the study, I showed that population coding in the times prior to licking maintains a dynamic activity pattern in the absence of M2 input, possibly due to other inputs. The increased decline of dynamic population responses after removing more sources of glutamate would support my prediction that remaining glutamatergic input supports striatal activity patterns in the absence of M2 input.

An alternative source of striatal dynamics could arise from local microcircuit activity. Although the striatal recurrent network is largely inhibitory, it has been demonstrated that these networks are able to generate complex activity patterns with only tonic excitatory input (Carrillo-Reid et al., 2008). In order to determine how much local interactions contribute to population codes related to anticipatory activity, one could infuse GABAergic antagonists, such as bicuculline or gabazine, into the striatum, thus suppressing these local signals. Importantly, glutamatergic signals would remain intact, and could still theoretically drive striatal output. Comparing spiking activity in MSNs in the absence of GABA blockade with activity in the presence of blockade could give some valuable insight into the role that local GABAergic interactions play in generating spiking output. Such approaches, in addition to using optogenetic manipulations, have proven to be highly valuable in mapping circuit organization in cortical areas (Isaacson and Scanziani, 2011). These results would also test a number of predictions made using neuronal network models, namely that striatal networks are capable of generating intrinsic activity patterns with only a tonic source of glutamate (Ponzi and Wickens, 2010).

I observed a diversity of effects of M2 input suppression on MSN activity. In some cases, M2 input only seemed to provide a simple linear contribution to individual MSN firing rates, suggesting that the glutamatergic input only served to amplify MSN output signals. However, these effects were largely an exception rather than a rule in terms of how M2 inputs contributed to MSN activity. Interestingly, these effects were not largely explained by grouping cells according to their response profile in the laser-off control condition. This result presents an

interesting challenge, as it suggests that MSNs with similar activity patterns do not necessarily generate these responses because of overlapping sources of glutamatergic drive. I describe in chapter 3 how perhaps MSN position within local microcircuit architecture may account for these discrepancies. Another potential contribution could be that the M2 inputs are highly diverse themselves, and that we were indiscriminant as to which kinds of input patterns we suppressed with our manipulation. An additional possibility is that different MSNs, in addition to their receipt of M2 signals, may also have many different glutamatergic synapses from diverse input areas. Because the M2 input happened to undergo the greatest level of plasticity, this is what is most represented in MSN output. The suppression of this input thus may reveal the diversity of synaptic connections that individual MSNs have with other areas that may not have been as involved in learning. One experiment that could address this possibility is one where inputs are suppressed during early stages of learning. Perhaps then, it would be possible to show that M2 contributions to individual MSN firing activity develops over time.

## References

- Adler, A., Katabi, S., Finkes, I., Israel, Z., Prut, Y., Bergman, H., 2012. Temporal convergence of dynamic cell assemblies in the striato-pallidal network. *Journal of Neuroscience* 32, 2473–2484. doi:10.1523/JNEUROSCI.4830-11.2012
- Alexander, G.E., DeLong, M.R., Strick, P.L., 1986. Parallel organization of functionally segregated circuits linking basal ganglia and cortex 9, 357–381.
- Allman, M.J., Meck, W.H., 2012. Pathophysiological distortions in time perception and timed performance. *Brain* 135, 656–677. doi:10.1093/brain/awr210
- Aoi, M.C., Lepage, K.Q., Kramer, M.A., Eden, U.T., 2015. Rate-adjusted spike–LFP coherence comparisons from spike-train statistics. *Journal of Neuroscience Methods* 240, 141–153. doi:10.1016/j.jneumeth.2014.11.012
- Aosaki, T., Tsubokawa, H., Ishida, A., Watanabe, K., Graybiel, A.M., Kimura, M., 1994. Responses of tonically active neurons in the primate's striatum undergo systematic changes during behavioral sensorimotor conditioning. *J Neurosci* 14, 3969–3984.
- Arieli, A., Sterkin, A., Grinvald, A., Aertsen, A.D., 1996. Dynamics of ongoing activity: explanation of the large variability in evoked cortical responses. *Science*.
- Atallah, H.E., McCool, A.D., Howe, M.W., Graybiel, A.M., 2014. Neurons in the ventral striatum exhibit cell-type-specific representations of outcome during learning. *Neuron* 82, 1145–1156. doi:10.1016/j.neuron.2014.04.021
- Averbeck, B.B., Latham, P.E., Pouget, A., 2006. Neural correlations, population coding and computation. *Nat Rev Neurosci* 7, 358–366. doi:10.1038/nrn1888
- Averbeck, B.B., Lee, D., 2006. Effects of noise correlations on information encoding and decoding. *J. Neurophysiol.* 95, 3633–3644. doi:10.1152/jn.00919.2005
- Bair, W., Zohary, E., Newsome, W.T., 2001. Correlated firing in macaque visual area MT: time scales and relationship to behavior. *Journal of Neuroscience* 21, 1676–1697.
- Bakhurin, K.I., Goudar, V., Shobe, J.L., Claar, L.D., Buonomano, D.V., Masmanidis, S.C., 2017. Differential Encoding of Time by Prefrontal and Striatal Network Dynamics. *J Neurosci* 37, 854–870. doi:10.1523/JNEUROSCI.1789-16.2017
- Bakhurin, K.I., Mac, V., Golshani, P., Masmanidis, S.C., 2016. Temporal correlations among functionally specialized striatal neural ensembles in reward-conditioned mice. *J. Neurophysiol.* 115, 1521–1532. doi:10.1152/jn.01037.2015
- Balleine, B.W., Delgado, M.R., Hikosaka, O., 2007. The role of the dorsal striatum in reward and decision-making. *Journal of Neuroscience* 27, 8161–8165. doi:10.1523/JNEUROSCI.1554-07.2007
- Balleine, B.W., Dickinson, A., 1998. Goal-directed instrumental action: contingency and incentive learning and their cortical substrates. *Neuropharmacology* 37, 407–419.

- Barbera, G., Liang, B., Zhang, L., Gerfen, C.R., Culurciello, E., Chen, R., Li, Y., Lin, D.-T., 2016. Spatially Compact Neural Clusters in the Dorsal Striatum Encode Locomotion Relevant Information. *Neuron* 92, 202–213. doi:10.1016/j.neuron.2016.08.037
- Barnes, T.D., Kubota, Y., Hu, D., Jin, D.Z., Graybiel, A.M., 2005. Activity of striatal neurons reflects dynamic encoding and recoding of procedural memories. *Nature* 437, 1158–1161. doi:10.1038/nature04053
- Bartolo, R., Prado, L., Merchant, H., 2014. Information processing in the primate basal ganglia during sensory-guided and internally driven rhythmic tapping. *Journal of Neuroscience* 34, 3910–3923. doi:10.1523/JNEUROSCI.2679-13.2014
- Beckstead, R.M., Domesick, V.B., Nauta, W.J., 1979. Efferent connections of the substantia nigra and ventral tegmental area in the rat. *Brain Res.* 175, 191–217.
- Beiser, D.G., Houk, J.C., 1998. Model of cortical-basal ganglionic processing: encoding the serial order of sensory events. *J. Neurophysiol.* 79, 3168–3188.
- Bennett, B.D., Wilson, C.J., 1999. Spontaneous activity of neostriatal cholinergic interneurons in vitro. *J Neurosci* 19, 5586–5596.
- Berke, J.D., 2011. Functional properties of striatal fast-spiking interneurons. *Front Syst Neurosci* 5, 45. doi:10.3389/fnsys.2011.00045
- Berke, J.D., Okatan, M., Skurski, J., Eichenbaum, H.B., 2004. Oscillatory entrainment of striatal neurons in freely moving rats. *Neuron* 43, 883–896. doi:10.1016/j.neuron.2004.08.035
- Bermudez, M.A., Schultz, W., 2014. Timing in reward and decision processes. *Philosophical Transactions of the Royal Society B: Biological Sciences* 369, 20120468–20120468. doi:10.1098/rstb.2012.0468
- Bliss, T.V., Lomo, T., 1973. Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path. *J Physiol (Lond)* 232, 331–356.
- Bolam, J.P., Hanley, J.J., Booth, P.A., Bevan, M.D., 2000. Synaptic organisation of the basal ganglia. *J. Anat.* 196 ( Pt 4), 527–542.
- Bonifazi, P., Goldin, M., Picardo, M.A., Jorquera, I., Cattani, A., Bianconi, G., Represa, A., Ben-Ari, Y., Cossart, R., 2009. GABAergic hub neurons orchestrate synchrony in developing hippocampal networks. *Science* 326, 1419–1424. doi:10.1126/science.1175509
- Boroud, T., Brown, P., Goldberg, J.A., Graybiel, A.M., Magill, P.J., 2006. Oscillations in the basal ganglia, in: Bolam, J.P., Ingham, C.A., Magill, P.J. (Eds.), *The Basal Ganglia VIII, The Good, the Bad, and the Unexpected*. pp. 3–24.
- Brody, C.D., Hernández, A., Zainos, A., Romo, R., 2003. Timing and neural encoding of somatosensory parametric working memory in macaque prefrontal cortex. *Cereb. Cortex* 13, 1196–1207. doi:10.1093/cercor/bhg100
- Brown, E.N., Kass, R.E., Mitra, P.P., 2004. Multiple neural spike train data analysis: state-of-

- the-art and future challenges. *Nature Neuroscience* 7, 456–461. doi:10.1038/nn1228
- Brown, L.L., 1992. Somatotopic organization in rat striatum: evidence for a combinational map. *Proc. Natl. Acad. Sci. U.S.A.* 89, 7403–7407.
- Brown, R.E., Milner, P.M., 2003. The legacy of Donald O. Hebb: more than the Hebb synapse., *Nature reviews. Neuroscience.* doi:10.1038/nrn1257
- Buhusi, C.V., Meck, W.H., 2005. What makes us tick? Functional and neural mechanisms of interval timing. *Nat Rev Neurosci* 6, 755–765. doi:10.1038/nrn1764
- Bullmore, E., Sporns, O., 2009. Complex brain networks: graph theoretical analysis of structural and functional systems. *Nat Rev Neurosci* 10, 186–198. doi:10.1038/nrn2575
- Buonomano, D.V., 2007. The biology of time across different scales. *Nat. Chem. Biol.* 3, 594–597. doi:10.1038/nchembio1007-594
- Buonomano, D.V., 2000. Decoding temporal information: A model based on short-term synaptic plasticity. *Journal of Neuroscience* 20, 1129–1141.
- Buonomano, D.V., Karmarkar, U.R., 2002. How do we tell time? *neuroscientist* 8, 42–51. doi:10.1177/107385840200800109
- Buonomano, D.V., Laje, R., 2010. Population clocks: motor timing with neural dynamics. *Trends Cogn. Sci. (Regul. Ed.)* 14, 520–527. doi:10.1016/j.tics.2010.09.002
- Buonomano, D.V., Maass, W., 2009. State-dependent computations: spatiotemporal processing in cortical networks. *Nat Rev Neurosci* 10, 113–125. doi:10.1038/nrn2558
- Buonomano, D.V., Merzenich, M.M., 1995. Temporal information transformed into a spatial code by a neural network with realistic properties. *Science.*
- Buzsáki, G., 2004. Large-scale recording of neuronal ensembles. *Nature Neuroscience* 7, 446–451. doi:10.1038/nn1233
- Cachope, R., Mateo, Y., Mathur, B.N., Irving, J., Wang, H.-L., Morales, M., Lovinger, D.M., Cheer, J.F., 2012. Selective activation of cholinergic interneurons enhances accumbal phasic dopamine release: setting the tone for reward processing. *Cell Rep* 2, 33–41. doi:10.1016/j.celrep.2012.05.011
- Cajal, S.R., 1899. *Texture of the Nervous System of Man and the Vertebrates (an annotated and edited translation of the original Spanish text with the additions of the French ....*
- Calabresi, P., Maj, R., Pisani, A., Mercuri, N.B., Bernardi, G., 1992a. Long-term synaptic depression in the striatum: physiological and pharmacological characterization. *J Neurosci* 12, 4224–4233.
- Calabresi, P., Misgeld, U., Dodt, H.U., 1987. Intrinsic membrane properties of neostriatal neurons can account for their low level of spontaneous activity. *Neuroscience* 20, 293–303.
- Calabresi, P., Pisani, A., Mercuri, N.B., Bernardi, G., 1992b. Long-term Potentiation in the

Striatum is Unmasked by Removing the Voltage-dependent Magnesium Block of NMDA Receptor Channels. *Eur J Neurosci* 4, 929–935.

- Cardinal, R.N., Parkinson, J.A., Lachenal, G., Halkerston, K.M., Rudarakanchana, N., Hall, J., Morrison, C.H., Howes, S.R., Robbins, T.W., Everitt, B.J., 2002. Effects of selective excitotoxic lesions of the nucleus accumbens core, anterior cingulate cortex, and central nucleus of the amygdala on autoshaping performance in rats. *Behav. Neurosci.* 116, 553–567. doi:10.1037//0735-7044.116.4.553
- Carelli, R.M., West, M.O., 1991. Representation of the body by single neurons in the dorsolateral striatum of the awake, unrestrained rat. *J. Comp. Neurol.* 309, 231–249. doi:10.1002/cne.903090205
- Carmichael, S.T., PRICE, J.L., 1995. Sensory and premotor connections of the orbital and medial prefrontal cortex of macaque monkeys. *J. Comp. Neurol.* 363, 642–664. doi:10.1002/cne.903630409
- Carnevale, F., de Lafuente, V., Romo, R., Barak, O., Parga, N., 2015. Dynamic Control of Response Criterion in Premotor Cortex during Perceptual Detection under Temporal Uncertainty. *Neuron* 86, 1067–1077. doi:10.1016/j.neuron.2015.04.014
- Carrillo-Reid, L., Hernández-López, S., Tapia, D., Galarraga, E., Bargas, J., 2011. Dopaminergic modulation of the striatal microcircuit: receptor-specific configuration of cell assemblies. *Journal of Neuroscience* 31, 14972–14983. doi:10.1523/JNEUROSCI.3226-11.2011
- Carrillo-Reid, L., Tecuapetla, F., Tapia, D., Hernández-Cruz, A., Galarraga, E., Drucker-Colin, R., Bargas, J., 2008. Encoding network states by striatal cell assemblies. 99, 1435–1450. doi:10.1152/jn.01131.2007
- Carter, A.G., Soler-Llavina, G.J., Sabatini, B.L., 2007. Timing and Location of Synaptic Inputs Determine Modes of Subthreshold Integration in Striatal Medium Spiny Neurons. *J Neurosci* 27, 8967–8977. doi:10.1523/JNEUROSCI.2798-07.2007
- Chang, C.C., Lin, C.J., 2011. LIBSVM: A library for support vector machines. *ACM Transactions on Intelligent Systems and ....*
- Chiba, A., Oshio, K.-I., Inase, M., 2008. Striatal neurons encoded temporal information in duration discrimination task. *Exp Brain Res* 186, 671–676. doi:10.1007/s00221-008-1347-3
- Cho, J., West, M.O., 1997. Distributions of single neurons related to body parts in the lateral striatum of the rat. *Brain Res.* 756, 241–246.
- Churchland, M.M., Cunningham, J.P., Kaufman, M.T., Foster, J.D., Nuyujukian, P., Ryu, S.I., Shenoy, K.V., 2012. Neural population dynamics during reaching. *Nature* 487, 51–56. doi:10.1038/nature11129
- Churchland, M.M., Cunningham, J.P., Kaufman, M.T., Ryu, S.I., Shenoy, K.V., 2010. Cortical Preparatory Activity: Representation of Movement or First Cog in a Dynamical Machine? *Neuron* 68, 387–400. doi:10.1016/j.neuron.2010.09.015

- Churchland, M.M., Shenoy, K.V., 2007. Delay of Movement Caused by Disruption of Cortical Preparatory Activity. *J. Neurophysiol.* 97, 348–359. doi:10.1152/jn.00808.2006
- Clark, A., 2015. Radical Predictive Processing. *The Sou. Jour. of Phil* 53, 3–27. doi:10.1111/sjp.12120
- Cohen, M.R., Kohn, A., 2011. Measuring and interpreting neuronal correlations. *Nature Neuroscience* 14, 811–819. doi:10.1038/nn.2842
- Cohen, M.R., Maunsell, J.H.R., 2009. Attention improves performance primarily by reducing interneuronal correlations. *Nature Neuroscience* 12, 1594–1600. doi:10.1038/nn.2439
- Cole, M.W., Bassett, D.S., Power, J.D., Braver, T.S., Petersen, S.E., 2014. Intrinsic and Task-Evoked Network Architectures of the Human Brain. *Neuron* 83, 238–251. doi:10.1016/j.neuron.2014.05.014
- Connolly, C.I., Burns, J.B., 1993. A model for the functioning of the striatum. *Biol Cybern* 68, 535–544.
- Corbit, L.H., Janak, P.H., 2007. Inactivation of the lateral but not medial dorsal striatum eliminates the excitatory impact of Pavlovian stimuli on instrumental responding. *Journal of Neuroscience* 27, 13977–13981. doi:10.1523/JNEUROSCI.4097-07.2007
- Corbit, L.H., Muir, J.L., Balleine, B.W., 2001. The role of the nucleus accumbens in instrumental conditioning: Evidence of a functional dissociation between accumbens core and shell. *Journal of Neuroscience* 21, 3251–3260.
- Costa, R.M., Cohen, D., Nicoletis, M.A.L., 2004. Differential corticostriatal plasticity during fast and slow motor skill learning in mice. *Curr Biol* 14, 1124–1134. doi:10.1016/j.cub.2004.06.053
- Coull, J.T., Cheng, R.-K., Meck, W.H., 2011. Neuroanatomical and neurochemical substrates of timing. *Neuropsychopharmacology* 36, 3–25. doi:10.1038/npp.2010.113
- Coull, J.T., Cheng, R.-K., Meck, W.H., 2010. Neuroanatomical and Neurochemical Substrates of Timing. *Neuropsychopharmacology* 36, 3–25. doi:10.1038/npp.2010.113
- Courtemanche, R., Fujii, N., Graybiel, A.M., 2003. Synchronous, focally modulated beta-band oscillations characterize local field potential activity in the striatum of awake behaving monkeys. *Journal of Neuroscience* 23, 11741–11752.
- Cowan, R.L., Wilson, C.J., 1994. Spontaneous firing patterns and axonal projections of single corticostriatal neurons in the rat medial agranular cortex. *J. Neurophysiol.*
- Cowan, R.L., Wilson, C.J., Emson, P.C., Heizmann, C.W., 1990. Parvalbumin-containing GABAergic interneurons in the rat neostriatum. *J. Comp. Neurol.* 302, 197–205. doi:10.1002/cne.903020202
- Crowe, D.A., Averbeck, B.B., Chafee, M.V., 2010. Rapid sequences of population activity patterns dynamically encode task-critical spatial information in parietal cortex. *Journal of Neuroscience* 30, 11640–11653. doi:10.1523/JNEUROSCI.0954-10.2010

- Crowe, D.A., Zarco, W., Bartolo, R., Merchant, H., 2014. Dynamic representation of the temporal and sequential structure of rhythmic movements in the primate medial premotor cortex. *Journal of Neuroscience* 34, 11972–11983. doi:10.1523/JNEUROSCI.2177-14.2014
- Czubayko, U., Plenz, D., 2002. Fast synaptic transmission between striatal spiny projection neurons. *Proc. Natl. Acad. Sci. U.S.A.* 99, 15764–15769. doi:10.1073/pnas.242428599
- Damodaran, S., Evans, R.C., Blackwell, K.T., 2014. Synchronized firing of fast-spiking interneurons is critical to maintain balanced firing between direct and indirect pathway neurons of the striatum. 111, 836–848. doi:10.1152/jn.00382.2013
- Dietrich, A., Allen, J.D., 1998. Functional dissociation of the prefrontal cortex and the hippocampus in timing behavior. *Behav. Neurosci.* 112, 1043–1047.
- Dudai, Y., Carruthers, M., 2005. The Janus face of Mnemosyne. *Nature Neuroscience* 434, 567–567. doi:10.1038/434567a
- Engel, A.K., König, P., Kreiter, A.K., Singer, W., 1991. Interhemispheric synchronization of oscillatory neuronal responses in cat visual cortex. *Science* 252, 1177–1179.
- English, D.F., Ibanez-Sandoval, O., Stark, E., Tecuapetla, F., Buzsáki, G., Deisseroth, K., Tepper, J.M., Koos, T., 2012. GABAergic circuits mediate the reinforcement-related signals of striatal cholinergic interneurons. *Nature Neuroscience* 15, 123–130. doi:10.1038/nn.2984
- Fanselow, M.S., Wassum, K.M., 2016. The Origins and Organization of Vertebrate Pavlovian Conditioning. *Cold Spring Harb Perspect Biol* 8, a021717–28. doi:10.1101/cshperspect.a021717
- Featherstone, R.E., McDonald, R.J., 2005. Lesions of the dorsolateral striatum impair the acquisition of a simplified stimulus-response dependent conditional discrimination task. *Neuroscience* 136, 387–395. doi:10.1016/j.neuroscience.2005.08.021
- Feierstein, C.E., Quirk, M.C., Uchida, N., Sosulski, D.L., Mainen, Z.F., 2006. Representation of Spatial Goals in Rat Orbitofrontal Cortex. *Neuron* 51, 495–507. doi:10.1016/j.neuron.2006.06.032
- Finch, D.M., 1996. Neurophysiology of converging synaptic inputs from the rat prefrontal cortex, amygdala, midline thalamus, and hippocampal formation onto single neurons of the caudate/putamen and nucleus accumbens. *Hippocampus* 6, 495–512. doi:10.1002/(SICI)1098-1063(1996)6:5<495::AID-HIPO3>3.0.CO;2-I
- Fino, E., Venance, L., 2011. Spike-timing dependent plasticity in striatal interneurons. *Neuropharmacology* 60, 780–788. doi:10.1016/j.neuropharm.2011.01.023
- Flaherty, A.W., Graybiel, A.M., 1994. Input-output organization of the sensorimotor striatum in the squirrel monkey. *J Neurosci* 14, 599–610.
- Flaherty, A.W., Graybiel, A.M., 1991. Corticostriatal transformations in the primate somatosensory system. Projections from physiologically mapped body-part representations. *J. Neurophysiol.* 66, 1249–1263.

- Fregnac, Y., 2003. Hebbian synaptic plasticity, in: *The Handbook of Brain Theory and Neural Networks*. MIT Press, p. 1290.
- Fujisawa, S., Amarasingham, A., Harrison, M.T., Buzsáki, G., 2008. Behavior-dependent short-term assembly dynamics in the medial prefrontal cortex. *Nature Neuroscience* 11, 823–833. doi:10.1038/nn.2134
- Fukai, T., Tanaka, S., 1997. A simple neural network exhibiting selective activation of neuronal ensembles: from winner-take-all to winners-share-all. *Neural Comput* 9, 77–97.
- Fukuda, T., 2009. Network Architecture of Gap Junction-Coupled Neuronal Linkage in the Striatum. *J Neurosci* 29, 1235–1243. doi:10.1523/JNEUROSCI.4418-08.2009
- Furuyashiki, T., HOLLAND, P.C., Gallagher, M., 2008. Rat Orbitofrontal Cortex Separately Encodes Response and Outcome Information during Performance of Goal-Directed Behavior. *J Neurosci* 28, 5127–5138. doi:10.1523/JNEUROSCI.0319-08.2008
- Fusi, S., Miller, E.K., Rigotti, M., 2016. Why neurons mix: high dimensionality for higher cognition. *Curr Opin Neurobiol* 37, 66–74. doi:10.1016/j.conb.2016.01.010
- Fuster, J.M., 2001. The prefrontal cortex--an update: time is of the essence. *Neuron* 30, 319–333.
- Gage, G.J., Stoetzner, C.R., Wiltschko, A.B., Berke, J.D., 2010. Selective Activation of Striatal Fast-Spiking Interneurons during Choice Execution. *Neuron* 67, 466–479. doi:10.1016/j.neuron.2010.06.034
- Genovesio, A., Tsujimoto, S., Wise, S.P., 2009. Feature- and Order-Based Timing Representations in the Frontal Cortex. *Neuron* 63, 254–266. doi:10.1016/j.neuron.2009.06.018
- Gerfen, C.R., Surmeier, D.J., 2011. Modulation of striatal projection systems by dopamine. *34*, 441–466. doi:10.1146/annurev-neuro-061010-113641
- Gibbon, J., 1977. Scalar expectancy theory and Weber's law in animal timing. *Psychol Rev.*
- Gittis, A.H., Hang, G.B., LaDow, E.S., Shoenfeld, L.R., Atallah, B.V., Finkbeiner, S., Kreitzer, A.C., 2011. Rapid target-specific remodeling of fast-spiking inhibitory circuits after loss of dopamine. *Neuron* 71, 858–868. doi:10.1016/j.neuron.2011.06.035
- Gittis, A.H., Nelson, A.B., Thwin, M.T., Palop, J.J., Kreitzer, A.C., 2010. Distinct roles of GABAergic interneurons in the regulation of striatal output pathways. *Journal of Neuroscience* 30, 2223–2234. doi:10.1523/JNEUROSCI.4870-09.2010
- Goel, A., Buonomano, D.V., 2014. Timing as an intrinsic property of neural networks: evidence from in vivo and in vitro experiments. *Philos. Trans. R. Soc. Lond., B, Biol. Sci.* 369, 20120460. doi:10.1098/rstb.2012.0460
- Goto, Y., Grace, A.A., 2005. Dopaminergic modulation of limbic and cortical drive of nucleus accumbens in goal-directed behavior. *Nature Neuroscience* 8, 805–812. doi:10.1038/nn1471

- Gouvêa, T.S., Monteiro, T., Motiwala, A., Soares, S., Machens, C., Paton, J.J., 2015. Striatal dynamics explain duration judgments. *Elife* 4, 2473. doi:10.7554/eLife.11386
- Graybiel, A.M., 2008. Habits, rituals, and the evaluative brain. *31*, 359–387. doi:10.1146/annurev.neuro.29.051605.112851
- Graybiel, A.M., 2000. The basal ganglia. *Curr Biol* 10, R509–11.
- Graybiel, A.M., Aosaki, T., Flaherty, A.W., Kimura, M., 1994. The basal ganglia and adaptive motor control. *Science* 265, 1826–1831.
- Gremel, C.M., Costa, R.M., 2013. Premotor cortex is critical for goal-directed actions. *Frontiers in Computational Neuroscience* 7, 110. doi:10.3389/fncom.2013.00110
- Haber, S.N., Fudge, J.L., McFarland, N.R., 2000. Striatonigrostriatal pathways in primates form an ascending spiral from the shell to the dorsolateral striatum. *Journal of Neuroscience* 20, 2369–2382.
- Hahnloser, R.H.R., Kozhevnikov, A.A., Fee, M.S., 2002. An ultra-sparse code underlies the generation of neural sequences in a songbird. *Nature* 419, 65–70. doi:10.1038/nature00974
- Harrington, D.L., Boyd, L.A., Mayer, A.R., Sheltraw, D.M., Lee, R.R., Huang, M., Rao, S.M., 2004. Neural representation of interval encoding and decision making. *Cognitive Brain Research* 21, 193–205. doi:10.1016/j.cogbrainres.2004.01.010
- Harvey, C.D., Coen, P., Tank, D.W., 2012. Choice-specific sequences in parietal cortex during a virtual-navigation decision task. *Nature* 484, 62–68. doi:10.1038/nature10918
- Hintiryan, H., Foster, N.N., Bowman, I., Bay, M., Song, M.Y., Gou, L., Yamashita, S., Bienkowski, M.S., Zingg, B., Zhu, M., Yang, X.W., Shih, J.C., Toga, A.W., Dong, H.-W., 2016. The mouse cortico-striatal projectome. *Nature Neuroscience* 19, 1100–1114. doi:10.1038/nn.4332
- Hjorth, J., Blackwell, K.T., Kotaleski, J.H., 2009. Gap junctions between striatal fast-spiking interneurons regulate spiking activity and synchronization as a function of cortical activity. *Journal of Neuroscience* 29, 5276–5286. doi:10.1523/JNEUROSCI.6031-08.2009
- Hoffer, Z.S., Alloway, K.D., 2001. Organization of corticostriatal projections from the vibrissal representations in the primary motor and somatosensory cortical areas of rodents. *J. Comp. Neurol.* 439, 87–103. doi:10.1002/cne.1337
- Hopfield, J.J., 1982. Neural networks and physical systems with emergent collective computational abilities. *Proc. Natl. Acad. Sci. U.S.A.* 79, 2554–2558.
- Hsu, C.W., Lin, C.J., 2002. A comparison of methods for multiclass support vector machines. *Neural Networks*.
- Humphries, M.D., Stewart, R.D., Gurney, K.N., 2006. A Physiologically Plausible Model of Action Selection and Oscillatory Activity in the Basal Ganglia. *J Neurosci* 26, 12921–12942. doi:10.1523/JNEUROSCI.3486-06.2006

- Humphries, M.D., Wood, R., Gurney, K., 2009. Dopamine-modulated dynamic cell assemblies generated by the GABAergic striatal microcircuit. *Neural Netw* 22, 1174–1188. doi:10.1016/j.neunet.2009.07.018
- Isaacson, J.S., Scanziani, M., 2011. How Inhibition Shapes Cortical Activity. *Neuron* 72, 231–243. doi:10.1016/j.neuron.2011.09.027
- Jadhav, S.P., Kemere, C., German, P.W., Frank, L.M., 2012. Awake hippocampal sharp-wave ripples support spatial memory. *Science* 336, 1454–1458. doi:10.1126/science.1217230
- Jaeger, D., Kita, H., Wilson, C.J., 1994. Surround inhibition among projection neurons is weak or nonexistent in the rat neostriatum. *J. Neurophysiol.* 72, 2555–2558.
- Jaeger, H., Haas, H., 2004. Harnessing nonlinearity: predicting chaotic systems and saving energy in wireless communication. *Science* 304, 78–80. doi:10.1126/science.1091277
- Jaidar, O., Carrillo-Reid, L., Hernandez, A., Drucker-Colin, R., Bargas, J., Hernandez-Cruz, A., 2010. Dynamics of the Parkinsonian Striatal Microcircuit: Entrainment into a Dominant Network State. *J Neurosci* 30, 11326–11336. doi:10.1523/JNEUROSCI.1380-10.2010
- Janssen, P., Shadlen, M.N., 2005. A representation of the hazard rate of elapsed time in macaque area LIP. *Nature Neuroscience* 8, 234–241. doi:10.1038/nn1386
- Jin, D.Z., Fujii, N., Graybiel, A.M., 2009. Neural representation of time in cortico-basal ganglia circuits. *Proc Natl Acad Sci USA* 106, 19156–19161. doi:10.1073/pnas.0909881106
- Jin, X., Costa, R.M., 2010. Start/stop signals emerge in nigrostriatal circuits during sequence learning. *Nature* 466, 457–462. doi:10.1038/nature09263
- Jo, S., Kim, K.-U., Lee, D., Jung, M.W., 2013. Effect of orbitofrontal cortex lesions on temporal discounting in rats. *Behav Brain Res* 245, 22–28. doi:10.1016/j.bbr.2013.02.014
- Joel, D., Weiner, I., 2000. The connections of the dopaminergic system with the striatum in rats and primates: an analysis with respect to the functional and compartmental organization of the striatum. *Neuroscience* 96, 451–474.
- Joel, D., Weiner, I., 1994. The organization of the basal ganglia-thalamocortical circuits: open interconnected rather than closed segregated. *Neuroscience* 63, 363–379.
- Jog, M.S., Kubota, Y., Connolly, C.I., Hillegaart, V., Graybiel, A.M., 1999. Building neural representations of habits. *Science* 286, 1745–1749. doi:10.1126/science.286.5445.1745
- Kasanez, F., Riquelme, L.A., O'Donnell, P., Murer, M.G., 2006. Turning off cortical ensembles stops striatal Up states and elicits phase perturbations in cortical and striatal slow oscillations in rat in vivo. *J Physiol (Lond)* 577, 97–113. doi:10.1113/jphysiol.2006.113050
- Kawaguchi, Y., Wilson, C.J., Augood, S.J., Emson, P.C., 1995. Striatal interneurons: chemical, physiological and morphological characterization. *Trends Neurosci.*
- Kawaguchi, Y., Wilson, C.J., Emson, P.C., 1990. Projection subtypes of rat neostriatal matrix cells revealed by intracellular injection of biocytin. *J Neurosci* 10, 3421–3438.

- Kawai, R., Markman, T., Poddar, R., Ko, R., Fantana, A.L., Dhawale, A.K., Kampff, A.R., Ölveczky, B.P., 2015. Motor cortex is required for learning but not for executing a motor skill. *Neuron* 86, 800–812. doi:10.1016/j.neuron.2015.03.024
- Kemp, J.M., Powell, T.P., 1971. The site of termination of afferent fibres in the caudate nucleus. *Philosophical Transactions of the Royal Society B: Biological Sciences* 262, 413–427.
- Kemp, J.M., Powell, T.P., 1970. The cortico-striate projection in the monkey. *Brain* 93, 525–546.
- Kim, J., Ghim, J.-W., Lee, J.H., Jung, M.W., 2013. Neural correlates of interval timing in rodent prefrontal cortex. *Journal of Neuroscience* 33, 13834–13847. doi:10.1523/JNEUROSCI.1443-13.2013
- Kim, J., Jung, A.H., Byun, J., Jo, S., Jung, M.W., 2009. Inactivation of medial prefrontal cortex impairs time interval discrimination in rats. *Front. Behav. Neurosci.* 3, 38. doi:10.3389/neuro.08.038.2009
- Kincaid, A.E., Wilson, C.J., 1996. Corticostriatal innervation of the patch and matrix in the rat neostriatum. *J. Comp. Neurol.* 374, 578–592. doi:10.1002/(SICI)1096-9861(19961028)374:4<578::AID-CNE7>3.0.CO;2-Z
- Kincaid, A.E., Zheng, T., Wilson, C.J., 1998. Connectivity and convergence of single corticostriatal axons. *J Neurosci* 18, 4722–4731.
- Kita, H., Kitai, S.T., 1988. Glutamate decarboxylase immunoreactive neurons in rat neostriatum: their morphological types and populations. *Brain Res.* 447, 346–352.
- Kita, H., Kosaka, T., Heizmann, C.W., 1990. Parvalbumin-immunoreactive neurons in the rat neostriatum: a light and electron microscopic study. *Brain Res.* 536, 1–15.
- Knerr, S., Personnaz, L., Dreyfus, G., 1990. Single-layer learning revisited: a stepwise procedure for building and training a neural network. *Neurocomputing.*
- Knowlton, B.J., Mangels, J.A., Squire, L.R., 1996. A neostriatal habit learning system in humans. *Science* 273, 1399–1402.
- Ko, H., Hofer, S.B., Pichler, B., Buchanan, K.A., Sjöström, P.J., Mrsic-Flogel, T.D., 2011. Functional specificity of local synaptic connections in neocortical networks. *Nature* 473, 87–91. doi:10.1038/nature09880
- Ko, H., Mrsic-Flogel, T.D., Hofer, S.B., 2014. Emergence of feature-specific connectivity in cortical microcircuits in the absence of visual experience. *Journal of Neuroscience* 34, 9812–9816. doi:10.1523/JNEUROSCI.0875-14.2014
- Kocsis, J.D., Sugimori, M., Kitai, S.T., 1977. Convergence of excitatory synaptic inputs to caudate spiny neurons. *Brain Res.* 124, 403–413.
- Komiyama, T., Sato, T.R., O'Connor, D.H., Zhang, Y.-X., Huber, D., Hooks, B.M., Gabitto, M., Svoboda, K., 2010. Learning-related fine-scale specificity imaged in motor cortex circuits of behaving mice. *Nature* 464, 1182–1186. doi:10.1038/nature08897

- Koos, T., Tepper, J.M., 2002. Dual cholinergic control of fast-spiking interneurons in the neostriatum. *Journal of Neuroscience* 22, 529–535.
- Koós, T., Tepper, J.M., 1999. Inhibitory control of neostriatal projection neurons by GABAergic interneurons. *Nature Neuroscience* 2, 467–472. doi:10.1038/8138
- Koralek, A.C., Costa, R.M., Carmena, J.M., 2013. Temporally precise cell-specific coherence develops in corticostriatal networks during learning. *Neuron* 79, 865–872. doi:10.1016/j.neuron.2013.06.047
- Koralek, A.C., Jin, X., Long, J.D., Costa, R.M., Carmena, J.M., 2012. Corticostriatal plasticity is necessary for learning intentional neuroprosthetic skills. *Nature* 483, 331–335. doi:10.1038/nature10845
- Kraus, B.J., Robinson, R.J., White, J.A., Eichenbaum, H., Hasselmo, M.E., 2013. Hippocampal “time cells”: time versus path integration. *Neuron* 78, 1090–1101. doi:10.1016/j.neuron.2013.04.015
- Kravitz, A.V., Tye, L.D., Kreitzer, A.C., 2012. Distinct roles for direct and indirect pathway striatal neurons in reinforcement. *Nature Neuroscience* 15, 816–818. doi:10.1038/nn.3100
- Kreitzer, A.C., Berke, J.D., 2011. Investigating striatal function through cell-type-specific manipulations. *Neuroscience* 198, 19–26. doi:10.1016/j.neuroscience.2011.08.018
- Kreitzer, A.C., Malenka, R.C., 2008. Striatal plasticity and basal ganglia circuit function. *Neuron* 60, 543–554. doi:10.1016/j.neuron.2008.11.005
- Kreitzer, A.C., Malenka, R.C., 2007. Endocannabinoid-mediated rescue of striatal LTD and motor deficits in Parkinson's disease models. *Nature* 445, 643–647. doi:10.1038/nature05506
- Kreßel, U.H.G., 1999. Pairwise classification and support vector machines, in: *Advances in Kernel Methods*. MIT Press, pp. 255–268.
- Kubota, Y., Liu, J., Hu, D., DeCoteau, W.E., Eden, U.T., Smith, A.C., Graybiel, A.M., 2009. Stable encoding of task structure coexists with flexible coding of task events in sensorimotor striatum. *Neuron* 102, 2142–2160. doi:10.1016/j.neuron.2009.05.022
- Künzle, H., 1977. Projections from the primary somatosensory cortex to basal ganglia and thalamus in the monkey. *Exp Brain Res* 30, 481–492.
- Künzle, H., 1975. Bilateral projections from precentral motor cortex to the putamen and other parts of the basal ganglia. An autoradiographic study in *Macaca fascicularis*. *Brain Res.* 88, 195–209.
- la Rocha, de, J., Doiron, B., Shea-Brown, E., Josić, K., Reyes, A., 2007. Correlation between neural spike trains increases with firing rate. *Nature* 448, 802–806. doi:10.1038/nature06028
- Laje, R., Buonomano, D.V., 2013. Robust timing and motor patterns by taming chaos in recurrent neural networks. *Nature Neuroscience* 16, 925–933. doi:10.1038/nn.3405

- Lansink, C.S., Jackson, J.C., Lankelma, J.V., Ito, R., Robbins, T.W., Everitt, B.J., Pennartz, C.M.A., 2012. Reward cues in space: commonalities and differences in neural coding by hippocampal and ventral striatal ensembles. *Journal of Neuroscience* 32, 12444–12459. doi:10.1523/JNEUROSCI.0593-12.2012
- Lau, T., Gage, G.J., Berke, J.D., Zochowski, M., 2010. Local dynamics of gap-junction-coupled interneuron networks. *Phys Biol* 7, 16015. doi:10.1088/1478-3975/7/1/016015
- Lee, D., Port, N.L., Kruse, W., Georgopoulos, A.P., 1998. Variability and correlated noise in the discharge of neurons in motor and parietal areas of the primate cortex. *J Neurosci* 18, 1161–1170.
- Lee, K., Holley, S.M., Shobe, J.L., Chong, N.C., Cepeda, C., Levine, M.S., Masmanidis, S.C., 2017. Parvalbumin Interneurons Modulate Striatal Output and Enhance Performance during Associative Learning. *Neuron* 93, 1451–1463.e4. doi:10.1016/j.neuron.2017.02.033
- Lee, S.-H., Kwan, A.C., Zhang, S., Phoumthipphavong, V., Flannery, J.G., Masmanidis, S.C., Taniguchi, H., Huang, Z.J., Zhang, F., Boyden, E.S., Deisseroth, K., Dan, Y., 2012. Activation of specific interneurons improves V1 feature selectivity and visual perception. *Nature* 488, 379–383. doi:10.1038/nature11312
- Lei, W., Jiao, Y., Del Mar, N., Reiner, A., 2004. Evidence for differential cortical input to direct pathway versus indirect pathway striatal projection neurons in rats. *Journal of Neuroscience* 24, 8289–8299. doi:10.1523/JNEUROSCI.1990-04.2004
- Levesque, M., Charara, A., Gagnon, S., Parent, A., Deschenes, M., 1996. Corticostriatal projections from layer V cells in rat are collaterals of long-range corticofugal axons. *Brain Res.* 709, 311–315.
- Levesque, M., Parent, A., 1998. Axonal arborization of corticostriatal and corticothalamic fibers arising from prelimbic cortex in the rat. *Cereb. Cortex* 8, 602–613.
- Li, N., Daie, K., Svoboda, K., Druckmann, S., 2016. Robust neuronal dynamics in premotor cortex during motor planning. *Nature*. doi:10.1038/nature17643
- Long, M.A., Jin, D.Z., Fee, M.S., 2010. Support for a synaptic chain model of neuronal sequence generation. *Nature* 468, 394–399. doi:10.1038/nature09514
- Lovinger, D.M., Tyler, E.C., Merritt, A., 1993. Short- and long-term synaptic depression in rat neostriatum. *J. Neurophysiol.* 70, 1937–1949.
- Maass, W., Natschläger, T., Markram, H., 2002. Real-time computing without stable states: a new framework for neural computation based on perturbations. *Neural Comput* 14, 2531–2560. doi:10.1162/089976602760407955
- Malach, R., Graybiel, A.M., 1986. Mosaic architecture of the somatic sensory-recipient sector of the cat's striatum. *J Neurosci* 6, 3436–3458.
- Mallet, N., Le Moine, C., Charpier, S., Gonon, F., 2005. Feedforward inhibition of projection neurons by fast-spiking GABA interneurons in the rat striatum in vivo. *Journal of Neuroscience* 25, 3857–3869. doi:10.1523/JNEUROSCI.5027-04.2005

- Mante, V., Sussillo, D., Shenoy, K.V., Newsome, W.T., 2013. Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature* 503, 78–84. doi:10.1038/nature12742
- Matell, M.S., Meck, W.H., 2004. Cortico-striatal circuits and interval timing: coincidence detection of oscillatory processes. *Cognitive Brain Research* 21, 139–170. doi:10.1016/j.cogbrainres.2004.06.012
- Matell, M.S., Meck, W.H., Nicolelis, M.A.L., 2003. Interval timing and the encoding of signal duration by ensembles of cortical and striatal neurons. *Behav. Neurosci.* 117, 760–773. doi:10.1037/0735-7044.117.4.760
- Mathis, M.W., Mathis, A., Uchida, N., 2017. Somatosensory Cortex Plays an Essential Role in Forelimb Motor Adaptation in Mice. *Neuron* 93, 1493–1503.e6. doi:10.1016/j.neuron.2017.02.049
- Mauk, M.D., Buonomano, D.V., 2004. The neural basis of temporal processing. *27*, 307–340. doi:10.1146/annurev.neuro.27.070203.144247
- Mauk, M.D., Donegan, N.H., 1997. A model of Pavlovian eyelid conditioning based on the synaptic organization of the cerebellum. *Learning & Memory* 4, 130–158.
- McGeorge, A.J., Faull, R.L., 1989. The organization of the projection from the cerebral cortex to the striatum in the rat. *Neuroscience* 29, 503–537.
- Meck, W.H., 2006. Neuroanatomical localization of an internal clock: A functional link between mesolimbic, nigrostriatal, and mesocortical dopaminergic systems. *Brain Res.* 1109, 93–107. doi:10.1016/j.brainres.2006.06.031
- Meck, W.H., 1996. Neuropharmacology of timing and time perception. *Brain Res Cogn Brain Res* 3, 227–242.
- Meck, W.H., 1983. Selective adjustment of the speed of internal clock and memory processes. *J Exp Psychol Anim Behav Process* 9, 171–201.
- Medina, J.F., Garcia, K.S., Nores, W.L., Taylor, N.M., Mauk, M.D., 2000. Timing mechanisms in the cerebellum: testing predictions of a large-scale computer simulation. *J Neurosci* 20, 5516–5525.
- Mello, G.B.M., Soares, S., Paton, J.J., 2015. A scalable population code for time in the striatum. *Curr Biol* 25, 1113–1122. doi:10.1016/j.cub.2015.02.036
- Merchant, H., Harrington, D.L., Meck, W.H., 2013a. Neural Basis of the Perception and Estimation of Time 36, 313–336. doi:10.1146/annurev-neuro-062012-170349
- Merchant, H., Perez, O., Zarco, W., Gamez, J., 2013b. Interval Tuning in the Primate Medial Premotor Cortex as a General Timing Mechanism. *J Neurosci* 33, 9082–9096. doi:10.1523/JNEUROSCI.5513-12.2013
- Merchant, H., Pérez, O., Bartolo, R., Méndez, J.C., Mendoza, G., Gámez, J., Yc, K., Prado, L., 2015. Sensorimotor neural dynamics during isochronous tapping in the medial premotor cortex of the macaque. *Eur J Neurosci* 41, 586–602. doi:10.1111/ejn.12811

- Merchant, H., Zarco, W., Pérez, O., Prado, L., Bartolo, R., 2011. Measuring time with different neural chronometers during a synchronization-continuation task. *Proc Natl Acad Sci USA* 108, 19784–19789. doi:10.1073/pnas.1112933108
- Milner, B., Squire, L.R., Kandel, E.R., 1998. Cognitive neuroscience and the study of memory. *Neuron* 20, 445–468.
- Mink, J.W., 1996. The basal ganglia: focused selection and inhibition of competing motor programs. *Prog. Neurobiol.*
- Mitchell, J.F., Sundberg, K.A., Reynolds, J.H., 2009. Spatial attention decorrelates intrinsic activity fluctuations in macaque area V4. *Neuron* 63, 879–888. doi:10.1016/j.neuron.2009.09.013
- Mogenson, G.J., Jones, D.L., Yim, C.Y., 1980. From motivation to action: functional interface between the limbic system and the motor system. *Prog. Neurobiol.* 14, 69–97.
- Moorman, D.E., Aston-Jones, G., 2014. Orbitofrontal Cortical Neurons Encode Expectation-Driven Initiation of Reward-Seeking. *J Neurosci* 34, 10234–10246. doi:10.1523/JNEUROSCI.3216-13.2014
- Morris, R.G., Garrud, P., Rawlins, J.N., O'Keefe, J., 1982. Place navigation impaired in rats with hippocampal lesions. *Nature* 297, 681–683.
- Murray, J.M., 2017. Recurrent inhibition in striatum enables transfer of time-flexible skills to basal ganglia. *bioRxiv*. doi:10.1101/110072
- Nelson, A.B., Hammack, N., Yang, C.F., Shah, N.M., Seal, R.P., Kreitzer, A.C., 2014. Striatal cholinergic interneurons Drive GABA release from dopamine terminals. *Neuron* 82, 63–70. doi:10.1016/j.neuron.2014.01.023
- Nicola, S.M., Yun, I.A., Wakabayashi, K.T., Fields, H.L., 2004. Cue-evoked firing of nucleus accumbens neurons encodes motivational significance during a discriminative stimulus task. *J Neurosci* 24, 1840–1865. doi:10.1523/jn.00657.2003
- Nirenberg, S., Carcieri, S.M., Jacobs, A.L., Latham, P.E., 2001. Retinal ganglion cells act largely as independent encoders. *Nature* 411, 698–701. doi:10.1038/35079612
- Nisenbaum, E.S., Xu, Z.C., Wilson, C.J., 1994. Contribution of a slowly inactivating potassium current to the transition to firing of neostriatal spiny projection neurons. *J. Neurophysiol.* 71, 1174–1189.
- Oh, S.W., Harris, J.A., Ng, L., Winslow, B., Cain, N., Mihalas, S., Wang, Q., Lau, C., Kuan, L., Henry, A.M., Mortrud, M.T., Ouellette, B., Nguyen, T.N., Sorensen, S.A., Slaughterbeck, C.R., Wakeman, W., Li, Y., Feng, D., Ho, A., Nicholas, E., Hirokawa, K.E., Bohn, P., Joines, K.M., Peng, H., Hawrylycz, M.J., Phillips, J.W., Hohmann, J.G., Wahnoutka, P., Gerfen, C.R., Koch, C., Bernard, A., Dang, C., Jones, A.R., Zeng, H., 2014. A mesoscale connectome of the mouse brain. *Nature* 508, 207–214. doi:10.1038/nature13186
- Onoe, H., Komori, M., Onoe, K., Takechi, H., Tsukada, H., Watanabe, Y., 2001. Cortical networks recruited for time perception: a monkey positron emission tomography (PET)

- study. *Neuroimage* 13, 37–45. doi:10.1006/nimg.2000.0670
- Oshio, K.-I., Chiba, A., Inase, M., 2008. Temporal filtering by prefrontal neurons in duration discrimination. *European Journal of Neuroscience* 28, 2333–2343. doi:10.1111/j.1460-9568.2008.06509.x
- O’Hare, J.K., Ade, K.K., Sukharnikova, T., Van Hooser, S.D., Palmeri, M.L., Yin, H.H., Calakos, N., 2016. Pathway-Specific Striatal Substrates for Habitual Behavior. *Neuron* 89, 472–479. doi:10.1016/j.neuron.2015.12.032
- Packard, M.G., Hirsh, R., White, N.M., 1989. Differential effects of fornix and caudate nucleus lesions on two radial maze tasks: evidence for multiple memory systems. *J Neurosci* 9, 1465–1472.
- Parkinson, J.A., Dalley, J.W., Cardinal, R.N., Bamford, A., Fehner, B., Lachenal, G., Rudarakanchana, N., Halkerston, K.M., Robbins, T.W., Everitt, B.J., 2002. Nucleus accumbens dopamine depletion impairs both acquisition and performance of appetitive Pavlovian approach behaviour: implications for mesoaccumbens dopamine function. *Behav Brain Res* 137, 149–163.
- Parthasarathy, H.B., Graybiel, A.M., 1997. Cortically driven immediate-early gene expression reflects modular influence of sensorimotor cortex on identified striatal neurons in the squirrel monkey. *J Neurosci* 17, 2477–2491.
- Pastalkova, E., Itskov, V., Amarasingham, A., Buzsáki, G., 2008. Internally generated cell assembly sequences in the rat hippocampus. *Science* 321, 1322–1327. doi:10.1126/science.1159775
- Perin, R., Berger, T.K., Markram, H., 2011. A synaptic organizing principle for cortical neuronal groups. *Proc Natl Acad Sci USA* 108, 5419–5424. doi:10.1073/pnas.1016051108
- Pérez-Ramírez, M.B., Laville, A., Tapia, D., Duhne, M., Lara-González, E., Bargas, J., Galarraga, E., 2015. KV7 Channels Regulate Firing during Synaptic Integration in GABAergic Striatal Neurons. *Neural Plast.* 2015, 472676. doi:10.1155/2015/472676
- Phillips, E., Hasenstaub, A.R., 2016. Asymmetric effects of activating and inactivating cortical interneurons. *Elife*. doi:10.7554/eLife.18383.001
- Pinto, L., Dan, Y., 2015. Cell-Type-Specific Activity in Prefrontal Cortex during Goal-Directed Behavior. *Neuron* 1–15. doi:10.1016/j.neuron.2015.06.021
- Plenz, D., 2003. When inhibition goes incognito: feedback interaction between spiny projection neurons in striatal function. *Trends Neurosci* 26, 436–443. doi:10.1016/S0166-2236(03)00196-6
- Plenz, D., Kitai, S.T., 1998. 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. *J Neurosci* 18, 266–283.
- Plenz, D., Kitai, S.T., 1999. A basal ganglia pacemaker formed by the subthalamic nucleus and external globus pallidus. *Nature* 400, 677–682. doi:10.1038/23281

- Plotkin, J.L., Day, M., Surmeier, D.J., 2011. Synaptically driven state transitions in distal dendrites of striatal spiny neurons. *Nature Neuroscience*. doi:10.1038/nn.2848
- Ponzi, A., Wickens, J., 2012. Input dependent cell assembly dynamics in a model of the striatal medium spiny neuron network. *Front Syst Neurosci* 6. doi:10.3389/fnsys.2012.00006/abstract
- Ponzi, A., Wickens, J., 2010. Sequentially switching cell assemblies in random inhibitory networks of spiking neurons in the striatum. *Journal of Neuroscience* 30, 5894–5911. doi:10.1523/JNEUROSCI.5540-09.2010
- Popescu, A.T., Popa, D., Paré, D., 2009. Coherent gamma oscillations couple the amygdala and striatum during learning. *Nature Neuroscience* 12, 801–807. doi:10.1038/nn.2305
- Prinz, A.A., Bucher, D., Marder, E., 2004. Similar network activity from disparate circuit parameters. *Nature Neuroscience* 7, 1345–1352. doi:10.1038/nn1352
- Raichle, M.E., 2010. Two views of brain function. *Trends Cogn. Sci. (Regul. Ed.)*. doi:10.1016/j.tics.2010.01.008
- Rajan, K., Harvey, C.D., Tank, D.W., 2016. Recurrent Network Models of Sequence Generation and Memory. *Neuron* 90, 128–142. doi:10.1016/j.neuron.2016.02.009
- Ramanathan, S., Hanley, J.J., Deniau, J.-M., Bolam, J.P., 2002. Synaptic convergence of motor and somatosensory cortical afferents onto GABAergic interneurons in the rat striatum. *Journal of Neuroscience* 22, 8158–8169.
- Ramirez, S., Liu, X., Lin, P.A., Suh, J., Pignatelli, M., Redondo, R.L., Ryan, T.J., Tonegawa, S., 2013. Creating a False Memory in the Hippocampus. *Science* 341, 387–391. doi:10.1126/science.1239073
- Redgrave, P., Prescott, T.J., Gurney, K., 1999. The basal ganglia: a vertebrate solution to the selection problem? *Neuroscience* 89, 1009–1023.
- Reig, R., Silberberg, G., 2014. Multisensory integration in the mouse striatum. *Neuron* 83, 1200–1212. doi:10.1016/j.neuron.2014.07.033
- Reimer, J., Froudarakis, E., Cadwell, C.R., Yatsenko, D., Denfield, G.H., Tolias, A.S., 2014. Pupil fluctuations track fast switching of cortical states during quiet wakefulness. *Neuron* 84, 355–362. doi:10.1016/j.neuron.2014.09.033
- Rigotti, M., Barak, O., Warden, M.R., Wang, X.-J., Daw, N.D., Miller, E.K., Fusi, S., 2013. The importance of mixed selectivity in complex cognitive tasks. *Nature* 497, 585–590. doi:10.1038/nature12160
- Ringach, D.L., 2009. Spontaneous and driven cortical activity: implications for computation. *Curr Opin Neurobiol* 19, 439–444. doi:10.1016/j.conb.2009.07.005
- Roesch, M.R., Taylor, A.R., Schoenbaum, G., 2006. Encoding of Time-Discounted Rewards in Orbitofrontal Cortex Is Independent of Value Representation. *Neuron* 51, 509–520. doi:10.1016/j.neuron.2006.06.027

- Roitman, M.F., Wheeler, R.A., Carelli, R.M., 2005. Nucleus accumbens neurons are innately tuned for rewarding and aversive taste stimuli, encode their predictors, and are linked to motor output. *Neuron* 45, 587–597. doi:10.1016/j.neuron.2004.12.055
- Rosenblatt, F., 1958. The perceptron: A probabilistic model for information storage and organization in the brain. *Psychol Rev.*
- Rueda-Orozco, P.E., Robbe, D., 2015. The striatum multiplexes contextual and kinematic information to constrain motor habits execution. *Nature Neuroscience* 18, 453–460. doi:10.1038/nn.3924
- Russo, G., Nieus, T.R., Maggi, S., Taverna, S., 2013. Dynamics of action potential firing in electrically connected striatal fast-spiking interneurons. *Front Cell Neurosci* 7, 209. doi:10.3389/fncel.2013.00209
- Santos, F.J., Oliveira, R.F., Jin, X., Costa, R.M., 2015. Corticostriatal dynamics encode the refinement of specific behavioral variability during skill learning. *Elife.* doi:10.7554/eLife.09423.001
- Saunders, B.T., Robinson, T.E., 2012. The role of dopamine in the accumbens core in the expression of Pavlovian-conditioned responses. *Eur J Neurosci* 36, 2521–2532. doi:10.1111/j.1460-9568.2012.08217.x
- Schacter, D.L., Addis, D.R., 2007. Constructive memory: the ghosts of past and future. *Nature Neuroscience* 445, 27–27. doi:10.1038/445027a
- Schneidman, E., Bialek, W., Berry, M.J., 2003. Synergy, redundancy, and independence in population codes. *Journal of Neuroscience* 23, 11539–11553.
- Schultz, W., Dayan, P., Montague, P.R., 1997. A neural substrate of prediction and reward. *Science* 275, 1593–1599.
- Selemon, L.D., Goldman-Rakic, P.S., 1985. Longitudinal topography and interdigitation of corticostriatal projections in the rhesus monkey. *J Neurosci* 5, 776–794.
- Seo, M., Lee, E., Averbeck, B.B., 2012. Action selection and action value in frontal-striatal circuits. *Neuron* 74, 947–960. doi:10.1016/j.neuron.2012.03.037
- Setlow, B., Schoenbaum, G., Gallagher, M., 2003. Neural encoding in ventral striatum during olfactory discrimination learning. *Neuron* 38, 625–636.
- Shadlen, M.N., Newsome, W.T., 1998. The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *J Neurosci* 18, 3870–3896.
- Shan, Q., Ge, M., Christie, M.J., Balleine, B.W., 2014. The acquisition of goal-directed actions generates opposing plasticity in direct and indirect pathways in dorsomedial striatum. *Journal of Neuroscience* 34, 9196–9201. doi:10.1523/JNEUROSCI.0313-14.2014
- Shen, W., Flajolet, M., Greengard, P., Surmeier, D.J., 2008. Dichotomous dopaminergic control of striatal synaptic plasticity. *Science* 321, 848–851. doi:10.1126/science.1160575

- Shen, W., Hamilton, S.E., Nathanson, N.M., Surmeier, D.J., 2005. Cholinergic suppression of KCNQ channel currents enhances excitability of striatal medium spiny neurons. *Journal of Neuroscience* 25, 7449–7458. doi:10.1523/JNEUROSCI.1381-05.2005
- Shiflett, M.W., Balleine, B.W., 2011. Molecular substrates of action control in cortico-striatal circuits. *Prog. Neurobiol.* 95, 1–13. doi:10.1016/j.pneurobio.2011.05.007
- Shobe, J.L., Claar, L.D., Parhami, S., Bakhurin, K.I., Masmanidis, S.C., 2015. Brain activity mapping at multiple scales with silicon microprobes containing 1,024 electrodes. *J. Neurophysiol.* 114, 2043–2052. doi:10.1152/jn.00464.2015
- Silberberg, G., Bolam, J.P., 2015. Local and afferent synaptic pathways in the striatal microcircuitry. *Curr Opin Neurobiol* 33, 182–187. doi:10.1016/j.conb.2015.05.002
- Silva, A.J., Zhou, Y., Rogerson, T., Shobe, J., Balaji, J., 2009. Molecular and cellular approaches to memory allocation in neural circuits. *Science* 326, 391–395. doi:10.1126/science.1174519
- Simon, N.W., Wood, J., Moghaddam, B., 2015. Action-outcome relationships are represented differently by medial prefrontal and orbitofrontal cortex neurons during action execution. *J. Neurophysiol.* 114, 3374–3385. doi:10.1152/jn.00884.2015
- Smith, Y., Galvan, A., Ellender, T.J., Doig, N., Villalba, R.M., Huerta-Ocampo, I., Wichmann, T., Bolam, J.P., 2014. The thalamostriatal system in normal and diseased states. *Front Syst Neurosci* 8, 5. doi:10.3389/fnsys.2014.00005
- Somogyi, P., Bolam, J.P., Smith, A.D., 1981. Monosynaptic cortical input and local axon collaterals of identified striatonigral neurons. A light and electron microscopic study using the golgi-peroxidase transport- .... *Journal of Comparative* ....
- Sompolinsky, H., Crisanti, A., Sommers, H., 1988. Chaos in random neural networks. *Phys. Rev. Lett.* 61, 259–262. doi:10.1103/PhysRevLett.61.259
- Stern, E.A., Jaeger, D., Wilson, C.J., 1998. Membrane potential synchrony of simultaneously recorded striatal spiny neurons in vivo. *Nature* 394, 475–478. doi:10.1038/28848
- Stokes, M.G., 2015. “Activity-silent” working memory in prefrontal cortex: a dynamic coding framework. *Trends Cogn. Sci. (Regul. Ed.)* 1–12. doi:10.1016/j.tics.2015.05.004
- Stokes, M.G., Kusunoki, M., Sigala, N., Nili, H., Gaffan, D., Duncan, J., 2013. Dynamic coding for cognitive control in prefrontal cortex. *Neuron* 78, 364–375. doi:10.1016/j.neuron.2013.01.039
- Surmeier, D.J., Bargas, J., Kitai, S.T., 1988. Voltage-clamp analysis of a transient potassium current in rat neostriatal neurons. *Brain Res.* 473, 187–192.
- Sussillo, D., Abbott, L.F., 2009. Generating coherent patterns of activity from chaotic neural networks. *Neuron* 63, 544–557. doi:10.1016/j.neuron.2009.07.018
- Tai, L.-H., Lee, A.M., Benavidez, N., Bonci, A., Wilbrecht, L., 2012. Transient stimulation of distinct subpopulations of striatal neurons mimics changes in action value. *Nature*

Neuroscience 15, 1281–1289. doi:10.1038/nn.3188

- Tang, C.C., Root, D.H., Duke, D.C., Zhu, Y., Teixeira, K., Ma, S., Barker, D.J., West, M.O., 2009. Decreased Firing of Striatal Neurons Related to Licking during Acquisition and Overtraining of a Licking Task. *Journal of Neuroscience* 29, 13952–13961. doi:10.1523/JNEUROSCI.2824-09.2009
- Taverna, S., Canciani, B., Pennartz, C.M.A., 2007. Membrane properties and synaptic connectivity of fast-spiking interneurons in rat ventral striatum. *Brain Res.* 1152, 49–56. doi:10.1016/j.brainres.2007.03.053
- Taverna, S., Ilijic, E., Surmeier, D.J., 2008. Recurrent collateral connections of striatal medium spiny neurons are disrupted in models of Parkinson's disease. *Journal of Neuroscience* 28, 5504–5512. doi:10.1523/JNEUROSCI.5493-07.2008
- Taverna, S., van Dongen, Y.C., Groenewegen, H.J., Pennartz, C.M.A., 2004. Direct physiological evidence for synaptic connectivity between medium-sized spiny neurons in rat nucleus accumbens in situ. 91, 1111–1121. doi:10.1152/jn.00892.2003
- Tepper, J.M., Bolam, J.P., 2004. Functional diversity and specificity of neostriatal interneurons. *Curr Opin Neurobiol* 14, 685–692. doi:10.1016/j.conb.2004.10.003
- Tepper, J.M., Tecuapetla, F., Koos, T., Ibanez-Sandoval, O., 2010. Heterogeneity and diversity of striatal GABAergic interneurons. *Front Neuroanat* 4, 150. doi:10.3389/fnana.2010.00150
- Tepper, J.M., Wilson, C.J., Koos, T., 2008. Feedforward and feedback inhibition in neostriatal GABAergic spiny neurons. *Brain Research Reviews* 58, 272–281. doi:10.1016/j.brainresrev.2007.10.008
- Thorn, C.A., Atallah, H., Howe, M., Graybiel, A.M., 2010. Differential dynamics of activity changes in dorsolateral and dorsomedial striatal loops during learning. *Neuron* 66, 781–795. doi:10.1016/j.neuron.2010.04.036
- Thorn, C.A., Graybiel, A.M., 2014. Differential entrainment and learning-related dynamics of spike and local field potential activity in the sensorimotor and associative striatum. *Journal of Neuroscience* 34, 2845–2859. doi:10.1523/JNEUROSCI.1782-13.2014
- Threlfell, S., Lalic, T., Platt, N.J., Jennings, K.A., Deisseroth, K., Cragg, S.J., 2012. Striatal dopamine release is triggered by synchronized activity in cholinergic interneurons. *Neuron* 75, 58–64. doi:10.1016/j.neuron.2012.04.038
- Tremblay, L., Hollerman, J.R., Schultz, W., 1998. Modifications of reward expectation-related neuronal activity during learning in primate striatum. 80, 964–977.
- Tremblay, S., Pieper, F., Sachs, A., Martinez-Trujillo, J., 2015. Attentional filtering of visual information by neuronal ensembles in the primate lateral prefrontal cortex. *Neuron* 85, 202–215. doi:10.1016/j.neuron.2014.11.021
- Tunstall, M.J., Oorschot, D.E., Kean, A., Wickens, J.R., 2002. Inhibitory interactions between spiny projection neurons in the rat striatum. *J. Neurophysiol.* 88, 1263–1269. doi:10.1152/jn.00886.2001

- Vinck, M., van Wingerden, M., Womelsdorf, T., Fries, P., Pennartz, C.M.A., 2010. The pairwise phase consistency: A bias-free measure of rhythmic neuronal synchronization. *Neuroimage* 51, 112–122. doi:10.1016/j.neuroimage.2010.01.073
- Voorn, P., Vanderschuren, L.J.M.J., Groenewegen, H.J., Robbins, T.W., Pennartz, C.M.A., 2004. Putting a spin on the dorsal–ventral divide of the striatum. *Trends Neurosci* 27, 468–474. doi:10.1016/j.tins.2004.06.006
- Wang, A.Y., Miura, K., Uchida, N., 2013. The dorsomedial striatum encodes net expected return, critical for energizing performance vigor. *Nature Neuroscience* 16, 639–647. doi:10.1038/nn.3377
- Wang, Z., Kai, L., Day, M., Ronesi, J., Yin, H.H., Ding, J., Tkatch, T., Lovinger, D.M., Surmeier, D.J., 2006. Dopaminergic Control of Corticostriatal Long-Term Synaptic Depression in Medium Spiny Neurons Is Mediated by Cholinergic Interneurons. *Neuron* 50, 443–452. doi:10.1016/j.neuron.2006.04.010
- Wickens, J.R., Kotter, R., Alexander, M.E., 1995. Effects of local connectivity on striatal function: stimulation and analysis of a model. *Synapse* 20, 281–298. doi:10.1002/syn.890200402
- Wiener, M., Turkeltaub, P., Coslett, H.B., 2010. The image of time: A voxel-wise meta-analysis. *Neuroimage* 49, 1728–1740. doi:10.1016/j.neuroimage.2009.09.064
- Wilson, C.J., 2013. Active decorrelation in the basal ganglia. *Neuroscience* 250, 467–482. doi:10.1016/j.neuroscience.2013.07.032
- Wilson, C.J., Chang, H.T., Kitai, S.T., 1990. Firing patterns and synaptic potentials of identified giant aspiny interneurons in the rat neostriatum. *J Neurosci* 10, 508–519.
- Wilson, C.J., Groves, P.M., 1980. Fine structure and synaptic connections of the common spiny neuron of the rat neostriatum: a study employing intracellular injection of horseradish peroxidase. *Journal of Comparative Neurology*.
- Wilson, C.J., Groves, P.M., Kitai, S.T., Linder, J.C., 1983. Three-dimensional structure of dendritic spines in the rat neostriatum. *J Neurosci* 3, 383–388.
- Wilson, M.A., McNaughton, B.L., 1993. Dynamics of the hippocampal ensemble code for space. *Science* 261, 1055–1058.
- Xiong, Q., Znamenskiy, P., Zador, A.M., 2015. Selective corticostriatal plasticity during acquisition of an auditory discrimination task. *Nature* 521, 348–351. doi:10.1038/nature14225
- Xu, M., Zhang, S.-Y., Dan, Y., Poo, M.-M., 2014. Representation of interval timing by temporally scalable firing patterns in rat prefrontal cortex. *Proc Natl Acad Sci USA* 111, 480–485. doi:10.1073/pnas.1321314111
- Xu, S., Jiang, W., Poo, M.-M., Dan, Y., 2012. Activity recall in a visual cortical ensemble. *Nature Neuroscience* 15, 449–55– S1–2. doi:10.1038/nn.3036

- Yassin, L., Benedetti, B.L., Jouhanneau, J.-S., Wen, J.A., Poulet, J.F.A., Barth, A.L., 2010. An embedded subnetwork of highly active neurons in the neocortex. *Neuron* 68, 1043–1050. doi:10.1016/j.neuron.2010.11.029
- Yim, M.Y., Aertsen, A., Kumar, A., 2011. Significance of input correlations in striatal function. *PLoS Comp Biol* 7, e1002254. doi:10.1371/journal.pcbi.1002254
- Yin, H.H., 2016. The Basal Ganglia in Action. *neuroscientist* 1–15. doi:10.1177/1073858416654115
- Yin, H.H., 2014. How basal ganglia outputs generate behavior. *Advances in neuroscience*. doi:10.1155/2014/768313
- Yin, H.H., Knowlton, B.J., 2004. Contributions of striatal subregions to place and response learning. *Learning & Memory* 11, 459–463. doi:10.1101/lm.81004
- Yin, H.H., Knowlton, B.J., Balleine, B.W., 2005. Blockade of NMDA receptors in the dorsomedial striatum prevents action-outcome learning in instrumental conditioning. *Eur J Neurosci* 22, 505–512. doi:10.1111/j.1460-9568.2005.04219.x
- Yin, H.H., Knowlton, B.J., Balleine, B.W., 2004. Lesions of dorsolateral striatum preserve outcome expectancy but disrupt habit formation in instrumental learning. *Eur J Neurosci* 19, 181–189. doi:10.1046/j.1460-9568.2003.03095.x
- Yin, H.H., Mulcare, S.P., Hilário, M.R.F., Clouse, E., Holloway, T., Davis, M.I., Hansson, A.C., Lovinger, D.M., Costa, R.M., 2009. Dynamic reorganization of striatal circuits during the acquisition and consolidation of a skill. *Nature Neuroscience* 12, 333–341. doi:10.1038/nn.2261
- Yttri, E.A., Dudman, J.T., 2016. Opponent and bidirectional control of movement velocity in the basal ganglia. *Nature* 533, 402–406. doi:10.1038/nature17639
- Yuste, R., 2015. From the neuron doctrine to neural networks. *Nat Rev Neurosci* 16, 487–497. doi:10.1038/nrn3962
- Zandvakili, A., Kohn, A., 2015. Coordinated Neuronal Activity Enhances Corticocortical Communication. *Neuron* 87, 827–839. doi:10.1016/j.neuron.2015.07.026
- Zheng, T., Wilson, C.J., 2002. Corticostriatal combinatorics: the implications of corticostriatal axonal arborizations. *J. Neurophysiol.* 87, 1007–1017. doi:10.1152/jn.00519.2001
- Zohary, E., Shadlen, M.N., Newsome, W.T., 1994. Correlated neuronal discharge rate and its implications for psychophysical performance. *Nature* 370, 140–143. doi:10.1038/370140a0