Title
Extracting Information: Characterizing neuronal cell types in the GPh by their activity profile.

Permalink
https://escholarship.org/uc/item/8xc747cv

Authors
Li, Danxun
Stephenson-Jones, Marcus
Li, Bo

Peer reviewed|Undergraduate

Extracting Information: Characterizing neuronal cell types in the GPh by their activity profile.

Danxun Li[1,2], Marcus Stephenson-Jones[3], Bo Li[3]
[1]Undergraduate Research Program, CSHL
[2]University of California, Berkeley
[3]Cold Spring Harbor Laboratory

Abstract

The globus pallidus is a major output station for the basal ganglia, a subcortical region of the brain that is heavily implicated in action selection and decision making. A subpopulation of neurons in the internal segment (GPi) projects to the laternal habenula (LHb), often associated with the limbic system and known to encode for negative motivational value. Dysfunction in these structures have been implicated in neurological diseases, such as depression and schizophrenia, which are ultimately disruptions in the ability to evaluate environmental cues and regulate motor output. In order to gain more information about the neurons which encode for this behavior, we conducted extracellular recordings while the mice are carrying out a set of reward learning tasks and analyzed the collected spike trains. We detail here the methods of information extraction from the neuronal populations that we have classified. We also present preliminary results of their activity profile for various outcomes as well as for reward history and prediction error. With collection of information from a larger set of cells, we might be able to more definitively gain an understanding of the methods by which these neurons encode motivation, action selection and outcome evaluation.

Introduction

How is an organism motivated to act on a value-based decision? How does an organism learn to perform a specific task, to make an appropriate choice, to make the right associations? How do neurons encode initiation of movement and assign values to an expected outcome?

Basal Ganglia Circuitry

Circuits governing reward and decision-making have been the subject of extensive studies.[1] The basal ganglia emerges as a structure of primary importance in action selection and outcome evaluation, in addition to its role in motor control via the more established basal ganglia-thalamocortical circuitry.[2] The thalamus forms the central core of the brain and may be divided based on the spatial location of various nuclei, which receives input from distinct pathways and projects to well-defined cortical areas. Many critical functions such as sensory and motor mechanisms and cognitive functions are relayed via the thalamus. Specifically, the medial nucleus sends projections to the prefrontal cortex and is heavily associated with higher cognitive functions, whereas the ventrolateral nucleus of the thalamus sends information to the motor and somatosensory cortices and is associated with motor tasks.[3]

The primary function of the basal ganglia is likely to control and regulate activities of the motor and premotor cortical areas so that voluntary movements can be performed smoothly.[4, 5] Stimulating the motor cortex of monkeys at various locations results in stereotyped sequences of movements. Thus, motor control may require the activation of these elemental motor programs in the precise temporal order to accomplish a sophisticated motor plan.[6, 7] These motor programs involve inhibitory networks across various cortical and subcortical structures such that a release of this inhibition permits a motor system to become active.[8]
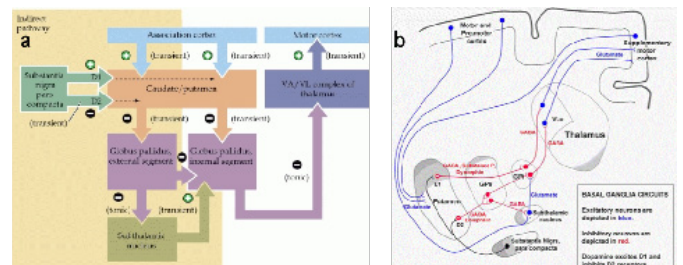

Figure 1: Basal ganglia circuits. a. Block diagram of circuits.[9] b. Schematic of circuits connecting various basal ganglia nuclei.[10]

The basal ganglia is comprised of the striatum, consisting of the caudate and putamen, the internal and external segments of the globus pallidus (GPi and GPe), the subthalamic nucleus (STN), and the substantia nigra pars reticulata and pars compacta (SNr and SNc).[3] The striatum is the input center of the basal ganglia, and receives excitatory afferents from the cerebral cortex such that along the extent of the caudate and the putamen, inputs from cortical regions vary by their relative proximities. In particular, the primary motor cortex projects mainly to the putamen, and the topography of projections is maintained in the intrinsic circuitry of the basal ganglia. The GPi serves as the major output station of the basal ganglia, along with SNr. These structures are tonically active, and impose inhibitory afferents onto the thalamus, which relays excitatory signals onto the

primary cortex associated with the relevant thalamic nuclei.[6]

There are two distinct pathways in the basal ganglia: the direct pathway, which has a net excitatory effect on its targets and the indirect pathway, which has a net inhibitory effect on its targets (Fig. 1) A model for achieving an appropriate motor response to a task is that these direct and indirect pathways coordinate the execution of stored elemental motor programs, and an upset of this balance results in motor dysfunction.

The direct pathway begins with medium spiny neurons of the striatum, which send inhibitory projections to the GPi and SNr and serves to release upper motor neurons from tonic inhibition, thus activating them. The indirect pathway starts with another population of medium spiny neurons sending inhibitory inputs to the GPe, which lifts the tonic inhibition on the excitatory neurons of the STN projecting to the GPi. The GPi is activated and the level of tonic inhibition is increased.[9] The direct and indirect pathways create a complex sequence of excitation, inhibition and disinhibition. The direct pathway is a positive feedback loop that has a net excitation of the motor cortex, whereas the indirect pathway is a negative inhibitory feedback loop, and a delicate balance is necessary for adequate performance of various motor tasks.[3]

In addition to the direct and indirect pathways, the nigrostriatal pathway, connecting the SNc to the striatum, has complex inhibitory and excitatory effects on striatal neurons. This pathway is composed of dopaminergic neurons in the SNc and largely GABAergic neurons with dopamine receptors in the striatum: activation of D1-like receptors on striatal neurons in the direct pathway induce adenylyl cyclase-dependent increase in intracellular Ca^{2+}, which stimulates neurotransmitter (GABA) release by the medium spiny neurons in response to dopamine, whereas indirect pathway striatal neurons possess primarily dopaminergic D2-like receptors, which inhibit AC activity, and thus inactivates the neuron (no GABA release).[11] Thus, activation of the dopamine neurons of the nigrostriatal pathway activates the excitatory direct pathway but inhibits the net inhibitory effect of the indirect pathway.

Encoding Action Selection by Reward Learning

Previously we described the basic basal ganglia circuitry for the activation of a specific motor program and inhibition of competing motor programs (primitive programs stored in the cortex) for a precise sequence of movements that allow the organism to adequately respond to certain environmental cues. However, one can imagine that in order to perform at maximum efficiency, a system has to adopt some measure of learning, so that a familiar cue immediately calls up a stored motor plan, which when executed yields an expected set of rewards. In the case that the outcome deviates from expectation, the system should also have methods for evaluating this error in reward prediction and perhaps adjust its established motor plan. When faced with choices, the

system should also learn to assign values to these various choices based on expected outcome. Thus, behaviors should be affected by rewards, undergoing long-term changes when rewards are different than predicted but remaining unchanged when rewards occur exactly as predicted. Dopamine neurons in the substantia nigra are believed to be involved in reward-dependent behaviors, especially with the reinforcement mechanism involved in learning.[12] **They were found to encode reward prediction errors and were activated by the receipt of an unexpected reward and inhibited by an omission of an expected reward.** Dopamine neurons were also activated by rewards during early trials, when errors were frequent and rewards unpredictable, but activation was progressively reduced as performance was consolidated and rewards became more predictable.[13]

Returning to the nigrostriatal pathway, when activated the pathway has a net excitatory effect on the cortex. As previously described dopamine neurons were activated during the early reward learning stages as well as upon deviations from expected outcome. As such, the changes in pattern of dopamine firing does not simply increase or decrease movement, but rather fine tune the balance of the direct and indirect pathways, allowing for enhanced activation of the cortical motor programs responsible for producing rewarding outcomes and suppression of motor programs that do not result in reward.

The current model of motor movement is that many primitive motor programs are stored in the cortex, and the role of the basal ganglia is to release and inhibit these primitive motor actions in a precise temporal sequence[3] such that competing motor programs are suppressed by the indirect pathway and the appropriate program is disinhibited by the direct pathway.[6] Glutamatergic release by cortical neurons onto MSNs in the striatum with D1 receptors result in a GABAergic output to the GPi, thus lifting its inhibition on the thalamus and eliciting an action. Glutamate release onto MSNs with D2 receptors induces GABA release onto the GPe, thus allowing the STN to excite the GPi and SNr, inhibiting the thalamus and inhibiting action. However, dopamine release by the SNc onto MSNs in the putamen is rather more subtle. It appears that dopamine does not directly induce or inhibit firing in a cell. Rather, its release modulates the excitability of the neuron to glutamate.[14] It has been shown that dopamine D1 receptor signaling enhances dendritic excitability and glutamatergic signaling in striatonigral MSNs, whereas D2 receptor reduces the excitability of postsynaptic neurons to glutamate and release of glutamate by the presynaptic axon terminal. When D2 receptors are activated by dopamine binding, the excitability of the neuron to glutamate is greatly reduced, thus it does not release GABA onto GPe, and disinhibits a competing pathway that was previously inhibited by the indirect pathway. When D1 receptors bind dopamine, the neuron's excitability is enhanced and the direct pathway is more likely to be activated. This piece of information corroborates with the previous observations that dopamine

neurons in the substantia nigra were more active during the learning period of earlier trials, to respond to prediction errors and unexpected rewards. During early trials when an organism is starting to associate a certain motor program with a certain cue in order to receive a reward, release of dopamine onto D1 neurons in the striatum increases their excitability to signals from the cortex and thus promotes the activation and thus consolidation of an appropriate motor program. When an organism makes a prediction error, excitability of D2 neurons is reduced and thus competing motor programs are disinhibited and the mouse can explore different motor programs. The end result of such a modulatory system is that motor habits which are likely to result in reward are retained, whereas those which interfere or reduce the likelihood of reward are inhibited.[3]

By modulating the elementary motor programs, mid-brain dopamine neurons are key components of the brain's reward system. But how do the dopamine neurons know when to release dopamine to modulate the direct and indirect pathways? It had been unclear which brain areas provide dopamine neurons with the signals necessary for their response to sensory stimuli predicting reward until recent efforts identified the lateral habenula (LHb) as a major candidate for a source of negative reward-related signals in dopamine neurons of the SNc and ventral tegmental area (VTA).[15, 16, 17] Habenula neurons were activated by a no-reward-predicting, or a punishment-predicting target and inhibited by a reward-predicting target, especially when they were less predictable, whereas dopamine neurons were excited and inhibited by reward-predicting and no-reward-predicting targets, respectively. These results suggest that **LHb sends inhibitory input to dopaminergic neurons** in determining their reward-related activity, and has the potential to adaptively control both reward-seeking and punishment avoidance behaviors.[2, 18, 19] The positive reward prediction error encoding by dopaminergic neurons— activation and release of dopamine upon an unexpected reward and inhibition of dopaminergic neurons upon an unexpected lack of reward— allows for actions that result in reward to be reinforced while inhibiting actions that no longer result in reward. Increases in habenula activity correlated with the Hamilton Rating Scale for Depression.[20] This may be due to elevated activity of the LHb inhibiting the midbrain dopaminergic systems, resulting in decreased drive to seek rewards, which may contribute to depressed behavior.

What else is missing from the perspective constructed of the motor loop of the basal ganglia? The direct pathway releases the tonic inhibition imposed by the GPi and SNr on the thalamus, and activates a certain motor program and the indirect pathway increases the inhibitory activity of the GPi via the GPe and STN, which inhibits competing motor programs. Both pathways are modulated by dopaminergic neurons in the SNc, which is subject to inhibition by the LHb. How does one close this loop? How does the LHb receive

feedback to ultimately induce dopamine release to reinforce a motor program, or inhibit dopamine release in response to an omission of expected reward?
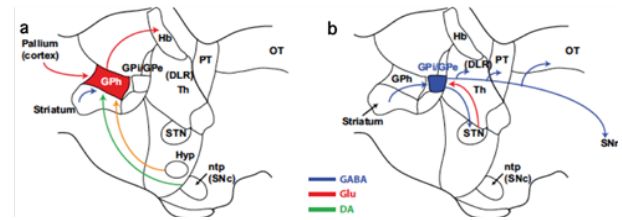

Figure 2: Globus pallidus circuits.21 a. GPh receives GABAergic input from the striatum and projects glutamatergic output to the lateral habenula. b. Primarily GPe and GPi neurons participate in the direct and indirect pathways of the basal ganglia, receiving glutamatergic input from the STN and projecting GABAergic input to the thalamus.

The component of this circuit that provides the signal to the LHb appears to originate in the pallidal region, close to the GPi. The GPi is classically considered to be related to the sensori-motor basal ganglia.[22] It receives inputs from the dorsal striatum, GPe, STN, and projects to the LHb and the ventral lateral thalamic motor nuclei, which, in turn, innervates the premotor and supplemental motor cortex and, thus, completes a basal ganglia-thalamocortical loop. Pallidal neurons projecting to the thalamus and those projecting to the LHb constitute separate neuron populations.[23] Some fibers originating in the GPi arborize extensively in the LHb, and exhibit numerous terminal-like specializations consistent with an important pallidal influence on lateral habenular functions. It was shown in lamprey that a separate evaluation circuit regulates habenula-projecting globus pallidus (GPh) neurons. These neurons are located in close proximity to components of the circuit that participate in the direct or indirect pathway and can thus integrate real time signals from the cortex to convey to the lateral habenula to call upon modulatory signals. They receive inhibitory input from the striatum but have glutamatergic output and can drive the activity of the lateral habenula, which then inhibits midbrain dopamine neurons in the SNc. The release of dopamine can provide feedback that reinforces a particular motor program.

What are the physical implications of such a network? From what we already know about the dopamine circuit and the communication between the GPh and the lateral habenula, we can derive the expected neural behavior of GPh neurons with various stimuli. When the GPh is inhibited, LHb does not extend inhibitory afferents to SNc and dopamine release onto the D1 and D2 receptors in the neurons of the striatum serves to lift the inhibition by the GPi and SNr and generate a net reinforcement of the motor plan. This should occur in tandem with the direct pathway so that when the cortex sends glutamatergic input to D1 neurons in the striatum, disinhibiting the GPi and executing a selected motor plan, the pathway is reinforced by activation of the dopamine pathways. On the other hand, activation of the GPh allows LHb to inhibit the dopaminergic neurons in the SNc, which
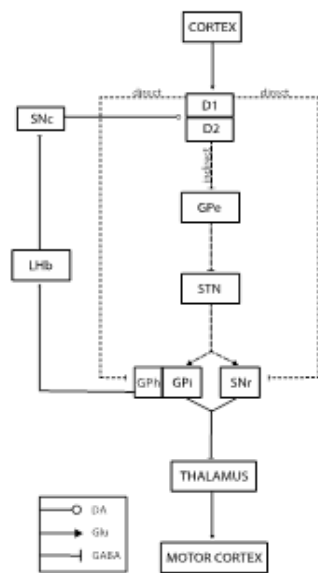
Figure 3: A rough sketch of the components of the basal ganglia: medium spiny neurons expressing either dopamine D1 or D2 receptors make up the majority of the striatum and govern the activation of the direct
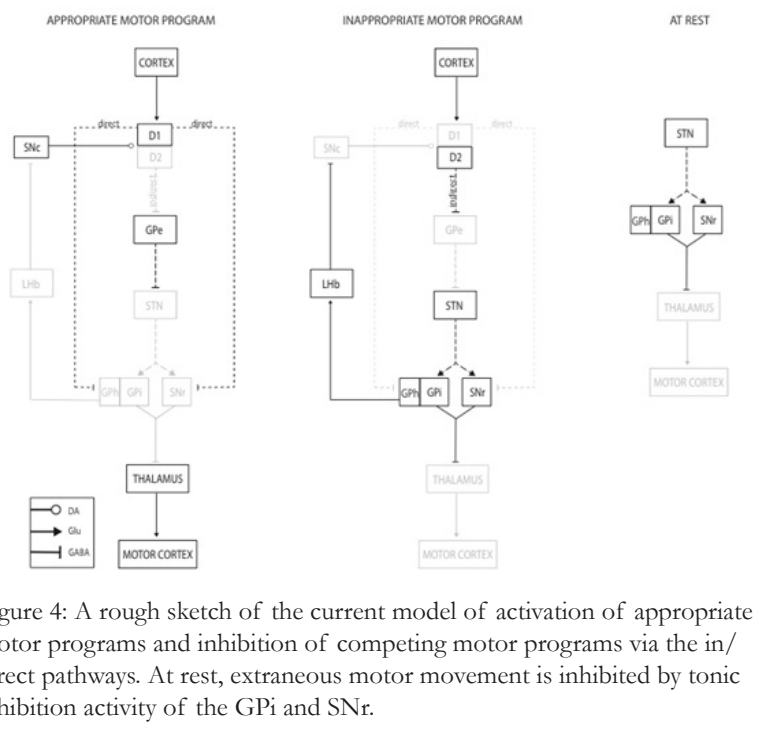

Figure 4: A rough sketch of the current model of activation of appropriate motor programs and inhibition of competing motor programs via the in/direct pathways. At rest, extraneous motor movement is inhibited by tonic inhibition activity of the GPi and SNr.

occurs upon an omission of an expected reward. This should occur alongside the indirect pathway, so that activation of D2 neurons inhibits the GPe, allowing tonic activity of the STN to activate the GPi and SNr, inhibiting thalamic/motor output. When the indirect pathway is activated, inappropriate motor programs are inhibited, or diminished, and signals from the GPh are conducted to inhibit the firing of dopamine neurons. Note that while both D1 and D2 neurons are activated by the cortex, it is the presence of dopamine signals that determine whether they release GABAergic output. We can then isolate the components of the circuit that are active at the activation of the various pathways.

Thus we have more comprehensively outlined the circuit that coordinates movement such that information about the external world is received in the prefrontal cortex, which curates a series of appropriate motor programs, conveyed to the appropriate effector systems via the basal ganglia circuits. Appropriate programs are reinforced and inappropriate ones are inhibited. This information gives us some insight into how an organism learns and how motivation to perform a certain set of tasks is computed: action sets are built largely by trial and error: ones that yield a reward (which can be either physical or psychological) are selected for by reinforcement and ones that fail to yield a reward are diminished by lack of reinforcement. The lack of motivation could be attributed to a deficiency in the reinforcement learning pathway, perhaps an overactive GPi/GPh, such that actions are not executed and dopamine release is inhibited so that even rewarding motor plans are not properly reinforced.

The circuit looks simplistic: there exists other

contributing components to the circuit which have not been incorporated, and whether or not they serve a direct or modulatory function remains to be seen. It is assumed, for instance, that most components of the circuit outlined above have some baseline activity, so that they are active when not inhibited. This may not be correct: there may exist other components that may be actively driving some parts of the circuit (and what drives them?) that are modulated by other parts of the circuit so that there is constant feedback between all components. It was not clear whether the separate population of the GPi that projects to the lateral habenula also receives a distinct set of signals from the STN or the cortex. This is important because it was assumed that the signal that originates in the PFC that sends inappropriate motor programs to the inhibitory pathway is conveyed via the same STN neurons that innervate GPi to innervate GPh neurons, which can then suppress the SNc, in order to coordinate inhibition with dopamine modulation. If this is not true, then the signal to activate GPh neurons must then come from some element in the circuit, perhaps even as far back as the original PFC cortical neurons. How does this other pathway complement the indirect pathway so that dopamine release is timed with inhibition of motion? How do the GPh neurons manage to receive distinct signals from those received by the GPi even though the two populations are so tightly interspersed? More research is needed to answer these questions. The current model of action selection is that signals from the PFC activates a motor program via the direct pathway and inactivates a competing program via the indirect pathway, but how are the direct and indirect pathways distinguished? The medium spiny neurons have D1 and D2 receptors, allowing them to have distinct activation patterns with dopamine release, but how do they get selectively activated to convey glutamatergic

signals form the cortex through either the D1 or D2 neurons? An appropriate motor movement in one context may be a competing motor movement in another; what are the changes that occur for them to be appropriately wired to either a direct or indirect pathway? There are recent reports that both pathways are concurrently activated during movement, and that all MSNs may facilitate or inhibit movement depending on synaptic plasticity.[24, 25] This may not undermine our current model of movement, as a motor program may call upon activating some motions and inhibiting others, but it may add another layer of complexity that is also important to consider.

It is perhaps good to remember some principles of pathways, that the complementation of pathways is important so that competing pathways are not on at the same time, and resources are not spent to activate competing pathways. This implicates the placement of various components in the circuit. Similarly, the brain exists to process external information, which is received via the cortex. Any signal must have originated in and must be ultimately conveyed back to the circuit, if the pathway exists to effect a systemic movement. This raises the larger question as to how information from other pathways is incorporated into the motion circuit: how is sensory information as to the receipt of a reward or punishment conveyed? How does memory about previous decisions factor into the action selection of a motor program? Upon encountering a reward prediction error, in addition to the more immediate update in motor motion, how does the circuit update learning and memory circuits for subsequent decisions? With more information about a circuit comes more questions. We may, in fact, be able to answer the last one, the cross-talk between circuits. We see that the lateral habenula may emerge as an important component in other analogous circuits in the mid-brain, and even a peripheral understanding of these other circuits can guide one's investigation of the motivation and reward circuit by studying the extensive networks of the LHb, and its afferent pathways.

The lateral habenula also features prominently in another important cortico-basal ganglia circuit: the limbic loop. The key structures in this network are the anterior cingulate cortex (ACC), the orbital prefrontal cortex (OFC), the nucleus accumbens in the ventral striatum (VS), the ventral pallidum (VP) and the midbrain dopamine neurons in the ventral tegmental area (VTA). As in the motor circuit, the thalamus (more specifically the medial dorsal thalamus, MD), is the final relay center that conveys signal from the loop to cortical regions.[28] Connectivity between these areas forms a complex neural network that mediates different aspects of reward processing.

Starting with the LHb, inhibitory projections are extended to dopaminergic neurons that reside in the ventral tegmental area (VTA), as depicted in Fig. 5b.[16] The VTA has been extensively studied in reward learning and fear circuits, and shows increased activation in response to stimuli that predict reward. Keep in mind that previously we had also identified the SNc to contain reward-positive dopamine neurons. The VTA, as well, is the seat of dopaminergic neurons in the reward circuit and projects to the ventral striatum, which contains D1 and D2 receptors.[29] The nucleus accumbens (NAcc), located in the VS, features prominently in the reward circuit. The VS receives a large glutamatergic input from the OFC and ACC and projects inhibitory input to the VP and to the VTA (Fig. 5c). The ventral pallidum projects inhibitory output to the medial dorsal thalamus, which has projections to the cortex. This is analogous to structures and functions in the motor loop, where the striatum (putamen) at the head of the direct and indirect pathways, receiving
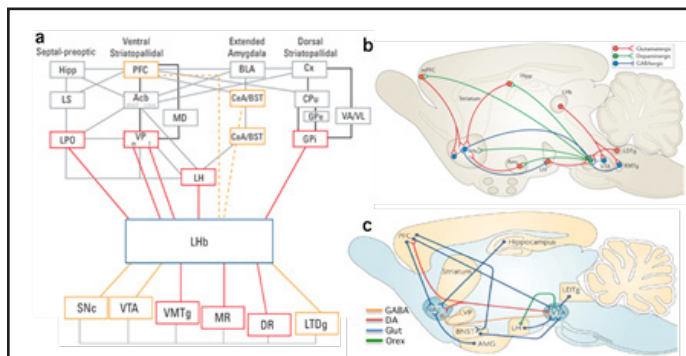

Figure 5: Lateral habenula in reward circuits. a. LHb projects to the SNc and receives input from the GPi, as previously described in the motor loop. It also projects to the VTA, and receives input from the VP, which are components of the limbic loop, suggesting that cross-talk between the two circuits could find a crucial link in the LHb. (The colors represent strength of connection.)[17] b. Activation of the LHb results in GABAergic output to the dopaminergic neurons in the VTA, which sends DA to the NAcc (ventral striatum). The mPFC sends glutamatergic output to the GABAergic neurons in the NAc.[26] c. Note that the NAcc projects GABAergic output to both the VP and the VTA.[27]
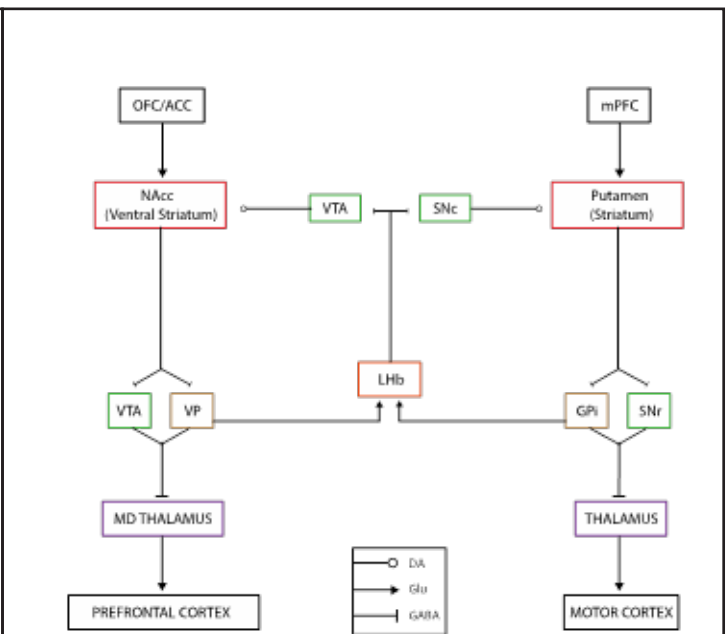

Figure 6: Analogous pathways in basal ganglia circuits. Note analogous functions and physical proximity of structures in same color boxes.

glutamatergic input from the sensorimotor cortices in the motor loop and projecting GABAergic output to the globus pallidus and substantia nigra, which in turn inhibit the ventral thalamus that projects out to the motor cortex.

The proximity of the complementary structures (cortical areas, NAcc/Putamen, VTA/SN, VP/GPi, dorsal/ventral thalamus) probably does not come as a surprise. Reward anticipation induces—and non-reward outcomes suppress—VS activation,[30] leading to proposed theories that VS activity tracks a reward prediction error.[31] The VP also has connections to the SNr and STN, and constitutes a major afferent of projections to LHb (Fig. 5a), in addition to the GPi.[32, 33] With the LHb having connections and receiving projections to major structures in both the motor and limbic loop, it, along with midbrain dopamine neurons in the VTA and SN assume important roles in the feedback between the two circuits.

The idea that VS (limbic loop) can influence the dorsal striatum (motor loop) through the midbrain dopamine cells originated in rodent studies, which demonstrated projections from the NAcc to the dorsal striatum, through the SN. Through this pathway, therefore, limbic regions could impact on the motor regions of the basal ganglia. The dopamine neurons in the VTA and medial SN are associated with limbic regions, and those in the central and ventrolateral SN are associated with the associative and motor striatal regions, respectively. Taken together, the interface between different striatal regions through the midbrain DA cells is organized in a loop interconnecting different functional regions of the striatum and creating a feed forward organization from reward-related regions of the striatum to cognitive and motor areas.[28]

Cortico-striatal Loops

It has not escaped attention that in addition to the motor loop, the basal ganglia supports other cortico-striatal loops that are organized in parallel both functionally and anatomically.[34] This should come as a relief as we gather information about the learning and motivation circuitry, as to develop an appropriate behavioral response to external environmental stimuli, information about motivation and reward needs to be combined with a strategy and an action plan for obtaining goals. The reward circuit comprises several cortical and subcortical regions forming a complex network that mediates different aspects of reward-based learning, leading to adaptive behaviors.

Simultaneous activation of seemingly unconnected regions (eg. mPFC and OFC) indicate that there must be some communication between the regions. This would likely route through the basal ganglia, which has evolved from a historically purely motor or sensory-motor function to a more complex set of functions that mediate the full range of goal-directed behaviors, including emotions, motivation, and cognition. The idea of separate cortical loops in the basal ganglia was

expanded to include several parallel and segregated circuits based on the finding that each general functional area of cortex (limbic, associative, and sensorimotor) is represented in specific regions in each basal ganglia structure.[34]

Reward pathways interface with circuits that mediate cognitive function to affect motor planning. Within each of the cortico-basal ganglia structures, there are convergence zones that can link the reward pathway with those associated with cognitive function. Through these interactive networks, information about reward can be channeled through cognitive circuits to influence motor control circuits. Fig. 7b is especially informative in depicting the cortico-striatal-basal ganglia circuits in parallel, with the major players of the circuits shown in their respective locations in the brain. The physical proximity of the structures that serve analogous functions is all the more elucidating. The length and direction of the path of a signal through the structures traces out a well-worn loop, and may have origins in the development of the brain, as structures differentiate and separate. Looking at how these tracts compare in more primitive organisms could be interesting. Temporal coordination could be key: while reward anticipation activates the NAcc in the ventral striatum, reward outcomes subsequently recruit the caudate and putamen, including the supplementary motor area, and most likely involves dopamine pathways. Thus, the ventral cortico-basal ganglia network, while at the heart of reward processing, does not work in isolation: there are pathways that allow communication between different parts of the reward circuit and between the reward circuit and the associative circuits.

Characterizing neuron types in the GPh by activity profile

Armed with the large amount of background information as presented above, it may be wise to have a starting point. The globus pallidus, as a major output center for the motor basoganglia circuits arises as a region that can potentially elucidate the inherent networks. We focus on the neurons within the globus pallidus that have projections to the lateral habenula. The lateral habenula has been shown to have extensive projections to structures in both the motor loop and the limbic system, which raises the possibility that the projection from the GPi to the LHb might be a key link between the basal ganglia and the limbic system, providing reward-related information and initiating motivation to
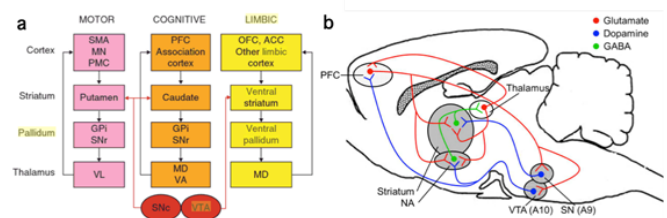

Figure 7: Analogous cortico-striatal circuits. a. Block diagram of circuits.35 b. Schematic of circuits showing the various basal ganglia circuits in parallel.[36]

move.[37] It has been shown in primates that this subpopulation of neurons, refered henceforth as GPh, exhibit negative reward behavior, similar to that of the laternal habenula. The negative reward signal may contribute to the well-known reward coding of neurons, during which constant inhibitory signals from structures in the basal ganglia inhibit various motor neurons. Inhibition of these neurons allow units downward from the circuit to fire and from there initiate various motions.[38] To better understand this circuitry, the wiring as shown in Figure 4 must be associated with real-time firing patterns in the neuron population of interest alongside the biochemical profile. We can collect such data by simultaneously obtaining electrophysiological and biochemical information from the same neurons. In addition, as a major objective of neuroscience is to understand how neural circuits give rise to behavior, it is optimal to obtain functional data on neural firing patterns from animals that are consciously perceiving stimuli and can respond consciously.

MATERIALS AND METHODS
Behavioral Training

We classically conditioned mice with different auditory cues that predicted appetitive or aversive outcomes. The possible outcomes were big reward, small reward (drop of water), no reward, small punishment, or a big punishment (a puff of air delivered to the animal's face). Each behavioral trial began with a conditioned stimulus (CS; a sound, 1 s), followed by a 0.5-s delay and an unconditioned stimulus (US; outcome). Upon the beginning of training sessions the mice are water deprived. In order to train the mice first to lick and to associate a sound with an outcome, drops of water was dispensed unconditionally and immediately following a sound. This also habituates the mice to an unfamiliar surrounding. Various frequencies were also associated with varying amounts of water dispensed. As the mice learn to lick for water, water

is then dispensed conditionally upon detection of licking. A delay period up to 0.5-ms was slowly introduced so that we could study associative behavior.

When the mice were trained for the reward segment, the punishment trials were introduced consisting of white noise preceding an air puff to the face. In order to better facilitate distinguishing of the appetitive and aversive auditory cues, they were delivered on left and right speakers individually. A lick by the animal closes a circuit and is detected by a lickometer,[39] which also detects the rate of licking by the animal during the interval between end of a CS and delivery of a US. Licking rate is one of the parameters by which we gauge how well and how much an animal anticipates an outcome.

Recording Device

As animal subjects, mice provide a very diverse platform for the investigation of behaviors ranging from learning to social performance. Viral expression targeting enables highly precise optogenetic investigation of mouse behavior. The ability to simultaneously record multiple channels of electrical activity during optogenetic manipulation in awake mice has been afforded by the optetrode, which combines electrophysiological recordings of multiple isolated units with optical hardware in awake freely moving mice.[40] At the heart of the device is an optical fiber, which conducted optical stimulation from a laser. It is surrounded by 16 microwires (channels) wound into four tetrode bundles, which records extracellular signals. These four bundles are separated by a width measuring the diameter of the optical fiber and extend beyond the optical fiber in order to better isolate the signals from individual neurons in the region of interest. These channels allowed us to deconstruct the spikes so that they could be represented in various amplitude spaces. Thus, the fiber also provides structural support for the tetrodes during vertical translation through brain tissue. The fiber-tetrode assembly was combined with a custom mechanical drive that allowed adjustment of depth in the brain region. The tetrode microwires were connected to an adaptor that then amplifies and displays the signal on a recording interface. Following
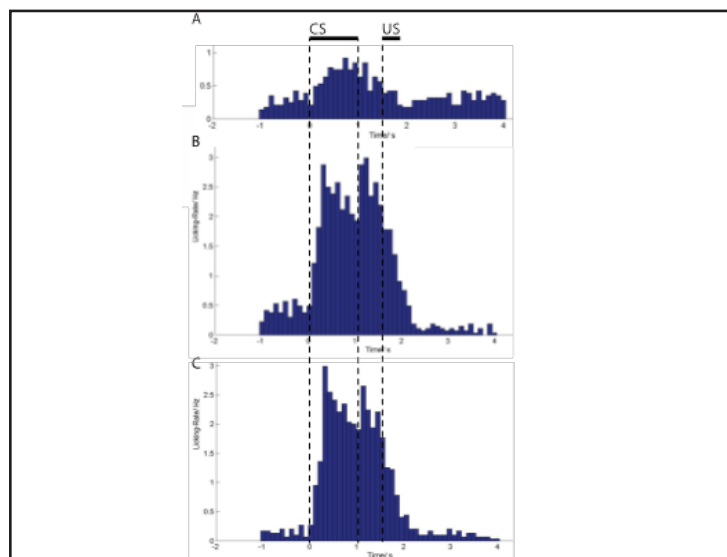

Figure 8: Licking rate during various reward segments. a. No reward. b. Small reward. c. Big reward. Notice licking rate is slightly faster than that in small reward.
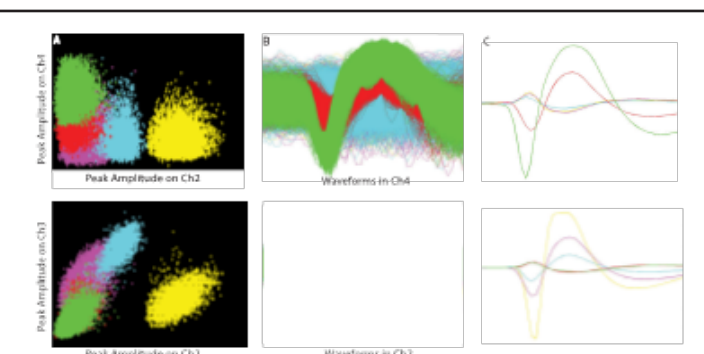

Figure 9: Clustering the spikes. a. Clusters of spikes represented in amplitude space of different channels: good separation between green and red cluster only in one of the channels. b. Actual waveforms of all spikes as seen in various channels. c. Average of the waveforms.

a surgery for implantation, the microdrive was permanently fixed onto the head of the animal. The entire contraption was very light and worn readily on freely moving adult mice, although during the experimental setup the mice were head-fixed, which confers an element of stability to the implanted device.

Spike Sorting

Use of multi-channel electrodes allow for multiple neurons to be recorded simultaneously. Depending on the physical relationship of neurons relative to the multi-channel electrode, the amplitude and extracellular waveform of a neuron on each channel will likely differ from that of a neuron in a different physical location. Clustering is normally accomplished by calculating a set of features of each spike waveform, such as the amplitude on each channel of a tetrode. Spikes presumed to come from the same neuron will form clusters in a high dimensional feature space which can be separated from other clusters representing other simultaneously recorded cells or noise events. Clusters are identified manually or by automatic clustering methods. After clustering spikes into units, it is important to ensure spikes assigned to one cluster are well-separated from other spikes recorded simultaneously, and for that we use two quantitative measures of cluster quality: Lratio and Isolation Distance. The former is an indication of how well separated the clusters are, and the latter is a measure of how well-contained they are. Measurements of isolation of the clusters are especially important in the recording of the globus pallidus because this area exhibits a high baseline activity, and visual clustering of neurons may not be very easy.

The spikes were sorted using a spike-sorting algorithm, MClust.[41] MClust represents each spike as a point in amplitude, energy and wave principle component space as recorded by each channel. Figure 9 is a rough depiction of clusters taken from tetrode data collected in mice globus pallidus. Some clusters are well-defined and isolated but other clusters required looking in other channels in order to find that they can be better separated. We also found that automatic clustering was oftentimes not very useful as clusters were found to fit the cluster quality parameters instead of the parameters used to judge how well-defined are the clusters, which resulted in very strange (and usually unreliable) clusters. After all spikes have been grouped into clusters from various tetrodes, in different recording regions (brain depths), across a span of several weeks across different training sessions, in different animals, we obtained a good number of functional units. In order to identify the types of neurons present in our region, we look at their activity profile and firing rates when correlated with time points of the CS and US.

Results
Classifying the clusters

We recorded the activity of GPh neurons while mice performed the conditioning task described previously. To characterize the responses of the population, we measured the temporal response profile of each unit (neuron) during the various reward trials by quantifying firing rate. Spike recorded from each channel and tetrode were clustered by their waveform properties, and the spikes in each cluster were then categorized by their preceding CS-reward or punishment- and was subsequently plotted against the time points during the various trials, obtaining first a raster plot of the spike time points (represented by a tick) spanning an entire trial. The tick marks were then binned and plotted as a histogram that represented the mean firing rate across a discrete set of time points during a trial. Henceforth activity of a neuron will be represented as firing rate (number of spikes/sec). Once we
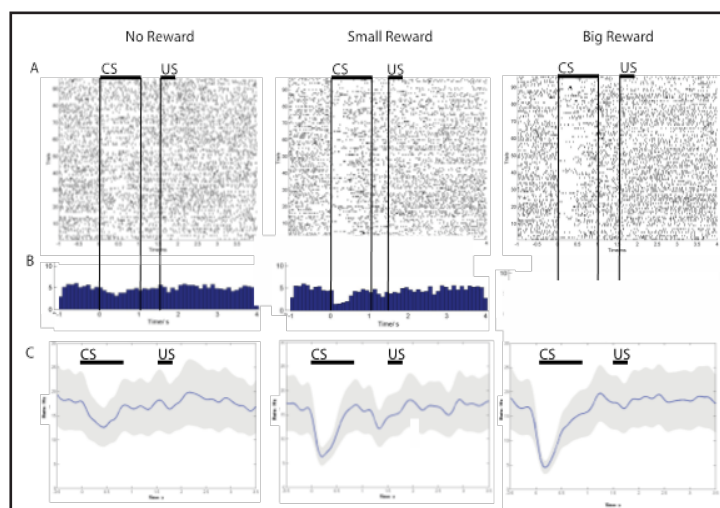

Figure 10: Firing rate during different time points, categorized by no reward, small reward and big reward. The first two panels, A and B, represent a single cluster. a. Raster plot of spikes spanning a trial for multiple trials. b. Histogram of binned spikes. c. Average of firing profile for units in the same tetrode exhibiting similar activity profile. Gray area indicates confidence range.
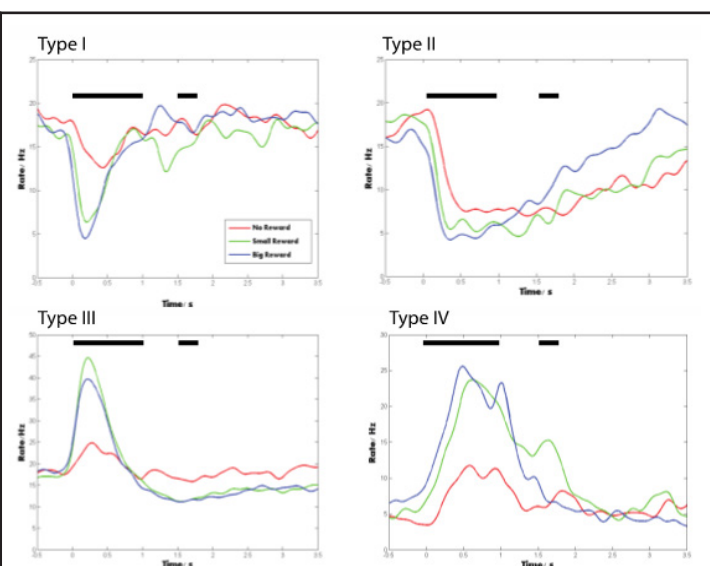

Figure 11: Types of neuron temporal activity profiles present in the recording region. Assigning of type numbers is arbitrary. As before, the long and short black segments indicate CS and US respectively.

obtain the histogram, the activity profile of a unit was average with another unit exhibiting the same activity profile to the same stimulus. One sample trial of a unit is presented. The units that were averaged were taken within the same tetrode, but as more units are collected and analyzed, units from different tetrodes, even different days may be averaged.

After having clustered all the recorded spikes, we can then group the units with the same activity profile. This yields four types of neuronal responses that were distinguished primarily according to the magnitude and length of response to the CS and US. Observing the temporal profiles of responses in trials with rewards, we found neurons that were excited or inhibited phasically by reward or punishment-predicting stimuli. It would be difficult to assign a biochemical character to these neuron types just based on activity profile, especially since the neurons within the GPi that project to the lateral habenula are not very well studied and characterized. With data from the tagging experiments however, we will be able to confirm the biochemical identity of these neurons. Within our recording region, we identified a large number of neurons which showed prolonged inhibition to the CS as well as to the US. Other neurons show a phasic excitation or inhibition to the CS. As the boundaries between the GPi and GPh are not too distinct, it was possible that amongst our recorded neuron population there were neurons from the GPi. The response profile below represents an average of the units which had similar response patterns.

There were other recorded responses which did not facilitate grouping into a distinct category, for which more recordings would be very useful. The majority of the neurons showed a much more pronounced response to the CS (auditory cue) rather than the US (actual reward), although in Type I neurons, there appears to be a slight dip in neuron firing rates preceding the US. Type I neurons are hypothesized to be the ones we are looking for: glutamatergic neurons in GPi that project to the lateral habenula, which exhibit reward negative firing patterns, which is similar to what has been identified.37 These would fit in the picture of the function of these neurons, as depicted in Figure 4: the inactivity of these neurons promotes firing of dopaminergic neurons in the substantia nigra, reinforcing an appropriate motor program. It should also make sense that these neurons also exhibit a certain degree of change in response to the US, as it is the receipt of the US that determines whether a motor program should be reinforced or altered. This consideration may be more important when computing reward prediction errors (RPE).

Reward Prediction Error

An important response property that supports RPE coding in dopaminergic neurons is their decrease in firing rate when an expected reward is omitted.[42] We omitted reward unexpectedly on about 10% of big reward trials. We also added some trials where a big reward was dispensed when the auditory cue predicted no reward. During the analysis
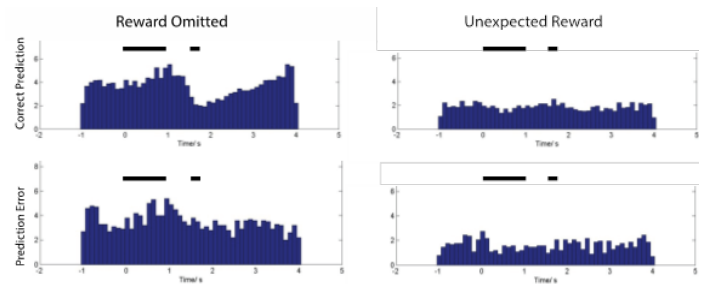

Figure 12: Reward Prediction Error. The panel on the left depicts a reward omission and that on the right depicts an unexpected reward. The reward prediction error response is on the lower panels.

of the clusters, the trials which were marked to be reward prediction error trials were extracted and their response plotted against the time points for CS and US as before and binned and plotted in as a histogram. Very few trials were run and only Type I neurons were analyzed, thus this is just a very preliminary presentation of results of RPE in target GPh neurons.

There is a large difference between the response profiles of correct prediction versus prediction error in the case of reward omission. In this case, the pronounced difference rests in the response to the US, which should be the case as both conditions receive the same auditory cue. It appears that there is an inhibition of the GPh neurons as reward is received, and an absence of inhibition when reward is omitted. When taken together with the dopaminergic neuron data, one can see how this might be a viable way for the globus pallidus to encode RPE with DA neurons in the substantia nigra. As previously established, DA neurons decrease in firing rate when an expected reward is omitted. In Figure 4, inactivity of DA neurons in SNc is correlated with activity in the GPh neurons. As hypothesized, glutamatergic neurons in the GPh activates GABAergic neurons in the lateral habenula, which inhibit the dopaminergic neurons in the SNc. Thus, a decrease in DA neuron activity should correlate with increase activity of the GPh. This is exactly what we see in Figure 12: when reward is omitted, there is a higher firing rate in the Type I GPh neurons as compared with the normal response pattern of a correct prediction of reward. In addition, a correct prediction of a large reward during the US block results in an inhibition of firing of GPh neurons. This phenomenon also agrees with our previous discussion of DA neurons: inactivity of GPh neurons removes the excitatory input to the lateral habenula, which is dis-inhibitory on the DA neurons in the SNc. Thus, DA neurons can release dopamine into the direct and indirect pathway. As previously discussed, this is a reinforcement mechanism which promotes an appropriate motor program resulting in an expected reward, and thus it should not be surprising to see an inhibition of GPh neurons upon a correct prediction.

What about in the case of the unexpected reward? This was only one instance (one cluster) and thus does not afford much in the area of significant differences, but we

can see that the firing rate during the prediction error trials was generally lower and more sparse than that of the correct prediction trials. In this case, the mice hear an auditory cue for nothing, but receive a large reward instead during the instances where we introduced a prediction error. A lower firing rate during the US onset in the prediction error compared with those during a correct prediction would serve the same purpose as that previously discussed: inhibition of GPh is strongly correlated with activation of DA neurons in the SNc, which serves a reinforcing purpose. This is the brain telling the circuits that there might need to be adjustment in the motor program upon hearing the sound for no reward as now there is reward. Thus, the mice might anticipate the no reward signal just slightly more. To summarize, inhibition of GPh neurons indicate a reinforcement of an anticipation (unexpected reward) by allowing the release of DA, whereas activation of GPh indicates a discouragement of an anticipation (reward omission) by inhibiting release of DA, which restates the finding mentioned previously that there is a decrease in DA neurons firing when reward is omitted. This is supported by our preliminary data on reward prediction error. Much more data collection would be crucial to arrive at an observation with greater confidence. Looking at licking rate as a measure of anticipation may also be elucidating.

Reward History

Reward history is essentially a measure of how the value assigned to a particular auditory cue-and thus reward-changed depending on the reward that was received in the preceding trial. It was an attempt to understand (1) if there were any changes (2) and if there were, in which neurons and (3) gain a rough understanding of how reward history was encoded in the neurons in the region of interest. The same experimental procedure as that outlined in classifying the neuron types was carried out: mice were given auditory cues which always preceded their associated reward and its delivery was conditional upon licking. The only difference was in the data analysis. For each cluster, not only were they separated by the auditory cue for that trial, all the spike time points for each auditory cue was also separated by their preceding outcome: no reward, small reward, big reward. We essentially had nine groups of spikes: 3 different rewards, and each reward following one of the three reward outcomes. Thus it was basically a trick with extracting the relevant information out of the huge amount of data that was collected. We analyzed Type I and Type III neuron firing profiles. Each type is represented by only one cluster as it seemed to be messy to simply average the reward history response profiles for multiple units. Preliminary results are depicted at the bottom of the page

First of all we note that there hardly seems to be any differences in reward history response profiles in Type III neurons. This may come as a slight surprise as a response that reacts so strongly to a CS cue does not have any reward history discrimination, even as it discriminates between the cues for the reward outcomes. The plots for the Type I neuron is slightly more erratic and definitely suffers from lack of a large pool of data to average out the noise. We look first at the no reward cues, and see that there is a significantly larger dip in the CS block for the no reward cues that follow a large reward (blue trace), with a higher baseline activity. Remembering that an inhibition of GPh neurons indicate a greater anticipation, it appears the neurons are telling us that even when the mice know that he is not getting a reward for that trial, because he had a large reward recently, he has a greater anticipation for the current trial. We next look at the large reward cues. Although there is almost no difference between the inhibition of firing during the CS block, we see that there is a lower
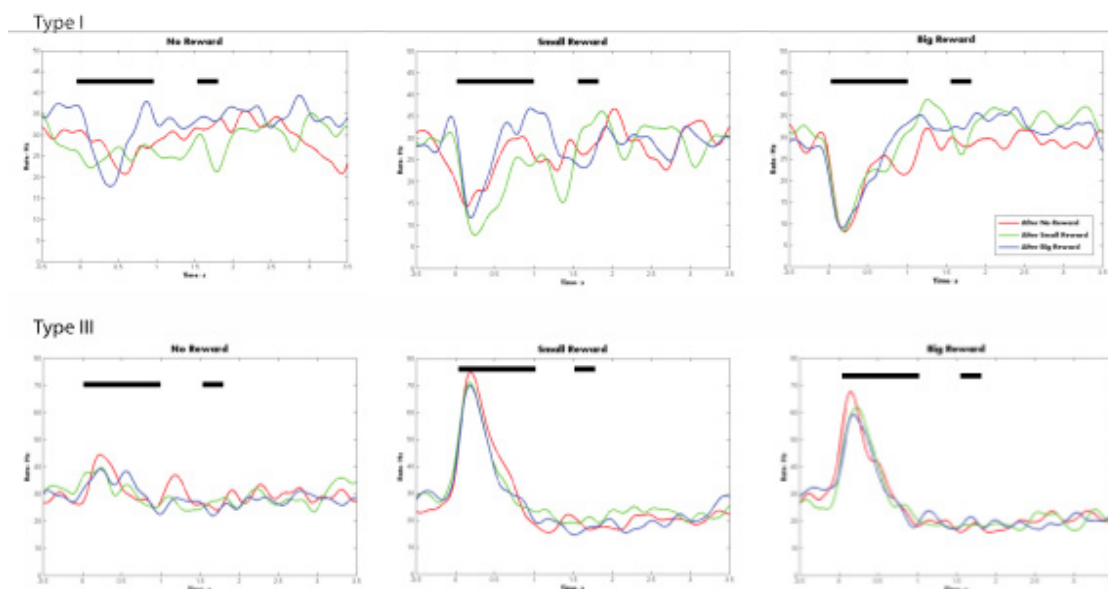

Figure 13: Reward History. As in Figure 11, Type I and Type III neurons response profiles were categorized into reward outcomes. Each reward outcome is further categorized by the outcome they follow.

baseline activity in the big rewards that follow a no reward outcome (red trace). Generally, inhibition of the GPh neurons promotes motor movement, indicating that in this case the mice are slightly more motivated to obtain the reward after not getting anything previously. The higher baseline activity in the neurons in the no rewards cue following a large reward outcome can also be interpreted as the mouse is slightly less motivated after having received a large reward. The small rewards plot does not bear a distinct trend from which we can draw significant conclusions. Reward history does appear to play a role in value-based decision making. It appears that reward history may be encoded in the baseline activity, and to some extent the level of inhibition achieved during a CS block, although that seems more to govern anticipation than motivation, which plays a larger role in reward history. The experimental data matches well with the current model for basal ganglia circuits, but as always, more collection and analysis needs to be done before any conclusions can be drawn.

DISCUSSION

There exists a subpopulation of the GPi neurons that projects to the lateral habenula, termed the GPh neurons. There has been significant evidence pointing to the significance of the lateral habenula in reward learning tasks. The GPi is a major output center in the basal ganglia, commonly associated with learning, reward circuits, motivation, decision-making, but most importantly with motor execution. A decision must be actualized by an action, and thus it behooves us to better understand these GPh neurons as they must serve as a critical link between the motor and limbic circuits. In this series of experiments and data analysis we have attempted first to record from and classify the types of neurons that exist in this region in mice, who have been trained to carry out a set of value-based tasks in which they lick for water. Delivery of water for these water-deprived mice (a reward) is contingent upon the act of licking, and is preceded by an auditory cue (CS) where different frequencies are associated with different outcomes (US: no water, small or large amounts of water). The recording device is similar to the optetrode design, in which an optical fiber is surrounded by 16 microwires wound into 4 tetrode bundles. Waveform properties of each spike event picked up by the tetrodes is represented in a feature space that ultimately allows those with similar properties to cluster together, presumably comprising of all the spikes belonging to a single unit across the entire number of trials. These spikes in each cluster are then separated by the reward outcome in order to observe the responses of different types of neurons in various situations.

We have arrived at least four different types of neurons in our recording region, with phasic and prolonged activation or inhibition to the CS. Type I, with phasic reward negative responses, appears to be the glutamatergic GPh neurons, with an inhibition in response to a cue that predicts reward. This agrees with the circuitry as depicted in Figure 4: inhibition

of the GPi promotes dopamine release and reinforces a lucrative motor program. These neurons also exhibit reward prediction error encoding, although the mechanism is not well-understood. The functional aspect in encoding for RPE appears to be the difference in firing rate during the onset of the US. This would make sense as one would imagine a difference in expectation, or an error in prediction is received when the outcome is different from expected and must be encoded in some form and fed back to the circuit that updates perception and value-based decision making for subsequent trials. We see this phenomenon further in reward history, where the main parameter for motivation after consideration of reward history is in the baseline firing rate of the response profile. Are the elements of the circuit that raises or lowers the baseline and thus motivation following an outcome inherent components of the basal ganglia or are they remote elements in a different, but related circuit?

Other parallel circuits to the motor loop involving the basal ganglia are investigated in the introduction, and increasingly it appears that nature works by analogies and patterns, from homologous limbs to conserved protein folds, to brain circuits. Already we see that analogous elements occupy the same space (Figure 6), and that activation and modulatory mechanisms are similar. What remains is to extrapolate any ideas from a better studied circuit and apply it to parallel circuits, which may facilitate understanding of both circuits. For example, social interactions, which are well studied in the limbic loop, have been known to great counteract the effects of addiction and depression.[43] What are the modulatory elements and how does it affect the structures in the basal ganglia? With input and output terminals, effectors that can activate, inhibit or modulate a downstream element, gating mechanisms, and loops that can run in parallel or in series in strategic topography, the neurons of the brain may be able to form a formidable logic circuit that can compute seemingly abstract concepts such as probability, effort, time, payout, etc., which are important considerations for an animal to assign a value to an outcome, and thus make a value-based decision as to whether or not to undertake an action. Thus we can begin to perhaps understand the brain circuitry in the language of the brain: dopamine, instead of being associated with feelings of satisfaction, more directly reinforces an expectation or a motor program, so that we are more likely to repeat the same actions, depression is not sadness, but perhaps more of a lack of motivation due to an overactive LHb or GPh, thus suppressing thalamic motor control, etc. This has strong implications in understanding the process of learning and treating mental disorders. With a basic understanding of electrical circuits an enterprising individual could potentially construct a simplistic functional reward circuit and simulate various reward learning situations, and a blueprint could be especially elucidating of the human brain circuitry. Following this line of thought, regions of the brain must be considered as elements in a complex circuit, rather than individual regions, much less a certain cell type. Thus it would not be very effective

to speak of a certain cell type in the brain, removed from their individual circuits as neurotransmitters are but effectors in a functional loop. One unit does not function alone; it belongs to a network and thus one should be wary of attributing an observed effect completely to a stimulated cell type or brain region, as effects upstream or downstream from that circuit element may not have been duly considered.

There are some experiments yet to be run and analyses yet to be carried out that would make this story more complete. One of the interesting pieces of information could be obtained as early as the first day of training. While for certain pieces information such as reward prediction error, it is imperative that the animal understands the cues well enough to know that there was supposed to be a reward during the trials when the reward is omitted, during the early stages it would be interesting to see how firing pattern of the neurons in the target region slowly evolves as the mice learn the task. As the reward circuit plays a big role in the process of learning, one can set the detectors to understand how the neurons form connections and ultimately adopt the final response profile. How does it correlate with the learning observed from the anticipatory licking? On the same lines, satiation would also be interesting to look at. We have already seen in the reward history that mice are less motivated after they have had a big reward, but what about when they are satiated up to the point where they would also ignore the big reward? How long before they stop licking does the brain register that it is satiated? Is it simply a change in baseline, thus a lack of motivation, or does the auditory cue not even induce inhibition of firing anymore? To obtain these information regarding satiation and learning, plotting the response across trials, rather than lining them up by trials and plotting over the CS/US time points might be interesting.

Another piece of the puzzle that appears to be blatantly missing is the tagging profile. The experiments were conducted, but the analyses had not been very comprehensive nor conclusive. Using genetic engineering (which is another area where ingenious tools are available to promote better experiment design), we can insert light-gated ion channels such as ChR2 in a specific neuron population (e.g. VGlut2 neurons in the GPi) using a Cre recombinase system. For each neuron, we can measure the response to light pulses and the wave shape of spontaneous and induced spikes. There should be a high correlation between the light pulse and the timing of the spike, and based on the waveform properties we can against represent them in a feature space, and the physical location of the cluster in the same two channel's amplitude space should overlap with a cluster found previously during the electrophysiology recording session. The criterion that the light-evoked waveform must look almost identical to the spontaneous waveform also ensures that the previously identified unit is correctly assigned a biochemical identity. Using retroviral tracing techniques one can also get a better idea as to the immediate connections to and from a target population of neurons.

Although there were many new techniques that I was exposed to through this project, including murine handling and surgery and building the microdrive, most of my time was spent on data analysis, from clustering to generating raster plots, from data collection to sorting, of a massive amount of information. It would be much more efficient to streamline the entire process of data collection to data analysis so that computations can run in the background while prepping and designing experiments can take place, and freshly collected data is automatically fed through the analysis machinery. Perhaps cleverer experimental design would need to be instituted in order to observe an isolated phenomenon. The brain undergoes and presents a lot of activity and careful extraction and sorting of the data could reveal so much about the brain machinery, and some thought should go into the code that sorts this information. Other analysis methods such as regression coefficients or changes in baseline could more clearly elucidate a trend. The clustering is achieved with MClust which works well most of the time. However, the automatic clustering program KlustaKwik often produces unreliable clusters and most clustering had to be done by hand, which greatly slows down the process. A more reliable automatic clustering software would greatly facilitate the data analysis process.

At last, we note that neural circuits are incredibly complex-and for good reason-but nature uses many analogies and the development of new tools is very promising. Keeping the analogous circuits in mind and using the available tools at hand, we have gotten glimpses of the role that the GPh plays in the basal ganglia circuits: the reward negative neurons which are most likely glutamatergic, encoding for anticipation upon receiving the cue and assigning a value to the associated outcome. Whether it is this value that changes, or the overall baseline motivation that fluctuates when computing reward history remains to be seen as more units are recorded and analyzed, which will also elucidate the changes in temporal firing rate when the neurons are encoding reward prediction error. For now the reward evaluation pathway consisting of the glutamatergic GPh neurons, GABAergic LHb neurons and DAergic SNc neurons stands up to experimental data, however how they interact to galvanize a decision and what roles they play in other circuits still remains to be studied.

it's taken me a long time to comprehend what he meant but slowly I am beginning to arrive at my own understanding. My own personal interest in making tools is I find more and more supplemented by my understanding of the population that needs it and the purpose it will serve.

References

[1] J-C. Dreher, L. Tremblay, Handbook of Reward and Decision Making, Academic Press, Jun 4, 2009. Print.

[2] M. DeLong, T. Wichmann, Update on models of basal ganglia function and dysfunction, Parkinsonism Relat Disord. 2009 Dec;15 Suppl 3:S237-40.

[3] A. S. Davis, Handbook of Pediatric Neuropsychology, Springer Publishing Company, Oct 25, 2010.

[4] A. Stocco, C. Lebiere, J. R. Anderson, Conditional Routing of Information to the Cortex: A Model of the Basal Ganglia's Role in Cognitive Coordination, Psychological Review 2010;117 (2): 541–74.

[5] M. Patestas and L. P. Gartner, A Textbook of Neuroanatomy, New York:Wiley-Blackwell; 2006. Print.

[6] J. Knierim, Neuroscience Online, http://neuroscience.uth.tmc.edu/.

[7] J. D. Fix, Basal Ganglia and the Striatal Motor System, Neuroanatomy (Board Review Series) (4th ed.). Baltimore: Wulters Kluwer & Lippincott Wiliams & Wilkins. pp. 274–281; 2008.

[8] I.G. Cameron, M. Watanabe, G. Pari, D. P. Munoz, Executive impairment in Parkinson's disease: response automaticity and task switching. Neuropsychologia. 2010 June;48 (7): 1948–57.

[9] D. Purves, G. J. Augustine, D. Fitzpatrick et al., Neuroscience 2nd Ed., Sunderland (MA):Sinauer Associates; 2001. Print.

[10] University of Fribourg Biochemistry Department, Striatum and Accumbens, August 31, 2014.

[11] J.-M. Beaulieu and R. R. Gainetdinov, The Physiology, Signaling, and Pharmacology of Dopamine Receptors, Pharm. Reviews. 2011 Mar; 63(1):182-217.

[12] E. S. Bromberg-Martin, M. Matsumoto, S. Hong and O. Hikosaka, A Pallidus-Habenula-Dopamine Pathway Signals Inferred Stimulus Values, J Neurophysiol 104:1068-1076, 2010.

[13] J. R. Hollerman, W. Schultz, Dopamine neurons report an error in the temporal prediction of reward during learning, Nat Neurosci. 1998 Aug;1(4):304-9.

[14] D. J. Surmeier, J. Ding, M. Day, Z. Wang, W. Shen, D1 and D2 dopamine-receptor modulation of striatal glutamatergic signaling in striatal medium spiny neurons, Cell. May 2007; 30(5):228–235.

[15] M. Matsumoto and O. Hikosaka, Lateral Habenula as a source of negative reward signals in dopamine neurons, Nature 447, 1111-1115 (28 June 2007).

[16] G. R. Christoph, R. J. Leonzio, K. S. Wilcox, Stimulation of the lateral habenula inhibits dopamine-containing neurons in the substantia nigra and ventral tegmental area of the rat., J Neurosci. 1986 Mar;6(3):613-9.

[17] S. Geisler, The Lateral Habenula: No Longer Neglected, CNS Spectr. 2008;13(6):484-489.

[18] M. Matsumoto and O. Hikosaka, Representation of negative motivational value in the primate lateral habenula, Nat Neurosci. 2008 Nov 30;12(1):77-84.

[19] S. Hong, T. C. Jhou, M. Smith, K. S. Saleem and O. Hikosaka, Negative Reward Signals from the Lateral Habenula to Dopamine Neurons Are Mediated by Rostromedial Tegmental Nucleus in Primates, J. Neurosci. 10 August 2011, 31(32): 11457-11471.

[20] J.S. Morris, K.A. Smith, P.J. Cowen, K.J. Friston, R.J. Dolan, Covariation of activity in habenula and dorsal raphe nuclei following tryptophan depletion, Neuroimage. 1999;10:163-172.

[21] M. Stephenson-Jones, A. A. Kardamakis, B. Robertson, and S. Grillner, Independent circuits in the basal ganglia for the evaluation and selection of actions, Proc Natl Acad Sci U S A. 2013 Sep 17;110(38)

[22] G.E. Alexander, M.D.Crutcher, M.R. DeLong, Basal ganglia-thalamocortical circuits: parallel substrates for motor, oculomotor, "prefrontal" and "limbic" functions, Prog Brain Res. 1990;85:119-146.

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

[24] P. Calabresi, B. Picconi, A. Tozzi, V. Ghiglieri and M. Di Filippo, Direct and indirect pathways of basal ganglia: a critical reappraisal, Nature Neuroscience 17, 1022–1030 (2014).

[25] G. H. Cui, S. B. Jun, X. Jin, M. D. Pham, S. S. Vogel, D. M. Lovinger and R. M. Costa, Concurrent activation of striatal direct and indirect pathways during action initiation, Nature 494, 238–242 (14 February 2013)

[26] S. J. Russo and E. J. Nestler, The brain reward circuitry in mood disorders, Nature Reviews Neuroscience 14, 609–625 (2013).

[27] J. A. Kauer and R. C. Malenka, Synaptic plasticity and Addiction, Nature Reviews Neuroscience 8, 844-858 (November 2007).

[28] S. N Haber and B. Knutson, The Reward Circuit: Linking Primate Anatomy and Human Imaging, Neuropharmacology. Jan 2010; 35(1): 4-26.

[29] F. W. Hopf, M. G. Cascini, A. S. Gordon, I. Diamond and A. Bonci, Cooperative Activation of Dopamine D1 and D2 Receptors Increases Spike Firing of Nucleus Accumbens Neurons via G-Protein-βγ Subunits, J. Neurosci. 23(12):5079 –5087, 2003.

[30] A. Beck, F. Schlagenhauf, T. Wüstenberg, J. Hein, T. Kienast, T. Kahnt, K. Schmack, C. Hägele, B. Knutson, A. Heinz, and J. Wrase, Ventral Striatal Activation During Reward Anticipation Correlates with Impulsivity in Alcoholics, Biol Psychiatry. 66(8):734-42, 2009.

[31] G. Juckel, F. Schlagenhauf, M. Koslowski, T. Wustenberg, A. Villringer, B. Knutson, J. Wrase and A. Heinz, Dysfunction of ventral striatal reward prediction in schizophrenia. NeuroImage 29 (2006) 409–416.

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

[33] K. S. Smith, A. J. Tindell, J. W. Aldridge, K. C. Berridge, Ventral Pallidum Roles in Reward and Motivation, Behav Brain Res. Jan 23, 2009; 196(2): 155–167.

[34] G. E. Alexander, M. R. DeLong, P. L. Strick, Parallel organization of functionally segregated circuits linking basal ganglia and cortex, Annu Rev Neurosci 9, 357-81, 1986.

[35] M. E. Shenton, BI. Turetsky, Understanding Neuropsychiatric Disorders, Cambridge University Press, Dec 9, 2010.

[36] E. M. Miller, T. C. Thomas, G. A. Gerhardt and P. E. A. Glaser, Dopamine and Glutamate Interactions in ADHD: Implications for the Future Neuropharmacology of ADHD, June 27, 2013.

[37] S. Hong and O. Hikosaka, The Globus Pallidus Send Reward-Related Signals to the Lateral Habenula, Neuron. 2008 Nov 26;60(4):720-9.

[38] J. P. Bolam, J. J. Hanley, P. A. C. Booth and M. D. Bevan, Synaptic organization of the basal ganglia, J Anat. 2000 May; 196(Pt 4): 527–542.

[39] B. Slotnick, A simple 2-transistor touch or lick detector circuit, J Exp Anal Behav. Mar 2009; 91(2): 253–255.

[40] P. Anikeeva, A. S. Andalman, I. Witten, M. Warden, I. Goshen, L. Grosenick, L. A. Gunaydin, L. M. Frank, K. Deisseroth, Optetrode: a multichannel readout for optogenetic control in freely moving mice, Nat Neurosci. 2011 Dec 4;15(1):163-70.

[41] N. Schmitzer-Torbert, J. Jackson, D. Henze, K. Harris, A. D. Redish. Quantitative measures of cluster quality for use in extracellular recordings, Neuroscience. 2005;131(1):1-11.

[42] J. Cohen, S. Haesler, L. Vong, B. Lowell, N. Uchida, Neuron-type-specific signals for reward and punishment in the ventral tegmental area, Nature. 2012 Jan 18;482(7383):85-8.

[43] S. A. Allsop, C. M. Vander Weele, R. Wichmann and K. M. Tye, Optogenetic insights on the relationship between anxiety-related behaviors and social deficits, Front. Behav. Neurosci., 16 July 2014.