This paper is now published. Please cite as:

Garr, E. (2019). Contributions of the basal ganglia to action sequence learning and performance. *Neuroscience & Biobehavioral Reviews*, 107, 279–295.
<https://doi.org/10.1016/j.neubiorev.2019.09.017>

Contributions of the basal ganglia to action sequence learning and
performance

Eric Garra^a

^aGraduate Center, City University of New York

365 5th Avenue, New York, NY 10016

egarr@gradcenter.cuny.edu

Abstract

Animals engage in intricately woven and choreographed action sequences that are constructed from trial-and-error learning. The mechanisms by which the brain links together individual actions which are later recalled as fluid chains of behavior are not fully understood, but there is broad consensus that the basal ganglia play a crucial role in this process. This paper presents a comprehensive review of the role of the basal ganglia in action sequencing, with a focus on whether the computational framework of reinforcement learning can capture key behavioral features of sequencing and the neural mechanisms that underlie them. While a simple neurocomputational model of reinforcement learning can capture key features of action sequence learning, this model is not sufficient to capture goal-directed control of sequences or their hierarchical representation. The hierarchical structure of action sequences, in particular, poses a challenge for building better models of action sequencing, and it is in this regard that further investigations into basal ganglia information processing may be informative.

Key words: action sequence; basal ganglia; dopamine; motor cortex; reinforcement learning; striatum; initiation; execution; termination

1. Introduction

The ability to learn and execute a complex series of movements in a particular order is a remarkable feat of any nervous system. Organisms are able to learn arbitrary sequences of actions through repeated practice, and once asymptotic performance has been reached, sequences can be recalled with impressive speed and precision. While much of the neuroscience of instrumental learning focuses on a single action that is reinforced or punished by an outcome, that single action is almost always part of a sequence of actions that are either not recorded or ignored by the experimenter. It is rarely the case that the sequence of events that constitute instrumental learning is action → outcome, as simple as it may be to depict it in that way. Often times the action that is to be reinforced or punished (i.e. the ‘target action’) is separated in time from an outcome by other actions, and the target action is always preceded by other actions. Moreover, the target action itself is usually composed of a series of muscle movements that follow a precise order. While there is much that is not known about the neural mechanisms of action sequence learning and performance in animals, there is nevertheless a rich literature that, when considered in full, narrows down the basal ganglia as a primary neural circuit involved in action sequencing. The aim of this review is to succinctly summarize that vast literature and filter out the main themes and unanswered questions across many studies.

In reviewing the literature on action sequencing, one finds a broad and loose definition of the term ‘action sequence’. Sometimes studies of action sequencing focus on different action types that are executed across different spatial locations (e.g. pressing a lever followed by pulling a chain), or actions of the same type that are executed across different spatial locations (e.g. pressing one lever followed by pressing another lever), or actions of the same type that are executed in the same spatial location (e.g. pressing the same lever repeatedly). Furthermore, the

number of actions available to the subject at any given point in time varies from study to study, and the ‘degrees of freedom’ of action will determine the difficulty of the task and thus the rate of learning. For example, an animal that is required to execute a set of actions in a specific order to obtain a reward will face a more challenging learning scenario if all of those actions are simultaneously available compared to a situation in which only one action is available at a time, or when each action is cued by an external stimulus. Action sequencing is a broadly defined phenomenon, and because it is not monolithic, any conclusions that are drawn from a single study may not generalize to other situations. It is important to bear this caveat in mind.

To start, I begin by reviewing one of the most dominant accounts of instrumental learning in behavioral neuroscience today: reinforcement learning (RL). RL dictates that actions are learned ‘in reverse’ from the moment of reinforcement, and this can help to explain some empirical phenomena in the animal learning literature. The mechanisms that allow for this type of learning are fairly simple, and fall under the heading of ‘model-free’ RL. Model-free RL does not capture some of the more complex aspects of action sequencing, and as an alternative another brand of RL (‘model-based’ RL) is sometimes appealed to. But even with model-based RL there are shortcomings, since it does not capture the hierarchical structure of action sequences. I then review how theoretical models of the basal ganglia aim to map RL variables to neural function, while also pointing out crucial shortcomings. Finally, the various components of sequence performance—initiation, execution, and termination—are reviewed with reference to their associated basal ganglia correlates.

2. Behavioral predictions and findings

2.1 Reinforcement learning as a guiding framework

One of the most influential accounts of action sequence learning comes from the computational theory of reinforcement learning (RL). RL is a collection of algorithms that describes how animals and artificial agents can maximize long-term reward through trial-and-error learning (Sutton & Barto, 2018). As an animal explores its environment and encounters appetitive and aversive outcomes, it needs a way of figuring out what specific actions or action sequences led to those outcomes so that it can repeat or avoid those actions in the future. The main idea behind RL is that agents assign values to actions, and then use the changes in these values to adjust the values of preceding actions. Action values are learned either by encountering rewards in the environment or stimuli that *predict* future rewards, with the ultimate goal being an incremental improvement of the agent's decision-making policy. While RL has its origins in artificial intelligence, it has also been championed by psychologists and neuroscientists as a way to understand the computations underlying animal learning and decision-making (e.g. Redish, Jensen, Johnson, & Kurth-Nelson, 2007; Mattar & Daw, 2018; Stachenfeld, Botvinick, & Gershman, 2017). What follows is a very basic overview of RL, intended only to introduce the reader to core concepts.

The main elements of an RL algorithm are states, actions, and rewards. Each action is executed in a state, and whatever value accrues to an action is conditional on that state. The term 'state' is loosely defined. It can be an observable set of stimuli, or it can be unobservable and subject to inference, such as a set of contingencies (Gershman, Blei, & Niv, 2010; Wilson et al., 2014). As an agent navigates through the state space of its environment, it will encounter rewards and learn the values of state-action pairs, which are conceptually similar to stimulus-response associations (Thorndike, 1898). When the agent arrives in a state, the probability of

selecting any given action is proportional to its learned value, and the probability distribution over actions defines the agent's policy.

One of the cornerstone assumptions of RL is that the values of actions are adjusted according to errors in prediction. When an agent executes an action in some state, the value of the outcome at a subsequent time step is compared to the agent's most recent value estimate of the state-action pair. This comparison is known as a prediction error. If the value of the outcome surpasses the learned state-action value, the prediction error is positive and the state-action value is incremented. If the outcome value is inferior to the state-action value, the prediction error is negative and the value is decremented. If the value of the outcome perfectly matches the state-action value, the prediction error is zero and there is no new learning. The prediction error reflects the degree of surprise associated with some outcome relative to the expectation represented by the state-action value, and larger discrepancies between expectation and outcome will create larger changes in the state-action value.

It is easy to understand how the value of a single action is learned when the outcome of the action is immediate. If the action results in reward, the state-action value is updated according to the size of the prediction error. But many pursuits in life are made up of *many* actions that shape the final outcome. These intermediate actions are delayed from the final outcome, and this delay presents a challenge for properly adjusting the values of actions within a sequence. This is called the temporal credit assignment problem (Minsky, 1961), and it is the *raison d'être* for all RL algorithms and a defining feature of action sequence learning.

How, then, can an agent solve the temporal credit assignment problem? An action may lead to a primary reinforcer that has intrinsic value, such as food or water. However, in a sequence of actions in which the occurrence of a primary reinforcer may be temporally distant

from actions further back in the sequence, there needs to be a way of updating action values in the absence of an immediate primary reinforcer. It is for this reason that states or actions can themselves serve as reinforcing outcomes. That is, an action may lead to a state or another action that has acquired value, and the value of the subsequent state or action updates the value of the state or action immediately preceding it by being incorporated into the prediction error computation. The idea that signals for future reward can serve a similar function to primary reward is known as conditioned reinforcement (Williams, 1994). Most algorithms assume that conditioned reinforcers are discounted such that states or actions that are further away in time from primary reinforcement will make for weaker conditioned reinforcers.

In addition, when a reinforcer is not available immediately following an action, the action value can still be updated some time later if an eligibility trace is active at the time of a future reinforcer. This eligibility trace decays with time and functions as a memory for an action, allowing for delayed reinforcement to adjust action values. Eligibility traces and conditioned reinforcement bridge the temporal gap between actions and outcomes, helping to solve the temporal credit assignment problem (Suri & Schultz, 1998; Walsh & Anderson, 2011).

This framework, while powerful and effective at solving action sequencing learning problems, assumes that learning and decision-making processes are *model-free*. This means that the agent does not learn an internal model of the environment or task. The state-action values that are learned by a model-free system are summary statistics that represent how good a given action is in a given state, based on how reliably that action led to rewards in the past. As a result, a model-free agent can evaluate an action based on an estimate of how much future reward it is associated with, but that estimate is a cached value that is relatively robust to sudden changes in

the environment. Given these limitations, it is of interest to know whether model-free RL mechanisms can accurately describe features of action sequence learning in biological agents.

2.2 Model-free RL explains a fundamental feature of action sequence learning in animals

Arguably, one of the strongest reasons for believing that model-free RL has any validity in terms of animal behavior is the documented importance of prediction errors in Pavlovian learning, although whether these prediction errors are necessarily model-free has come under scrutiny recently (Keiflin et al., 2019; Li & McNally, 2014; Sharpe et al., 2017; Steinberg et al., 2013; Takahashi et al., 2009; Waelti, Dickinson, & Schultz, 2001). In Pavlovian learning outcomes are not contingent on actions, but rather the presence of external cues. The importance of prediction errors for instrumental learning has been investigated using free operant, single response paradigms (Dickinson & Charnock, 1985; Hammond & Weinberg, 1984; St. Claire-Smith, 1979; Williams, 1999), but the evidence is not as strong as in the Pavlovian case, and it has even been suggested that instrumental learning may adhere to a different set of rules than Pavlovian learning (Dickinson, 1994; Morris et al., 2017; Pérez et al., 2016). This broaches the issue of whether action sequence learning is governed by model-free prediction errors, and whether action sequencing can be understood in terms of model-free RL more generally.

One prediction that model-free RL makes about action sequencing is that the values of proximal actions (relative to the timing of reinforcement) will be learned more rapidly relative to the values of more distal actions. This should happen for several reasons. First, the assumption of a decaying eligibility trace implies that the degree to which an action is eligible for update becomes weaker the further it is temporally removed from a reinforcer. Second, a distal action

separated from a primary reinforcer by subsequent actions relies on more proximal actions to serve as conditioned reinforcers, but this form of reinforcement is relatively weak because the value of a conditioned reinforcer is discounted. Third, the prediction error that is hypothesized to update action values propagates back from the time of a reinforcer, taking longer to reach more distal actions.

The prediction of slower learning of distal actions is borne out across a variety of behaviors. For example, when mice are required to press a left lever and then a right lever in sequence to earn food rewards, they perseverate on the right lever early in training before learning the serial order of the task (Rothwell et al., 2015; Yin, 2010). Similarly, when mice are trained to press levers in a left-left-right-right pattern, their performance accuracy decreases for actions further back in the sequence (Geddes, Li, & Jin, 2018). When rats learn to navigate mazes, the number of erroneous entries into a blind alley increases as a function of the alley's serial distance from the food box (Spence, 1932; Tolman & Honzik, 1930), and human subjects navigating an imaginary maze make a similar pattern of errors (Fu & Anderson, 2006). When pigeons are required to emit a specific sequence of pecks across two keys on a variable interval schedule of reinforcement, the serial distance of the key from the reward affects how frequently that key is pecked (Catania, 1971). For example, when pigeons were required to execute the sequence left-right-right-right and right-right-left-right on different training days, the proportion of left key pecks emitted across the entire training session was greater in the latter case, even though the number of left and right key pecks required for reward was identical across both sequences.

A similar pattern of behavior holds for extinction, during which an expected reward is repeatedly withheld following the performance of a previously reinforced sequence. Morgan

(1974) documented a number of instances during which sequential actions extinguish at different rates depending on what order they are performed in the sequence. For example, in an experiment in which rats were trained to run down an alleyway for food reward and then food was withheld, running speed slowed down at a faster rate in the goalbox, where the food reward was previously delivered, compared to the more distant startbox (Wagner, 1961).

These behavioral observations suggest that actions that are more proximal to the time of reward are learned about sooner than distal actions (Figure 1A; see also Hull, 1932; Killeen, 1994). The implication is that sequential actions are learned ‘in reverse’ from the moment of reinforcement, and this is predicted by model-free RL. The mechanisms that allow for action sequences to be learned in this way—prediction error, conditioned reinforcement, eligibility traces—while powerful, nonetheless imply an impoverished model of learning and decision-making that does not completely match up to what is known to be true about action sequence learning in animals. Specifically, research has shown that action sequence learning involves more than just the learning of model-free state-action associations.

[Insert Figure 1 here]

2.3 Model-free RL does not explain goal-directed control of action sequences

The description of RL up to this point has intentionally been restricted to model-free learning. A model-free learner computes state-action values as it encounters reinforcing outcomes (or the absence of expected outcomes), caches those values in memory, and then subsequently retrieves those values to choose actions in the future. At the time of a decision, the agent is assumed to know what state it is in, and executing the appropriate action is as simple as

choosing the action with the highest value. This makes the decision-making process relatively fast and efficient. In contrast, there are other RL algorithms that fall under *model-based* methods that compute state-action values differently (Sutton & Barto, 2018). In brief, model-based RL relies on anticipating the outcomes of actions in terms of the probability of transitioning to future states and the rewards expected in those states. Note, however, that both behavioral strategies are fundamentally concerned with the sequential structure of actions, because the mechanisms by which actions are learned and selected depend on the value of subsequent actions in a sequence.

This distinction has bearing on the sequential decision-making of biological agents. Humans and non-human animals are capable of both model-based and model-free control, also known as goal-directed and habitual control, respectively (Daw et al., 2011; Dickinson, Nicholas, & Adams, 1983; Gillan et al., 2016; Gremel & Costa, 2013; Killcross & Coutureau, 2003). There are two common methods of revealing one or the other modes of behavioral control. One method involves training subjects to perform a set of actions for a reinforcing outcome, followed by devaluing the reinforcing outcome offline (i.e. outside of the learning situation), and then testing the subject's propensity to perform those actions when the outcome is withheld. If performance is suppressed following devaluation, the conclusion is that the animal is goal-directed—it can anticipate the consequences of its actions. On the other hand, if performance is unaffected by devaluation, then the animal is considered habitual—actions are controlled by the cached value of a stimulus-response association. Most often, this procedure is used to analyze single actions, rather than the sequential structure of actions. However, some work has examined how outcome devaluation affects action sequencing.

For example, Dickinson et al. (1983) and Killcross & Coutureau (2003) trained rats to press a lever and then enter a magazine for food rewards on various schedules of reinforcement.

The food rewards were then devalued either by conditioned taste aversion or selective satiation. They found that, while the expression of goal-directed lever pressing depended on the training schedule, magazine entering was consistently goal-directed regardless of the schedule (see also Balleine & Dickinson, 1991). This suggests that actions in a sequence are differentially goal-directed depending on the temporal proximity of the action to the outcome, with more proximal actions showing greater goal-directed control. A similar finding was reported by Balleine, Paredes-Oley, and Dickinson (2005). In that experiment, rats were trained to perform two sequential actions for food reward (lever press followed by chain pull or vice versa). A subset of rats then underwent reward devaluation via conditioned taste version. During the test the following day, performance of the action most proximal to reward was reduced but performance of the distal action was relatively frequent and did not differ from a control group that had not undergone the aversion treatment. In another study in which rats were trained to perform discriminated action sequences (e.g. panel light 1 → lever press → panel light 2 → chain pull → magazine entry → food pellet), only the magazine entry action was sensitive to devaluation (Thrailkill & Bouton, 2017). In a more recent study, rats that were trained to press levers in a fixed order were able to suppress performance of the sequence when rewards were devalued, and greater goal-directed control of the proximal action was observed but only in the case where rats had been extensively trained (Garr & Delamater, 2019). These findings collectively support the notion that action sequencing involves more than the learning of model-free state-action values. Rather, animals demonstrate knowledge of the consequences of their actions, although, in some cases, this knowledge seems to decay for actions further back in the sequence (Figure 1B; see also Daw, Niv, & Dayan, 2005; Corbit & Balleine, 2003).

An alternate way of differentiating model-free from model-based control is to use a complex action sequence task, known as the ‘two-step task’ (Daw et al., 2011). In brief, this task can differentiate model-free from model-based learning by analyzing whether subjects will simply repeat actions that previously led to reinforcement or whether actions will also be sensitive to the anticipated future state transitions. In humans, the finding is often that performance is a mixture of model-free and model-based learning (Daw et al., 2011; Dezfouli & Balleine, 2013; Doll et al., 2015; Friedel et al., 2014), while in rats and mice, behavior is either a mixture of the two learning strategies (Akam et al., 2017; Hasz & Redish, 2018) or completely model-based (Miller, Botvinick, & Brody, 2017). These findings, once again, strengthen the idea that action sequence learning cannot be described as a simple form of model-free RL.

2.4 Action sequences are hierarchically organized

While models of action sequence learning are greatly improved by incorporating model-based planning, that is not sufficient to capture all the complexities of action sequencing. Adding to that complexity is the idea that sequencing is hierarchical (Botvinick, 2008, 2012; Botvinick, Niv, & Barto, 2009; Dezfouli & Balleine 2012, 2013; Geddes et al., 2018; Lashley, 1951; Rosenbaum et al., 2007; Sutton, Precup, & Singh, 1999). The description of RL up to this point has been restricted to non-hierarchical control. Specifically, an agent that behaves in a non-hierarchical manner selects and evaluates actions one at a time, independent of preceding actions. The agent performs an action, transitions to a new state, and whatever action it selects in that state is independent of all the preceding actions. A hierarchical representation, in contrast, implies that actions are organized and represented as coherent subunits, such that the performance of a single action depends on the performance of a set of preceding actions.

Specifically, a hierarchical representation of an action sequence implies that individual actions are nested within higher order representations, which are referred to by a number of names—chunks (Graybiel, 1998), macro-actions (Dezfouli & Balleine, 2012), nested subroutines (Botvinick, 2008), and options (Sutton et al., 1999). This higher order representation of action sequences is hypothesized to form over the course of learning, with agents starting out evaluating and selecting actions piecemeal, and then eventually evaluating and selecting *groups* of actions (Dezfouli & Balleine, 2012).

What is the evidence that actions become chunked over the course of learning? Researchers often rely on ‘slips of action’ to indicate that chunking has occurred. For example, in a study conducted by Matsumoto and colleagues (1999), a monkey was trained to push three different buttons in sequence for a terminal water reward. Each time the correct button was pushed, the next button to be pressed in the sequence was illuminated to cue the monkey. After thousands of training trials on the same sequence of pushes, the experimenters abruptly changed the task so that reward was now delivered after the second button press, the third button press was no longer required as a response, and it stopped illuminating after the second button was pushed. The monkey, however, continued to press the third button after pressing the second button for tens of trials, and the latency of the third button push was similar before and after the task switch. This is also despite the fact that the reward and reward-associated clicker were presented immediately after the second button push. This result implies that, following extensive practice, sequential actions can become independent of the external cues that may have originally guided them. The authors interpreted this to mean that the entire button push sequence was pre-programmed, or chunked. While it is tempting to construe such slips of action as being the result of a chunking mechanism, it is also possible that such errors stem from strong stimulus-response

associations, in which the stimulus is no longer an external cue but proprioceptive feedback from the previous movement in the sequence (Horvitz, 2009).

Another approach that has been used to reveal chunking is to analyze how the structure of simple sequences changes over the course of training. For example, in one experiment mice were trained to press a single lever four times in a row with the constraint that mice were required to make progressively higher within-bout press rates across training sessions (Jin, Tecuapetla, & Costa, 2014). Across training sessions, the mean and variability of the inter-press intervals decreased, and the number of ultra-fast bouts increased. Based on this pattern of data, the authors concluded that mice learned to chunk their lever press actions together. Similarly, in another study in which mice were trained to press a single lever repeatedly, the mean and variability of the inter-press intervals once again decreased across training (Matamales et al., 2017). In addition to analyzing the inter-press intervals, the authors also examined the sequence boundary intervals—those intervals separating a terminal lever press from a magazine check and a magazine check from an initiating lever press. They found that the boundary intervals were, on average, much longer than the inter-press intervals in young mice, but in aged mice these two intervals were very similar. The authors concluded that, whereas young mice only chunked together the lever presses that fell within an uninterrupted sequence, aged mice chunked together lever presses and sequence boundary elements. Once again, however, these findings cannot definitively implicate chunking as playing a role in the observed behavior because improvements in response speed and variability alone can be explained by other, equally likely mechanisms (Niv et al., 2007; Pérez et al., 2016).

How, then, can the hierarchical nature of action sequencing be revealed? Recall that a defining feature of hierarchical control is that the sequence is “pre-programmed,” such that the

order of each action is determined prior to the initiation of the sequence. This means that the performance of an action within a sequence should be a function of the action that precedes it. To test this idea, Dezfouli and Balleine (2013) ran human participants in a version of the two-step task. The authors reasoned that, if a reward serves to reinforce an action chunk ($A1 \rightarrow A2$), then performing A1 in the future should entail a high probability of also performing A2 immediately after—even when participants are given information in the middle of the sequence that signals a more optimal action to be performed. In contrast, if a reward serves to reinforce single actions independently of one another, then selection of the second action should be independent of the first action performed. The pattern of data were consistent with the prediction that participants learned to chunk their actions together: performing A1 immediately after a rewarded $A1 \rightarrow A2$ sequence increased the probability of subsequently performing A2 even when it was not optimal.

While the task and analyses developed by Dezfouli and Balleine (2013) are able to reveal behavior that is consistent with a hierarchical representation of action chunks, the use of a non-hierarchical stimulus-response mechanism still cannot be ruled out. The finding that the second action in the sequence was conditional on the first action could either result from the sequence being pre-programmed (hierarchical representation) or the first action serving as a stimulus that triggers the second action (non-hierarchical representation). Thus, an additional method is needed for differentiating hierarchical from non-hierarchical representations of sequences. In one study, rats were trained to perform two sequences of lever presses, in which one sequence consisted of pressing a left lever followed by a right lever for one reward type (sucrose pellets or Polycose solution), while the other sequence consisted of pressing right followed by left for the other reward type (Ostlund, Winterbauer, & Balleine, 2009). Reward devaluation tests were then conducted after training. If rats had learned to chunk their lever presses together, then the entire

sequence associated with the devalued reward should have been selectively suppressed. This is what was found for normal rats. Interestingly, rats with secondary motor cortex lesions were unable to suppress the entire sequence, but instead selectively suppressed responding on the lever that was most temporally proximal to the devalued reward. The implication is that these rats treated each lever press as an independent behavioral unit, rather than part of a unified chunk.

2.5 Integrating hierarchical representations with model-free and model-based RL

The idea that actions can be chunked together implies that each action within the chunk is pre-determined as part of an entire set of actions and is therefore not decided upon individually. This idea is made clear in an RL model constructed by Dezfouli and Balleine (2012, 2013), in which they assume that all instrumental learning and decision-making is model-based, but rather than being limited to single actions, model-based RL can operate on entire sequences of actions. This means that, once an action sequence is selected in a model-based manner, the individual actions that make up the sequence are pre-determined and executed as part of a chunk. The implication is that the initiation of an action sequence is subject to goal-directed control, but that once the sequence is initiated, the individual actions that make up the sequence are executed habitually (see also Dezfouli, Lingawi, & Balleine, 2014). The data from the study by Ostlund et al. (2009) are consistent with goal-directed control of sequence initiation, because in that study rats were capable of selectively suppressing the initiation of a sequence whose outcome was devalued. This result does not constitute a full affirmation of the thesis put forth by Dezfouli and Balleine (2013), but it should be possible to use outcome devaluation to investigate this idea further.

To address this issue, Garr and Delamater (2019) trained animals to perform a sequence in order to explore how the extent of training would impact goal-directed control of action sequences. Rats were trained to press a left lever followed by a right lever for food rewards, and rats received either a moderate or extensive amount of training followed by reward devaluation tests. While overall sequence performance was clearly goal-directed, the locus of goal-directed control appeared to shift over the course of training. Specifically, rats given a moderate amount of training showed a greater latency to initiate a sequence when the associated reward was devalued, but the time to complete the sequence was unaffected by devaluation. The opposite pattern of behavior was observed in rats given extensive training: completion latencies, but not initiation latencies, were sensitive to devaluation. These findings run counter to the idea that, following extensive training of an action sequence, the execution of the component parts will become habitual. Rather, rats appeared to hesitate during the execution of a well-practiced sequence. It is possible, however, that this finding does not generalize across all action sequence learning situations. In the study by Garr and Delamater (2019), rats were required to complete a sequence by traversing a non-trivial distance—each lever was separated by approximately half the subject’s body length. Greater distances between actions lead to greater latencies, during which time goal-directed control could manifest (Hardwick et al., 2017).

The hierarchical structure of action sequences poses a challenge to the non-hierarchical RL algorithms commonly used to model sequential action learning and decision-making. Attempts have been made to invent new hierarchical RL models that frame the initiation and execution of sequences as being under the control of different computations (e.g. Dezfouli & Balleine, 2012, 2013), but it is uncertain whether these models accurately describe the computations underlying real animal behavior. While there is support for the notion that the

initiation and execution of a sequence are controlled by separate decision-making processes (Garr & Delamater, 2019), more work needs to be done to understand the details of those processes.

3. Contributions of the basal ganglia to action sequence learning

I next turn to an overview of the role of the basal ganglia in action sequence learning and performance. The basal ganglia are a set of evolutionary conserved subcortical nuclei that consist of a set of cortico-striato-pallido-thalamo-cortical loops, and the striatum in particular has been implicated in the control of movement and movement disorders for decades (Kish, Shannak, & Hornykiewicz, 1988; Marsden, 1980). Over time, though, it has become clear that the basal ganglia also play a role in the learning of movement, and there have been attempts to map RL algorithms onto basal ganglia circuits. In the following review, I have chosen to focus exclusively on circuits that interact with the dorsal striatum.

3.1 Mammalian basal ganglia circuit organization and function

The main input nucleus of the basal ganglia, the striatum, receives excitatory glutamatergic projections from almost all areas of cortex and some regions of the thalamus (Díaz-Hernández et al., 2018; Guo et al., 2015; Wall et al., 2013). The striatum is part of two prominent neural pathways: the direct pathway and the indirect pathway (Figure 2). A large body of evidence suggests that these two pathways function in antagonistic ways (Calabresi et al., 2014; Nelson & Kreitzer, 2014). In the direct pathway, D1 dopamine receptor-expressing medium spiny neurons (MSNs) in the striatum receive glutamatergic projections from cortex and

thalamus and send GABAergic projections to the two basal ganglia output nuclei: the substantia nigra pars reticulata (SNr) and the entopeduncular nucleus (EP) as it is known in rodents, or the internal segment of the globus pallidus (GPi) as it is known in primates. The SNr and EP/GPi, in turn, send GABAergic projections to the ventral thalamus, and the excitatory thalamic projections back to cortical motor areas complete the loop. The net result of activating the D1-expressing MSNs is the disinhibition of neurons in the motor cortex that facilitate movement (Oldenburg & Sabatini, 2015). In the indirect pathway, D2 dopamine receptor-expressing MSNs also receive excitatory signals from the cortex and thalamus, but connect with the basal ganglia output nuclei in an indirect way. The GABAergic D2 receptor-expressing MSNs synapse onto the external segment of the globus pallidus (GPe), which makes GABAergic connections with the subthalamic nucleus, which makes glutamatergic connections with the SNr and EP/GPi. The net result of activating the D2-expressing MSNs is the inhibition of neurons in the motor cortex and suppression of movement (Oldenburg & Sabatini, 2015). In addition to participating in the looped circuitry, the SNr also send outputs directly to brainstem nuclei that control behavior in primates, rodents, and even lamprey (Grillner & Robertson, 2015; Hikosaka et al., 2006; Lalive et al., 2018; Roseberry et al., 2016).

[Insert Figure 2 here]

The two pathways are able to work in synchrony due to the opposite effects that dopamine has on each MSN type. Stimulation of D1 receptors on MSNs of the direct pathway promotes depolarization, while stimulating D2 receptors on MSNs of the indirect pathway stunts depolarization (Gerfen & Surmeier, 2011; Tritsch & Sabatini, 2012). Postsynaptic D1 receptor activation at corticostriatal synapses of the direct pathway has the effect of inducing long-term potentiation (LTP) so long as postsynaptic NMDA receptors are activated, presynaptic firing

precedes postsynaptic firing, and dopamine release quickly follows presynaptic firing (Calabresi et al., 2007; Shen et al., 2008; Wickens, Bregg, & Arbuthnott, 1996; Yagishita et al., 2014). If the preceding conditions are met but dopamine is absent, these synapses undergo long-term depression (LTD; Pawlak & Kerr, 2008; Shen et al., 2008; Wickens et al., 1996). In contrast, pairing presynaptic with postsynaptic activation during postsynaptic D2 receptor activation results in LTD (Kreitzer & Malenka, 2007; Lerner & Kreitzer, 2012; Shen et al., 2008). Thus, when dopamine is released in the dorsal striatum from the ascending fibers of the substantia nigra pars compacta (SNc), the direct pathway is strengthened while the indirect pathway is weakened. When dopamine levels in the striatum drop below baseline, the direct pathway is weakened and the indirect pathway is strengthened.

While this model of basal ganglia circuit organization has proven useful, the circuitry is more complex than originally formulated. First, the various cortical inputs to the basal ganglia are unevenly distributed across the direct and indirect pathways, with the former receiving relatively more input from somatosensory and limbic areas and the latter receiving more input from motor cortices (Wall et al., 2013). Second, the dorsomedial and dorsolateral regions of the striatum (caudate and putamen in primates) play distinct roles in instrumental learning and receive projections from distinct SNc dopamine neurons in mice (Corbit, Nie, & Janak, 2012; Lerner et al., 2015; Yin et al., 2004, 2005). Third, the interneurons within the dorsal striatum, while comprising only about 5% of all neurons, play a critical role in regulating MSN excitability and plasticity (Burguière, Monteiro, Feng, & Graybiel, 2013; English et al., 2011; O'Hare et al., 2017; Schulz & Reynolds, 2013). Fourth, as Calabresi and colleagues (2014) point out, the direct and indirect pathways are not completely segregated. Many striatal neurons that project to the SNr also send collaterals to the GPe (Cazorla et al., 2014), and a small proportion

of striatal neurons co-express D1 and D2 receptors (Bertran-Gonzalez et al., 2010; Gagnon et al., 2017; Nadjar et al., 2006). They also point out that both pathways are subject to D2 receptor-mediated LTD via cholinergic interneurons, as well as D1 receptor-mediated LTD via interneurons containing nitric oxide synthase.

There are also findings that directly contradict model predictions. For example, destruction of SNc dopaminergic neurons in monkeys results in enhanced spontaneous MSN activity in both pathways, while the traditional model predicts that only D2 MSNs should become more excitable (Liang et al., 2008). Another problem arises from the finding that some SNr neurons become excited, rather than inhibited, by D1 MSN stimulation (Freeze et al., 2013). The same researchers also found that some SNr neurons become inhibited, rather than excited, by D2 MSN stimulation. Additionally, in one study, inhibitory input from the EP/GPi was found to induce excitatory, rather than inhibitory, responses in ventral thalamic neurons (Kim et al., 2017). It is clear from these findings that a more complete model of the basal ganglia is needed. Yet it is exactly because of this incompleteness that it is important to study how the basal ganglia function in a variety of behavioral contexts, including instrumental learning and action sequencing.

3.2 Basal ganglia contributions to movement and instrumental learning

A common conception of how the direct and indirect basal ganglia pathways contribute to behavior is the ‘go/no-go’ model, in which the direct pathway is identified as providing a ‘go’ signal and the indirect pathway a ‘no-go’ signal (e.g. Collins & Frank, 2014; Frank, Seeberger & O’Reilly, 2004). There is a set of simple behavioral predictions that come out of this model. One

prediction is that activation of D1 MSNs should lead to increased movement while activation of D2 MSNs should lead to decreased movement. In support of this hypothesis, Kravitz and colleagues (2010) found that excitatory optogenetic laser stimulation of D1 MSNs in mice increased the frequency of ambulation in an open arena, while stimulation of D2 MSNs increased the frequency of freezing. In a complementary study, optogenetic silencing of D1 MSNs during instrumental joystick movements in mice resulted in reduced movement velocity (Panigrahi et al., 2015). In accordance with these findings, ablation of D1 MSNs results in bradykinesia and dystonia (Drago et al., 1998), while ablation of D2 MSNs produces hyperactivity (Durieux et al., 2009).

The ‘go/no-go’ model can also be expanded to make predictions about instrumental learning, since this form of learning entails knowing when (‘go’) and when not (‘no-go’) to perform specific movements. The specific hypothesis is that D1 MSNs mediate reinforcement while D2 MSNs mediate punishment. In support of this hypothesis, Kravitz and colleagues (2012) found that mice given the opportunity to press a lever for excitatory optogenetic D1 MSN stimulation in the dorsomedial striatum (DMS) pressed significantly more often compared to a control lever, while mice given the opportunity to press for excitatory optogenetic D2 MSN stimulation showed the opposite preference. Notably, these preferences gradually disappeared when stimulation was withheld during an extinction phase. Similarly, in mice that were trained to move a joystick for water rewards, optogenetically stimulating D1 MSNs in the DMS only during fast movements caused mice to gradually increase their movement velocity, but movement velocity gradually decreased when D2 MSNs were stimulated under the same conditions (Yttri & Dudman, 2016). Conversely, when stimulation was given only during slow movements, mice gradually decreased movement velocity if D1 MSNs were stimulated and

gradually increased velocity if D2 MSNs were stimulated. These effects also disappeared when stimulation was withheld. Additionally, in a two-alternative forced-choice task with mice, Tai and colleagues (2012) found that optogenetically stimulating D1 MSNs in the DMS of mice during a choice period biased choices toward a reward port contralateral to the stimulated hemisphere while D2 MSN stimulation biased choices away from the contralateral port. These finding provides strong support for the role of the direct and indirect pathways via the DMS in reinforcement and punishment, respectively.

A more complicated account holds for optogenetic stimulation in the dorsolateral striatum (DLS). Vicente and colleagues (2016) demonstrated that mice will press a lever for optogenetic D1 MSN stimulation in the DLS more than an inactive lever (and more than non-stimulated control mice), but that D2 MSN stimulation also reinforces lever pressing above non-stimulated controls. Notably, however, mice were very slow to acquire lever pressing for D2 stimulation, the frequency of lever pressing was very low, and these mice pressed the laser-paired lever just as often as the inactive lever. A more recent study showed that D1 but not D2 receptors in the DLS mediate LTP during early rotarod learning in mice (Giordano et al., 2018).

Aside from simple acquisition of responding, the dorsal striatum is also involved in more complex forms of instrumental learning, such as goal-directed control. For example, in the DMS, lesions, NMDA receptor antagonism, and disconnections with prelimbic cortex disrupt instrumental sensitivity to reward devaluation (Gremel & Costa, 2013; Hart, Bradfield, & Balleine, 2018; Yin et al., 2005a, 2005b). Conversely, DLS lesions, rather than interfering with goal-directed control, result in a disruption of habit formation (Gremel & Costa, 2013; Yin, Knowlton, & Balleine, 2004). The dissociation between DMS and DLS function has led to the hypothesis that the cortico-basal ganglia loops that pass through these different regions of the

dorsal striatum are involved in distinct computations related to model-based and model-free RL, respectively (Bornstein & Daw, 2011; Daw et al., 2005; Khamassi & Humphries, 2012).

Since the dorsal striatum is clearly involved in instrumental learning, and some aspects of instrumental learning can theoretically be described by a set of RL computations (see section 2.1), it is of interest to know whether activity in the striatum and/or its afferents resembles those computations. In the section below, I review the potential significance of striatal dopamine in conveying prediction error signals, which could possibly serve as teaching signals to adjust the strength of corticostriatal synapses and thereby adjust the values of state-action pairs just like an RL algorithm. The implication is that, if the basal ganglia implement RL-like computations, then it could serve as a neural substrate for action sequence learning.

3.3 Mapping RL variables to striatal physiology

It is commonly thought that the role of the dorsal striatum in instrumental learning is mediated by dopamine input from the SNc. The assumption is that dopamine release serves as a teaching signal that modulates corticostriatal synapses, the strength of which determines the strength of instrumental learning (Collins & Frank, 2014; Wickens, Reynolds, & Hyland, 2003; Wickens et al., 2007). This assumption is based on the so-called “three-factor Hebbian rule,” which posits that instrumental learning depends on three requirements: presynaptic cortical activation, followed by postsynaptic striatal activation, and coincident dopamine release onto the striatal neuron. Some of the strongest support for this theory comes from an experiment in which it was shown that (1) simultaneously stimulating SNc and striatal MSNs potentiated corticostriatal synapses, (2) injections of a D1 receptor antagonist attenuated this plasticity, and

(3) the degree to which corticostriatal synapses were potentiated was negatively correlated with the time taken to learn to respond for SNc self-stimulation (Reynolds, Hyland, & Wickens, 2001). It is therefore plausible that dopaminergic input to the dorsal striatum could modulate the connection weights between cortical and striatal neurons, which consequently could facilitate instrumental learning.

The three-factor rule has led theorists to hypothesize that the three neurophysiological factors contributing to instrumental learning—coincidental cortical, striatal, and dopaminergic activation—can be identified with three variables from RL: states, actions, and prediction errors, respectively. The idea is that striatal dopamine release encodes prediction errors, which are used to update the values of state-action pairs by modifying corticostriatal synapses, with the incoming cortical inputs representing state information and the striatal neurons representing specific actions. Perhaps the most well-supported of these assumptions is that midbrain dopamine activity reflects prediction error computations. A copious number of studies have shown that the activity of midbrain dopamine neurons resembles a prediction error (Bayer & Glimcher, 2005; Eshel et al., 2015; Schultz, Dayan, & Montague, 1997; Starkweather et al., 2017; Takahashi et al., 2017; Waelti et al., 2001), and that optogenetically modulating dopamine neuron activity mimics the effects of endogenous prediction errors during learning (Chang et al., 2016; Keiflin et al., 2019; Steinberg et al., 2013). There is also evidence that cortical inputs, particularly from the orbitofrontal cortex, signal state information (Sharpe et al., 2019; Wilson et al., 2014), and that the activity of individual striatal MSNs correlates with specific actions (Klaus et al., 2017). Thus, when an animal performs an action in some state and earns an unexpected reward, a set of cortical, striatal, and dopaminergic neurons are co-active at that moment. Assuming that the dopaminergic terminals in the striatum release dopamine in proportion to the

cell firing rate, the synapses between the co-active cortical neurons and D1 MSNs will be strengthened in proportion to the size of the prediction error, while cortico-D2 MSN synapses will be weakened owing to the opposing effects dopamine has on each striatal MSN type. The net result of this neuromodulation is that, when the animal encounters that same state in the future, the corresponding cortical neurons will more readily activate the set of D1 MSNs that drive the action that led to reward in the past while weakly activating the set of D2 MSNs that suppress the action (Frank, 2005). Moreover, since dopaminergic prediction errors are computed during all phases of an action sequence, and not just at the time of reward, this mechanism could theoretically support the learning of a whole sequence of actions.

3.4 Questioning the validity of RL models of basal ganglia function

The neurocomputational framework described above, while theoretically useful and intriguing, faces significant uncertainties. The primary uncertainty is whether dopamine release in the dorsal striatum reflects a prediction error computation. Much of the evidence for dopaminergic prediction error signals comes from studies of the ventral tegmental area (VTA), which sends dopaminergic projections to the ventral striatum and receives different inputs compared to SNc dopamine neurons (Watabe-Uchida et al., 2016). It is possible that the signals conveyed by VTA dopamine neurons and their terminals in the ventral striatum may differ from those conveyed by SNc dopamine neurons and their terminals in the dorsal striatum. In support of this notion, Parker and colleagues (2016) found that the activity of dopamine terminals in the dorsal striatum was strongly correlated with contralateral movements and weakly correlated with prediction errors, while the opposite was true of terminals in the ventral striatum. In another study, rats learned to press a lever for a cue that was previously paired with VTA dopamine cell

body stimulation or the terminals in the nucleus accumbens core, but stimulation of SNc dopamine cell bodies or the terminals in the dorsal striatum failed to imbue cues with this conditioned reinforcing property (Saunders et al., 2018). Further evidence for a functional difference between midbrain dopamine populations comes from a study showing that optogenetic stimulation of dopamine neurons in SNc failed to induce a prediction error as deduced by a failure to unblock a Pavlovian cue, while this was not the case for VTA dopamine neuron stimulation (Keiflin et al., 2019).

The evidence for dopamine prediction error computation comes almost exclusively from Pavlovian learning paradigms, which raises the question of how dopamine functions during instrumental learning. While it is known that rodents can learn to self-stimulate for VTA and SNc dopamine neuron excitation (Ilango et al., 2014; Keiflin et al., 2019; Saunders et al., 2018; Witten et al., 2011), self-stimulation does not unambiguously implicate a role for dopamine in prediction error computation. In a typical appetitive Pavlovian learning scenario, in which a reward is preceded by a cue or series of cues, the canonical neural signature of a prediction error computation is the initial increase in activity in response to the reward, followed by a gradual decrease over subsequent trials and a concurrent, gradual increase in response to the preceding cues (Pan et al., 2005; Schultz, 1998). One could argue that, if a similar pattern was observed during instrumental learning, this would go a long way toward verifying the role of dopamine in RL prediction error computation. To this end, two studies measured dopamine release using fast-scan cyclic voltammetry in the striatum during action sequence learning in rats and found that, consistent with a role for striatal dopamine in prediction error computation, dopamine release in response to reward diminished over training (Collins et al., 2016; Wassum, Ostlund, & Maidment, 2012). However, contrary to the prediction error hypothesis, the peak dopamine

response around the time of the initiating action did not change over training, and dopamine release frequently ramped up such that the onset of the ramp changed over the course of training—migrating from around the time of reward to before the initial action. These features of dopamine release, rather than reflecting prediction error encoding, may be better explained by time-discounted state value encoding (Berke, 2018).

Whatever the interpretation of the dopamine release ramps reported Wassum et al. (2012) and Collins et al. (2016), the fact is that voltammetry was performed only in the ventral striatum. It is unknown whether similar patterns would be observed in the dorsal striatum. Even if a similar pattern of neural activity was observed in the dopamine terminals in the dorsal striatum, it is unclear to what degree that pattern would reflect a learning process. One solution to this problem is to provide causal evidence for nigrostriatal dopamine in the acquisition of an action sequence. One such attempt was made in a preliminary study (Keiflin, Pribut, & Janak, 2018). In this experiment, rats were trained to perform a sequence of actions for optogenetic dopamine neuron stimulation, either in the VTA or SNc. Rats first began in phase 1 by pressing a lever (lever 1) for immediate stimulation. Then, in phase 2, access to lever 1 was made contingent on a single press of another lever (lever 2). Finally, in phase 3, access to lever 2 was made contingent on a nose-poke in the back of the chamber, so that, by the final phase of training, rats were required to perform a three item sequence (nose-poke, lever 2, lever 1) for dopamine neuron stimulation. For rats that received VTA stimulation, the mean number of earned stimulations declined from one training phase to the next, which probably reflects the difficulty in learning a lengthy action sequence, but nevertheless remained well above zero for all rats. But for rats that received SNc stimulation, the mean number of stimulations dropped to nearly zero from phase 1 to phase 2, and stayed at around zero from phase 2 to phase 3. In other words, VTA dopamine

stimulation supported the learning of an action sequence while SNc stimulation did not. This finding casts doubt on the idea that SNc dopamine provides the prediction error signal necessary for action sequence learning.

Finally, an additional problem for mapping RL variables to striatal physiology is the absence of plastic changes in the dorsal striatum following instrumental learning. In a recent study, mice were trained to nose-poke for optogenetic D1 MSN stimulation in the DMS, and this instrumental response was quickly acquired—responding was significantly higher compared to an inactive port, and mice without the opsin showed very few nose pokes at either active and inactive ports (Lalivie et al., 2018). Yet, when the researchers looked for signatures of changes in synaptic plasticity in the very same neurons that were stimulated during learning, they could not find any evidence of plastic changes. Rather, the authors' data suggest that the plasticity supporting instrumental acquisition of D1 MSN stimulation may reside in the synapses between the SNr and ventral thalamus, which suggests a dopamine-independent mechanism of instrumental learning.

Given the uncertainty surrounding prediction error coding in dorsal striatum-projecting dopamine neurons, along with the failure to uncover plastic changes in the dorsal striatum after the acquisition of D1 MSN self-stimulation, the exact role of the dorsal striatum during action sequence learning is unclear. The neurocomputational theory that combines the three factor rule with RL may, ultimately, be unsubstantiated. At minimum, if the dorsal striatum is required for action sequence learning then lesioning it should disrupt sequence acquisition. One study found that DLS lesions, but not DMS lesions, in mice interfered with the acquisition of an action sequence, but this learning deficit could have been caused by a motor deficit that lengthened the time between actions, which could have affected learning (Yin, 2010). Another study found that

temporary inactivation of the anterior caudate and putamen in monkeys (DMS in rodents) disrupted the learning of a new button-push sequence but spared the execution of a familiar sequence, while inactivation of the middle and posterior putamen (DLS in rodents) had the opposite effect (Miyachi et al., 1997). This study, having used only two subjects and precluded measurements of movement kinematics, deserves replication. One general conclusion, though, is that loss-of-function studies in the basal ganglia make it difficult to tease apart effects of learning versus performance, since proper performance is often required for normal learning to proceed. It is in this respect that detailed behavioral analysis and transient inactivation techniques, such as DREADDs (Smith et al., 2016), will be useful.

An alternative approach to modelling how the basal ganglia participates in action sequence learning may require incorporating temporal estimation as a key computation of dopaminergic and striatal circuits. There is an abundance of data that argue in favor of the dorsal striatum and its dopaminergic afferents as critical nodes in the circuitry supporting interval timing (De Corte et al., 2019; Gouvêa et al., 2015; Howard et al., 2017; Matell, Meck, & Nicolelis, 2003; Meck, 2006; Mello, Soares, & Paton, 2015; Soares, Atallah, & Paton, 2016), and learning the correct ordering of actions may entail learning the relative temporal distance between actions and outcomes. It has been shown that dopamine release in the dorsal striatum correlates with temporal proximity to reward during maze running in rats (Howe et al., 2013), and SNc dopamine neuron firing rates in mice control the subjective estimation of time in interval discrimination tasks, with greater firing rates leading to shorter perceived intervals between trial start time and reward receipt (Howard et al., 2017; Soares et al., 2016). In the dorsal striatum, neurons are tuned to the relative time within a fixed interval and display precise dynamics that correlate with the accuracy of temporal judgements (Gouvêa et al., 2015; Mello et

al., 2015). An intriguing hypothesis is that, when an animal performs a sequence of movements, the striatum keeps track of action-outcome intervals by integrating information about the current action being performed and the estimated time relative to expected future reward—information potentially conveyed by the cortex and SNc, respectively. Indeed, one study with mice found that the firing rates of dorsal striatal neurons surrounding the initiation of a bout of lever presses during a fixed interval depended on the temporal distance from future reward (Mello et al., 2015). Learning when each action occurs relative to a rewarding outcome would allow the brain to decode the correct sequence of actions that previously lead to reward, thus reinforcing the relative timing of different actions.

4. Contributions of the basal ganglia to action sequence performance

Once action sequences have been learned they can be performed with astonishing fluidity. What are the neural mechanisms that enable the fluid performance of action sequences? In the following sections, I will review research that has illuminated how the basal ganglia circuitry controls the initiation, execution, and termination of action sequences.

4.1 Initiation

It is known that MSNs within the DLS fire phasically in response to the onset of a sequence of actions (Barnes et al., 2005; Jog et al., 1999; Thorn et al., 2010; Smith & Graybiel, 2013). This finding comes primarily from studies that use rats in a T-maze task, in which the opening of the start gate and the onset of locomotion occur very close in time and thus obscure the cause of the neural activity. However, in recent years there have been efforts to study

sequence initiation activity in free operant, self-paced tasks in which rodents are allowed to initiate sequences of lever presses at any time throughout a session under simple ratio schedules. These studies have confirmed the existence of phasic spiking activity in the dorsal striatum surrounding the initiation of a bout of lever presses (Cui et al., 2013; Geddes et al., 2018; Jin & Costa, 2010; Jin et al., 2014). Paradoxically, this phasic activity is reflected in both D1 and D2 MSNs (Cui et al., 2013; Jin et al., 2014), and optogenetically inhibiting or activating D1 or D2 MSN activity prior to lever pressing results in an increased latency to initiate a bout of lever presses (Tecuapetla et al., 2016).

These findings pose problems for the traditional model of basal ganglia function because activation of D1 and D2 MSNs are predicted to have opposite effects on movement, with activation of D1 MSNs promoting movement and activation of D2 MSNs inhibiting movement. One hypothesis is that the phasic D1 MSN spiking that occurs at the beginning of an action sequence selects the desired motor program while phasic D2 MSN spiking inhibits competing motor programs, and that any abnormal change in firing patterns—whether due to inhibition or activation—will disrupt sequence initiation (Jin & Costa, 2015; Tecuapetla et al., 2016). Consistent with this hypothesis, manipulations of D1 and D2 MSNs retard the initiation of lever press bouts for different reasons. While D1 MSN inhibition or excitation results in temporary freezing between the exiting of the food magazine and the pressing of the lever, D2 MSN inhibition or excitation causes mice to leave the area and explore the chamber (Tecuapetla et al., 2016). Consistent with these findings, Geddes et al. (2018) reported that optogenetic excitation of D1 and D2 MSNs prior to sequence initiation resulted in delayed initiation latencies, but while D1 MSN stimulation merely delayed the start of the sequences, D2 MSN stimulation abolished the first half of the sequence. Collectively, these findings argue against the idea that the basal

ganglia permit movement via a read-off of the relative firing rates of the direct and indirect pathways. That is, sequence initiation cannot simply be predicted by how active the direct pathway is relative to the indirect pathway. Instead, the basal ganglia may very well operate on a neural code that depends more on the specific spiking patterns and relative timing of spikes in each pathway (O'Hare et al., 2016; Schmidt et al., 2013).

Sequence initiation also appears to depend critically on nigrostriatal dopamine. In one recent study with mice, it was found that SNc dopamine neurons fire in a burst prior to the initiation of spontaneous movement in an open field (da Silva et al., 2018). Optogenetically inhibiting dopamine neurons during periods of immobility reduced the acceleration and probability of subsequent movement, while also increasing the latency to move. Remarkably, if the animal was already moving and the laser was activated, there was no detectable change in movement. A similar pattern was observed during optogenetic excitation: laser activation during periods of immobility resulted in quick increases in acceleration while laser activation during ongoing movement did not. Similar patterns of dopamine neuron activity were also detected during the initiation of a sequence of instrumental lever presses, and inhibiting dopamine neuron activity just prior to the initiation of the sequence resulted in increased initiation latencies, while activating the laser during sequence execution did not change pressing rates. The phasic activity of SNc dopamine neurons thus appears to drive the initiation, but not execution, of action sequences, and could possibly serve as the source of the aforementioned increase in MSN firing rates observed prior to sequence initiation (see also Howe & Domback, 2016).

Although there is a multitude of cortical regions that innervate the striatum, one of these regions appears to be particularly important for action sequence initiation. The secondary motor cortex (area M2) in rodents projects to the DLS and has been shown to be important in sequence

initiation. In an experiment reported by Bailey and Mair (2007) rats with either M2, M1, or sham lesions were trained to make a series of nose pokes in adjacent ports for a terminal reward. After pressing a lever, rats ran down an alley and crossed a photo beam that led into an arena containing five ports. Once the beam was crossed, one of the five ports was illuminated. If the rat entered the illuminated port, another port was illuminated, and the cycle continued until five nose pokes were made into the appropriate ports, at which time reward was delivered in the last port entered. While some sessions contained trials in which the order of the illuminated ports was random, other session contained trials in which the order was repeated and thus highly predictable. For all groups, the time to complete the sequence was significantly shorter during the sessions containing repeated sequences. For the M1 and sham groups, the time to initiate sequences remained constant across trial types. However, M2 lesioned animals took longer to initiate repeated sequences compared to random sequences. That is, M2 lesions interfered with sequence initiation while sparing the advantage in execution time conferred by repetition. However, since this study explored the effects of pretraining lesions it is difficult to know whether the effects reported in this study reflect impairments in sequence learning or performance.

In another study in which mice were required to press a left lever and then a right lever in sequence for reward, silencing neural activity specifically in the M2-DLS pathway (but not the M1-DLS pathway) after learning impaired the ability of mice to choose the correct first step of a two-lever sequence, resulting in higher instances of right-right sequences (Rothwell et al., 2015). A separate experiment from the same paper showed that optogenetic excitation of the M2-DLS pathway prior to sequence initiation improved initiation accuracy, resulting in a higher probability of choosing the correct first lever. These results suggest that M2 neurons may

communicate information to the DLS about the correct sequence to be performed. Consistent with this hypothesis, neurons in the primate SMA—which is homologous to rodent area M2—have been shown to fire prior to the initiation of a specific sequence of button presses (e.g. before push-turn-pull but not before push-pull-turn; Tanji, 2001). Given the finding that D1 and D2 MSNs are simultaneously active at the time of sequence initiation (Cui et al., 2013; Jin et al., 2014), it would be particularly interesting to investigate whether exciting or inhibiting D1 and D2 MSNs via M2 projection neurons just prior to sequence initiation results in similar improvements or decrements in sequence performance, respectively.

4.2 Execution

Once a well-learned sequence is initiated, the subsequent movements can be executed one after another without deliberation. William James postulated that a well-learned series of movements is “nothing but concatenated discharges in the nerve-centres, due to the presence there of systems of reflex paths, so organized as to wake each other up successively” (James, 1890). What James advocated was a chain theory of sequence execution in which neurons deterministically activate other neurons in a chain. Partial evidence for this ‘synaptic chain’ theory comes from studies of songbirds that rattle off learned sequences of song syllables in rapid succession. In zebra finches, individual neurons in area HVC—a premotor cortical area with connections to the striatum—produce short bursts of spikes at a specific time during a song motif, with different neurons firing at different times with almost no temporal overlap (Hahnloser et al., 2002; Okubo et al., 2015). There is tentative support for the hypothesis that the sequential bursting in HVC arises from the propagation of neural activity through a chain of connected neurons (Long et al., 2010).

One idea is that the HVC bursts are not motor signals that drive vocal sequence execution, but rather serve as state signals to be used in RL computations (Fee, 2014). An alternative idea is that the sequential activity in HVC is driven by short and temporally precise bursts of striatal MSNs, which influence the activity of individual HVC neurons by transmitting information about specific vocal gestures through the pallido-thalamo-cortical network (Tanaka et al., 2016). This implies that individual MSNs fire brief bursts of action potentials at one unique time within the vocal sequence, just as HVC neurons do.

The idea that individual MSNs in the striatum spike at one unique time during action sequence execution does not seem to hold up in the mammalian basal ganglia. For example, when mice perform fast bouts of lever presses under a fixed ratio schedule, most MSNs in the DLS show generally sustained or inhibited activity throughout the duration of the sequence rather than punctuated bursts (Jin et al., 2014). Not surprisingly, the MSNs that show sustained activation belong to the D1 class of MSNs while the MSNs showing inhibition belong to the D2 class. That the smooth execution of action sequences requires heightened D1 MSN activity and dampened D2 activity is corroborated by the finding that optogenetic excitation of D1 MSNs in the mouse DLS during lever press bouts increases the within-bout press rate, so long as stimulation matches the natural average MSN firing rate (Tecuapetla et al., 2016). As expected, optogenetic excitation of D2 MSNs during bout execution results in increased frequency of bout cancellations (Tecuapetla et al., 2016). Oddly, though, dampening the activity of D2 MSNs during bout execution also increases the frequency of bout cancellations (Tecuapetla et al., 2016). It has also been found that permanent inactivation of D1 and D2 MSNs in the mouse DLS after learning increases and decreases, respectively, the probability of making a correct second lever press given a correct initial lever press (Rothwell et al., 2015). These findings are generally

consistent with the traditional model of basal ganglia function, which posits that relatively high D1 MSN activity and low D2 MSN activity promotes movement via inhibition of the SNr. Correspondingly, SNr neurons show more inhibited activity during the execution of lever bouts than GPe neurons, which in turn show a higher prevalence of sustained activity (Jin et al., 2014).

The interpretation of the broadly sustained and inhibited activity of MSNs observed by Jin et al. (2014) during sequence execution suffers from the drawback that neural activity was recorded during the performance of homogenous sequences, during which identical actions are repeated (e.g. bouts of presses on a single lever). An MSN that shows broadly sustained or inhibited activity during a bout of identical movements may very well display a different activity profile during the performance of a heterogeneous sequence—perhaps activity similar to the precisely time-locked bursts seen in songbird HVC. When dorsal striatal neurons have been recorded during the performance of a heterogeneous sequence, activity has indeed been shown to be unique to specific parts of the sequence. In one study in which rats were trained to perform sequences of three lever presses across two separate levers (Martiros, Burgess, & Graybiel, 2018), the authors describe cases in which MSNs that burst around the time of a lever press during a correct sequence did not burst when that same lever press was executed as part of an incorrect sequence. While this result is consistent with a chaining mechanism, it is not clear whether actions were executed the same way across correct and incorrect sequences (e.g. similar kinematics and posture during movement).

An additional feature of striatal activity reported by Martiros et al. (2018) is that the MSN population showed ‘task bracketing’ activity during which the firing rates increased around the time of the first and last lever presses. Another study also confirmed that MSNs in the mouse dorsal striatum burst exclusively during the beginning or end of a heterogeneous sequence of

lever presses (left-left-right-right), but that this pattern of activity characterized only a subset of MSNs (Geddes et al., 2018). Other MSNs showed sustained or inhibited firing rates throughout the entire sequence, with most sustained neurons belong to the D1 class and most inhibited neurons to the D2 class. In addition, it was shown that a population of D2 MSNs changed their firing rates significantly during the transition from the left to right subsequences, suggesting that switching between elements of a sequence may require a ‘stop and switch’ signal. Indeed, optogenetically stimulating D2 MSNs after the first lever press in the sequence resulted in mice prematurely switching to the right lever. This finding raises the interesting possibility that action-related bursting during sequence execution may not causally support the generation of that action, but rather the stopping of the previous action.

The studies reviewed so far indicate that, if action sequence execution is governed by a chaining mechanism, in which sequential actions are generated by a string of sequentially active connected neural ensembles, it is unlikely that the synaptic chains that support sequence execution reside in the striatum. From a purely behavioral perspective, though, a chaining theory is not sufficient to capture the complexities of action sequence execution. This has been recognized since the time of von Holst (von Holst, 1937/1973). The most rudimentary version of a chaining theory posits that actions within a sequence cannot be executed unless the appropriate preceding action has been performed, since each action serves as a cue for its subsequent action. Several findings that contradict this idea come from the simultaneous chain procedure (Terrace, 2005). The procedure, usually conducted with primates or pigeons, involves displaying many different visual stimuli on a screen simultaneously. If the subject contacts each stimulus in the correct order, reward is earned. In one variant of the procedure (Chen, Swartz, & Terrace, 1997), monkeys were trained on four different simultaneous chains, composed as follows:

$$A_1 \rightarrow B_1 \rightarrow C_1 \rightarrow D_1$$

$$A_2 \rightarrow B_2 \rightarrow C_2 \rightarrow D_2$$

$$A_3 \rightarrow B_3 \rightarrow C_3 \rightarrow D_3$$

$$A_4 \rightarrow B_4 \rightarrow C_4 \rightarrow D_4$$

Each letter refers to a different visual stimulus on which the monkeys had to press, and each subscript refers to the list number. After training, new stimulus displays were created in which stimuli from each list were combined such that the correct sequence maintained the same ordinal relationship among the actions (e.g. $A_1 \rightarrow B_2 \rightarrow C_3 \rightarrow D_4$) or switched the previous ordinal positions (e.g. $B_1 \rightarrow A_2 \rightarrow D_3 \rightarrow C_4$). Monkeys learned chains that maintained ordinal relationships at a faster rate, indicating that they had not simply associated consecutive responses with one another. Rather, they were capable of learning about the serial order of actions within a sequence, which is a feature that is not captured by chaining theories.

Another finding opposing chaining theories is that actions in a sequence can be cued in advance of their immediately preceding action. Behmer and Crump (2017) instructed human participants to read paragraphs on a screen and type the letters on a keyboard, during which each letter to be typed was cued by red coloring. Often, the cued letter was the one immediately adjacent to the most recent letter, but sometimes more distant letters were cued (e.g. if the word is “explain,” and the most recently typed letter was “x,” then the participant had a higher probability of being cued to type the adjacent letter “p” than the distant letters “l” or “a”). The

question was, in the event a distant letter was cued, how long would it take the participants to type that letter? If well-learned action sequences like typing are governed by a chain of sequentially firing neurons or neural ensembles, then participants should be fast to type adjacent letters (i.e. the next action in the chain) but slow to type distant letters (i.e. actions further down the chain). Crucially, the latency to type a distant letter should be equally slow regardless of how distant it is. While participants were indeed faster to type adjacent letters, they also showed a graded increase in the latency to type progressively distant actions. This result strongly suggests that actions within a well-learned sequence are not serially elicited in a chain-like fashion, but are activated in parallel.

A model of action sequence execution that incorporates parallel activation of actions could inform models of basal ganglia control of action sequencing. For example, it has been proposed that actions in a sequence may be primed simultaneously by a top-down controller, with each action receiving inhibition in proportion to its ordinal position in the sequence and the correct ordering of actions controlled by a gradual ramping down of inhibition (Rumelhart & Norman, 1982). Within the brain, this “competitive queuing” process could possibly be implemented by a cortical ensemble exciting a set of striatal MSN ensembles and each ensemble receiving inhibition from interneurons or lateral connections with other MSNs. Two predictions that come out of this model are that a striatal MSN should show a ramp of excitation leading up to the performance of its corresponding action, and removing sources of inhibition—either by inactivating inhibitory interneurons or inducing broad excitation throughout the striatum—should interfere with the correct ordering of actions.

4.3 Termination

One idea is that in order for a sequence of actions to be terminated, the basal ganglia must produce a stop signal that inhibits motor activity (Roseberry & Kreitzer, 2017). This hypothesis is partially substantiated by the common observation that the striatum and its cortical afferents fire a burst of action potentials at the end of a sequence of actions (Barnes et al., 2005; Desrochers et al., 2015; Fuji & Graybiel, 2003; Fujimoto et al., 2011; Jin & Costa, 2010; Jin et al., 2014; Jog et al., 1999; Smith & Graybiel, 2013; Thorn et al., 2010). These studies cover a diverse set of behaviors ranging from maze running in rats, to lever pressing in mice, to saccades in monkeys. The traditional model of basal ganglia function predicts that this phasic activity should be exclusive to D2 MSNs, which inhibit movement through the indirect pathway. While studies have found that only a minority of D2 MSNs showed phasic spiking at the end of a sequence (Jin et al., 2014; Geddes et al., 2018), the activity of these neurons may convey a powerful ‘stop’ signal that is transmitted via the subthalamic nucleus to the SNr (Schmidt et al., 2013; Schmidt & Berke, 2017).

It is important to recognize that most of the studies that have documented phasic end activity have used appetitive conditioning paradigms in which the termination of an action sequence is signaled by food reward. Thus, the phasic firing that is observed in the striatum at the end of an action sequence may signal information about reward expectation rather than a motor-related stop signal. One way to test this hypothesis is to train animals to perform action sequences under different conditions in which the probability of reward is varied. If the phasic firing observed within the striatum at the end of the sequence conveys information about reward expectation, then one would expect the firing rate to monotonically scale with reward probability. However, if the phasic firing does indeed serve as a motor stop signal, then neural activity should be indifferent to reward probability. On the other hand, the expectation of reward

itself could serve as a stop signal to terminate a sequence and engage in consummatory behavior (e.g. food retrieval). This idea is explicit in RL algorithms, in which the delivery of a reward signals a new state in which certain actions (i.e. consummatory behaviors) are more predictive of reward receipt than others. The fact that consummatory behaviors quickly follow the termination of sequences is also a problem in itself. That is, it is unclear whether the phasic end activity in the striatum codes information about the end of a sequence or the initiation of food retrieval responses.

It could be argued that an alternative way to avoid the problems with appetitive conditioning paradigms is to study naturally occurring behaviors that do not require explicit reinforcement, such as self-grooming in rodents and spontaneous singing in birds. For example, rats display a stereotyped sequence of grooming actions that most often terminate with a bout of licks to the flank (Berridge et al., 1987). Lesions to the DLS interfere with the proper termination of grooming sequences, with rats frequently replacing the licking of the flank with another grooming action or simply terminating grooming prematurely (Cromwell & Berridge, 1996). In Bengalese finches, lesions of the striatum exacerbate the tendency to repeat the terminating syllable of a song motif (Kubikova et al., 2014). It would be especially useful to investigate the neural mechanisms that mediate sequence termination in the context of these naturally occurring behaviors, including the causal function of striatal end signals.

Rodent self-grooming studies have revealed that a potential mechanism for action sequence termination may originate from fast-spiking striatal GABAergic interneurons, which have been shown to provide an important source of MSN inhibition in transgenic mouse models of obsessive-compulsive disorder—a disorder that has been characterized as a problem of action sequence termination (Ahmari et al., 2013; Burguière et al., 2013; Hinds et al., 2012). Yet

another possible source of a sequence stop mechanism could come from arkypallidal cells in the GPe, which send GABAergic projections back to the striatum and exhibit increased firing rates in response to a ‘stop’ cue in a stop-signal task (Mallet et al., 2016). The potential problem, though, with using rodent self-grooming as a model for non-reinforced behavior is that grooming may be negatively reinforcing, serving to prevent filth or irritation. Other models of truly spontaneous and non-reinforced behaviors are needed. For example, it was observed that cessation of spontaneous wheel running in mice correlated with reduced dopamine transients in the dorsal striatum (Howe & Dombeck, 2016). The observed changes in striatal dopamine were generally correlated with the accelerations and decelerations of running behavior, but not with random deliveries of water rewards (at least for those axons originating in SNc). This study points to another potential neural mechanism for sequence termination that appears indifferent to reward delivery. In addition, hyperkinetic disorders involving the basal ganglia, which are characterized by a loss of synaptic depotentiation (Calabresi et al., 2016), could potentially serve as useful models of action sequence termination as these disorders are problems of stopping ongoing movement.

5. Conclusions

The study of action sequence learning and performance provides an excellent opportunity to test computational and neural models of basal ganglia function. Only recently has it been possible to study and manipulate different striatal cell types *in vivo*, and this avenue of research has proven to be fruitful in revealing the differential contributions of the direct and indirect pathways to different facets of action sequencing. However, technological advances need to be matched by better behavioral assays in which the study of action sequencing is more generally

extended to heterogeneous sequences of varying movements. While the study of repetitive actions on ratio schedules with rodents can serve as a useful starting point, it is not at all certain whether neurons in basal ganglia nuclei show similar patterns of activity during the learning and performance of heterogeneous sequences as they do during simple operant schedules.

It is also important to recognize that the exact role of the basal ganglia in action sequence learning is far from clear. While neurocomputational models identify the dorsal striatum as harboring representations of RL variables that are theoretically crucial for action sequence learning (e.g. Collins & Frank, 2014; Fee, 2014; Joel, Niv, & Ruppín, 2002), it is uncertain whether RL algorithms can be mapped onto dorsal striatal physiology. Models of the dorsal striatum as being a nexus of reinforcement learning have been mostly constrained to model-free control, in which corticostriatal weights are updated according to changes in the value of state-action pairs. This is problematic for two reasons. First, behavioral research has shown that action sequence learning transcends model-free computations as demonstrated by reward devaluation experiments, and an action's sensitivity to devaluation depends on its serial position in the sequence (Balleine et al., 2005; Dickinson et al., 1983; Garr & Delamater, 2019; Killcross & Coutureau, 2003; Thraillkill & Bouton, 2017). Second, the dorsal striatum—specifically, the DMS—is required for the expression of devaluation-sensitive actions in simple operant schedules (e.g. Gremel & Costa, 2013; Yin et al., 2005a, 2005b), which has led to the hypothesis that the DMS and its cortical afferents are part of a circuit that participates in model-based computations (Bornstein & Daw, 2011; Daw et al., 2005; Khamassi & Humphries, 2012). Given the hypothesized role of the DMS in model-based planning, along with the role of the hippocampus in tracking distance to goals (Patai et al., 2019; Spiers, Olafsdottir, & Lever, 2018), future research that focuses on the interaction between these structures in the context of goal-

directed action sequencing could provide insight into the neural mechanisms that underlie the “goal gradient” that governs many forms of sequential behavior (Hull, 1932).

A novel hypothesis for how the mammalian basal ganglia supports action sequence learning draws from interval timing studies, and proposes that encoding the intervals between actions and outcomes could be a critical computation of the dorsal striatum. A potentially useful paradigm for following up on this idea is to train action sequence based on previously learned temporal relationships. For example, one could initially train a set of actions separately, with each action separated from reward by a unique, but fixed, amount of time (Figure 3, top left). Following this training regime, one could present all action simultaneously and probe whether learning to perform actions in the correct order is facilitated by prior knowledge of the temporal distance between actions and outcomes. If the dorsal striatum encodes the temporal relationship between action and outcome necessary for learning an action sequence, then disrupting neural activity during action-outcome interval training should attenuate this learning advantage (Figure 3, right).

[Insert Figure 3 here]

A specific prediction that comes out of this framework is that the firing patterns of striatal neurons should change depending on the temporal distance of the action from its expected future outcome. This has been observed during a fixed interval timing task (Mello et al., 2015), but remains to be investigated during the performance of a fixed action sequence. Another prediction relates to the goal-directed nature of sequential actions. Based on findings showing that goal-directed control of sequential actions generally increases as a function of proximity to the outcome (Balleine et al., 2005; Garr & Delamater, 2019; Dickinson et al., 1983; Killcross & Coutureau, 2003; Thrailkill & Bouton, 2017), goal-directed control depends on the average

temporal distance between actions and outcomes (DeRusso et al., 2010), and the rodent DMS is required for the expression of goal-directed control of single actions (Gremel & Costa, 2013; Yin et al., 2005), it is possible that targeting neural excitation in the DMS to actions temporally distal in a sequence could bestow them with goal-directed control when, under normal conditions, it would otherwise be absent. These predictions can serve as a springboard for future investigations.

Beyond questions of simple acquisition and goal-directed control, the relation between basal ganglia function and hierarchical action sequence organization is largely a mystery, and deserves to be explored in more detail. This will require improvements in how hierarchical action sequencing is studied at the level of behavior. In the study of action chunking, researchers often rely on ‘slips of action’ (Dezfouli et al., 2014) to indicate that chunking has occurred—that is, when animals continue performing latter parts of an action sequence that no longer need to be performed (e.g. Matsumoto et al., 1999). Slips of action, however, could also be explained by a stimulus-response mechanism, where the stimulus is proprioceptive feedback from the previous action in the sequence. The problem with this explanation is that it precludes a hierarchical representation of the sequence, or, at the very least, makes it unnecessary. An action chunk is a unitized representation of a sequence, and the component parts that make up the sequence are thought to be determined prior to its initiation. A stimulus-response mechanism, however, implies that each action is determined only once the immediately preceding action is performed.

Stimulus-response mechanisms of action sequencing can be ruled out using sophisticated tasks in which successful performance does not rely on the association between consecutive responses (e.g. Chen et al., 1997). Alternatively, rather than relying on sophisticated and laborious behavioral procedures, neural recordings could be leveraged to reveal hierarchical

representations of action chunks. For example, there is some evidence that neurons in the primate motor cortex signal specific sequences of actions before they occur (Tanji, 2001), and there are neurons in the rodent striatum that fire around the time of an action only when it is performed as part of a particular sequence (Martiros et al., 2018). The optimal approach would be to combine sophisticated action sequence tasks, such as the two-step task developed for rodents (Miller et al., 2017), with *in vivo* recordings to construct a more accurate model of how cortico-basal ganglia loops support the learning and/or performance of action sequences, be them hierarchical or not. If the goal is to focus on hierarchical sequencing, two regions worth investigating further are the secondary motor cortex and SNc, both of which provide inputs to the dorsal striatum. Both of these regions appear crucial for the successful initiation, but not execution, of action sequences (Bailey & Mair, 2007; da Silva et al., 2018; Rothwell et al., 2015). Given that a hierarchical representation is thought to involve differential control of initiation and execution (Dezfouli & Balleine, 2013), these regions are ideal for studying how hierarchical representations are constructed over the course of learning and how the neural circuits controlling initiation and execution interact. It is possible that monitoring neural activity in these regions during sophisticated behavioral tasks could yield insights that inform better models of action sequencing.

Acknowledgements: I thank Andrew Delamater and Jon Horvitz for extensive discussion and comments on an earlier version of this paper.

References

- Ahmari, S. E., Spellman, T., Douglass, N. L., Kheirbek, M. A., Simpson, H. B., Deisseroth, K., ... Hen, R. (2013). Repeated Cortico-Striatal Stimulation Generates Persistent OCD-Like Behavior. *Science*, 340(6137), 1234–1239. <https://doi.org/10.1126/science.1234733>
- Akam, T., Rodrigues-Vaz, I., Zhang, X., Pereira, M., Oliveira, R., Dayan, P., & Costa, R. M. (2017). Single-Trial Inhibition of Anterior Cingulate Disrupts Model-based Reinforcement Learning in a Two-step Decision Task . Abstract : Introduction : *BioRxiv*, 1–49. <https://doi.org/10.1101/126292>
- Bagetta, V., Picconi, B., Marinucci, S., Sgobio, C., Pendolino, V., Ghiglieri, V., ... Calabresi, P. (2011). Dopamine-Dependent Long-Term Depression Is Expressed in Striatal Spiny Neurons of Both Direct and Indirect Pathways: Implications for Parkinson's Disease. *Journal of Neuroscience*, 31(35), 12513–12522. <https://doi.org/10.1523/JNEUROSCI.2236-11.2011>
- Bailey, K. R., & Mair, R. G. (2007). Effects of frontal cortex lesions on action sequence learning in the rat. *European Journal of Neuroscience*, 25(9), 2905–2915. <https://doi.org/10.1111/j.1460-9568.2007.05492.x>
- Balleine, B. W., Paredes-Olay, C., & Dickinson, A. (2005). Effects of outcome devaluation on the performance of a heterogeneous instrumental chain. *International Journal of Comparative Psychology*, 18(4), 257–272. Retrieved from <http://escholarship.org/uc/item/5pd9x995.pdf>
- Balleine, B., & Dickinson, A. (1991). Instrumental Performance Following Reinforcer Devaluation Depends upon Incentive Learning. *Quarterly Journal of Experimental Psychology*, 43B(3), 279–296.
- 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(7062), 1158–1161. <https://doi.org/10.1038/nature04053>
- Bayer, H. M., & Glimcher, P. W. (2005). Midbrain dopamine neurons encode a quantitative reward prediction error signal. *Neuron*, 47(1), 129–41. <https://doi.org/10.1016/j.neuron.2005.05.020>
- Behmer, L. P., & Crump, M. J. C. (2017). The Dynamic Range of Response Set Activation During Action Sequencing. *Journal of Experimental Psychology: Human Perception and Performance*, 43(3), 537–554. <https://doi.org/10.1037/xhp0000335>

- Berke, J. D. (2018). What does dopamine mean? *Nature Neuroscience*, 21(June), 1. <https://doi.org/10.1038/s41593-018-0152-y>
- Berridge, K. C., Fentress, J. C., & Parr, H. (1987). Natural syntax rules control action sequences of rats. *Behavioural Brain Research*, 23, 59–68.
- Bertran-Gonzalez, J., Herve, D., Girault, J.-A., & Valjent, E. (2010). What is the degree of segregation between striatonigral and striatopallidal projections? *Frontiers in Neuroanatomy*, 4, 1–9. <https://doi.org/10.3389/fnana.2010.00136>
- Bornstein, A. M., & Daw, N. D. (2011). Multiplicity of control in the basal ganglia: Computational roles of striatal subregions. *Current Opinion in Neurobiology*, 21(3), 374–380. <https://doi.org/10.1016/j.conb.2011.02.009>
- Botvinick, M. M. (2008). Hierarchical models of behavior and prefrontal function. *Trends in Cognitive Sciences*, 12(5), 201–208. <https://doi.org/10.1016/j.tics.2008.02.009>
- Botvinick, M. M., Niv, Y., & Barto, A. C. (2009). Hierarchically organized behavior and its neural foundations: A reinforcement learning perspective. *Cognition*, 113(3), 262–280. <https://doi.org/10.1016/j.cognition.2008.08.011>
- Botvinick, M. M. (2012). Hierarchical reinforcement learning and decision making. *Current Opinion in Neurobiology*, 22(6), 956–962. <https://doi.org/10.1016/j.conb.2012.05.008>
- Burguière, E., Monteiro, P., Feng, G., & Graybiel, A. M. (2013). Optogenetic Stimulation of Lateral Orbitofronto-Striatal Pathway Suppresses Compulsive Behaviors. *Science*, 340(2013), 1243–1246. <https://doi.org/10.1126/science.1232380>
- Calabresi, P., Picconi, B., Tozzi, A., & Di Filippo, M. (2007). Dopamine-mediated regulation of corticostriatal synaptic plasticity. *Trends in Neurosciences*, 30(5), 211–219. <https://doi.org/10.1016/j.tins.2007.03.001>
- Calabresi, P., Picconi, B., Tozzi, A., Ghiglieri, V., & Di Filippo, M. (2014). Direct and indirect pathways of basal ganglia: a critical reappraisal. *Nature Neuroscience*, 17(8), 1022–1030. <https://doi.org/10.1038/nn.3743>
- Calabresi, P., Pisani, A., Rothwell, J., Ghiglieri, V., Obeso, J. A., & Picconi, B. (2016). Hyperkinetic disorders and loss of synaptic downscaling. *Nature Neuroscience*, 19(7), 868–875. <https://doi.org/10.1038/nn.4306>
- Catania, A. C. (1971). Reinforcement schedules: The role of responses preceding the one that produces the reinforcer. *Journal of the Experimental Analysis of Behavior*, 15(3), 271–287.
- Cazorla, M., DeCarvalho, F. D., Chohan, M. O., Shegda, M., Chuhma, N., Rayport, S., Ahmari, S. E., Moore, H., & Kellendonk, C. (2014). Dopamine D2 receptors regulate the anatomical

- and functional balance of basal ganglia circuitry. *Neuron*, 81(1), 153–164.
<https://doi.org/10.1016/j.neuron.2013.10.041>
- Chang, C. Y., Esber, G. R., Marrero-Garcia, Y., Yau, H. J., Bonci, A., & Schoenbaum, G. (2016). Brief optogenetic inhibition of dopamine neurons mimics endogenous negative reward prediction errors. *Nature Neuroscience*, 19, 111–116.
<https://doi.org/10.1038/nn.4191>
- Chen, S., Swartz, K. B., & Terrace, H. S. (1997). Knowledge of the Ordinal Position of List Items in Rhesus Monkeys. *Psychological Science*, 8(2), 80–86.
<https://doi.org/10.1111/j.1467-9280.1997.tb00687.x>
- Collins, A. G. E., & Frank, M. J. (2014). Opponent actor learning (OpAL): Modeling interactive effects of striatal dopamine on reinforcement learning and choice incentive. *Psychological Review*, 121(3), 337–366. <https://doi.org/10.1037/a0037015>
- Collins, A. L., Greenfield, V. Y., Bye, J. K., Linker, K. E., Wang, A. S., & Wassum, K. M. (2016). Dynamic mesolimbic dopamine signaling during action sequence learning and expectation violation. *Scientific Reports*, 6, 20231. <https://doi.org/10.1038/srep20231>
- Corbit, L. H., & Balleine, B. W. (2003). Instrumental and Pavlovian incentive processes have dissociable effects on components of a heterogeneous instrumental chain. *Journal of Experimental Psychology: Animal Behavior Processes*, 29(2), 99–106.
<https://doi.org/10.1037/0097-7403.29.2.99>
- Corbit, L. H., Nie, H., & Janak, P. H. (2012). Habitual alcohol seeking: Time course and the contribution of subregions of the dorsal striatum. *Biological Psychiatry*, 72(5), 389–395.
<https://doi.org/10.1016/j.biopsych.2012.02.024>
- Cromwell, H. C., & Berridge, K. C. (1996). Implementation of action sequences by a neostriatal site: a lesion mapping study of grooming syntax. *The Journal of Neuroscience*, 16(10), 3444–3458.
- Cui, G., Jun, S. B., Jin, X., Pham, M. D., Vogel, S. S., Lovinger, D. M., & Costa, R. M. (2013). Concurrent activation of striatal direct and indirect pathways during action initiation. *Nature*, 494(7436), 238–242. <https://doi.org/10.1038/nature11846>
- Da Silva, J. A., Tecuapetla, F., Paixão, V., & Costa, R. M. (2018). Dopamine neuron activity before action initiation gates and invigorates future movements. *Nature*, 554(7691), 244–248. <https://doi.org/10.1038/nature25457>
- Daw, N. D., Niv, Y., & Dayan, P. (2005). Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nature Neuroscience*, 8(12), 1704–11.
<https://doi.org/10.1038/nn1560>

- Daw, N. D., Gershman, S. J., Seymour, B., Dayan, P., & Dolan, R. J. (2011). Model-based influences on humans' choices and striatal prediction errors. *Neuron*, 69(6), 1204–1215. <https://doi.org/10.1016/j.neuron.2011.02.027>
- De Corte, B. J., Wagner, L. M., Matell, M. S., & Narayanan, N. S. (2019). Striatal dopamine and the temporal control of behavior. *Behavioural Brain Research*, 356(September 2018), 375–379. <https://doi.org/10.1016/j.bbr.2018.08.030>
- Desrochers, T. M., Amemori, K., & Graybiel, A. M. (2015). Habit Learning by Naive Macaques Is Marked by Response Sharpening of Striatal Neurons Representing the Cost and Outcome of Acquired Action Sequences. *Neuron*, 87(4), 853–868. <https://doi.org/10.1016/j.neuron.2015.07.019>
- Dezfouli, A., & Balleine, B. W. (2013). Actions, Action Sequences and Habits: Evidence That Goal-Directed and Habitual Action Control Are Hierarchically Organized. *PLoS Computational Biology*, 9(12), e1003364. <https://doi.org/10.1371/journal.pcbi.1003364>
- Dezfouli, A., & Balleine, B. W. (2012). Habits, action sequences, and reinforcement learning. *European Journal of Neuroscience*, 35(7), 1036–1051. <https://doi.org/10.1111/j.1460-9568.2012.08050.x>
- Dezfouli, A., Lingawi, N. W., & Balleine, B. W. (2014). Habits as action sequences : hierarchical action control and changes in outcome value. *Phil. Trans. R. Soc. B*, 369(1655), 20130482.
- Díaz-Hernández, E., Contreras-López, R., Sánchez-Fuentes, A., Rodríguez-Sibrián, L., Ramírez-Jarquín, J. O., & Tecuapetla, F. (2018). The Thalamostriatal Projections Contribute to the Initiation and Execution of a Sequence of Movements. *Neuron*, 100(3), 739–752.e5. <https://doi.org/10.1016/j.neuron.2018.09.052>
- Dickinson, A. (1994). Instrumental conditioning. In N. Mackintosh (Ed.), *Animal cognition and learning* (pp. 45–78). London: Academic Press.
- Dickinson, A., & Charnock, D. J. (1985). Contingency effects with maintained instrumental reinforcement. *Quarterly Journal of Experimental Psychology Section B: Comparative and Physiological Psychology*, 37(4), 397–416. <https://doi.org/10.1080/14640748508401177>
- Dickinson, A., Nicholas, D. J., & Adams, C. D. (1983). The effect of the instrumental training contingency on susceptibility to reinforcer devaluation. *The Quarterly Journal of Experimental Psychology Section B*, 35(1), 35–51. <https://doi.org/10.1080/14640748308400912>
- Doll, B. B., Duncan, K. D., Simon, D. A., Shohamy, D., & Daw, N. D. (2015). Model-based choices involve prospective neural activity. *Nature Neuroscience*, 18(5), 767–772. <https://doi.org/10.1038/nn.3981>

- Drago, J., Padungchaichot, P., Wong, J. Y., Lawrence, a J., McManus, J. F., Sumarsono, S. H., ... Finkelstein, D. I. (1998). Targeted expression of a toxin gene to D1 dopamine receptor neurons by cre-mediated site-specific recombination. *The Journal of Neuroscience*, 18(23), 9845–57. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/9822743>
- Durieux, P. F., Bearzatto, B., Guiducci, S., Buch, T., Waisman, A., Zoli, M., ... de Kerchove d'Exaerde, A. (2009). D2R striatopallidal neurons inhibit both locomotor and drug reward processes. *Nature Neuroscience*, 12(4), 393–5. <https://doi.org/10.1038/nn.2286>
- English, D. F., Ibanez-Sandoval, O., Stark, E., Tecuapetla, F., Buzsáki, G., Deisseroth, K., ... Koos, T. (2012). GABAergic circuits mediate the reinforcement-related signals of striatal cholinergic interneurons. *Nature Neuroscience*, 15(1), 123–130. <https://doi.org/10.1038/nn.2984>
- Eshel, N., Bukwich, M., Rao, V., Hemmelder, V., Tian, J., & Uchida, N. (2015). Arithmetic and local circuitry underlying dopamine prediction errors. *Nature*, 525(7568), 243–246. <https://doi.org/10.1038/nature14855>
- Fee, M. S. (2014). The role of efference copy in striatal learning. *Current Opinion in Neurobiology*, 25, 194–200. <https://doi.org/10.1016/j.conb.2014.01.012>
- Frank, M. J., Seeberger, L. C., & O'Reilly, R. C. (2004). By carrot or by stick: cognitive reinforcement learning in parkinsonism. *Science*, 306(5703), 1940–1943. <https://doi.org/10.1126/science.1102941>
- Frank, M. J. (2005). Dynamic Dopamine Modulation in the Basal Ganglia: A Neurocomputational Account of Cognitive Deficits in Medicated and Nonmedicated Parkinsonism. *Journal of Cognitive Neuroscience*, 17(1), 51–72.
- Freeze, B. S., Kravitz, A. V., Hammack, N., Berke, J. D., & Kreitzer, A. C. (2013). Control of basal ganglia output by direct and indirect pathway projection neurons. *The Journal of Neuroscience*, 33(47), 18531–9. <https://doi.org/10.1523/JNEUROSCI.1278-13.2013>
- Friedel, E., Koch, S. P., Wendt, J., Heinz, A., Deserno, L., & Schlagenhauf, F. (2014). Devaluation and sequential decisions: linking goal-directed and model-based behavior. *Frontiers in Human Neuroscience*, 8(August), 587. <https://doi.org/10.3389/fnhum.2014.00587>
- Fu, W. T., & Anderson, J. R. (2006). From recurrent choice to skill learning: A reinforcement-learning model. *Journal of Experimental Psychology: General*, 135(2), 184–206. <https://doi.org/10.1037/0096-3445.135.2.184>
- Fujii, N., & Graybiel, A. M. (2003). Representation of action sequence boundaries by macaque prefrontal cortical neurons. *Science*, 301(5637), 1246–1249. <https://doi.org/10.1126/science.1086872>

- Fujimoto, H., Hasegawa, T., & Watanabe, D. (2011). Neural coding of syntactic structure in learned vocalizations in the songbird. *The Journal of Neuroscience*, 31(27), 10023–10033. <https://doi.org/10.1523/JNEUROSCI.1606-11.2011>
- Gagnon, D., Petryszyn, S., Sanchez, M. G., Bories, C., Beaulieu, J. M., De Koninck, Y., ... Parent, M. (2017). Striatal Neurons Expressing D1 and D2 Receptors are Morphologically Distinct and Differently Affected by Dopamine Denervation in Mice. *Scientific Reports*, 7, 9–17. <https://doi.org/10.1038/srep41432>
- Garr, E., & Delamater, A. R. (2019). Exploring the relationship between actions, habits, and automaticity in an action sequence task. *Learning & Memory*, 26(4), 128–133. <https://doi.org/10.1101/lm.048645.118.26>
- Geddes, C. E., Li, H., & Jin, X. (2018). Optogenetic Editing Reveals the Hierarchical Organization of Learned Action Sequences. *Cell*, 174(1), 32–43.e15. <https://doi.org/10.1016/j.cell.2018.06.012>
- Gerfen, C. R., & Surmeier, D. J. (2011). Modulation of striatal projection systems by dopamine. *Annual Review of Neuroscience*, 34, 441–66. <https://doi.org/10.1146/annurev-neuro-061010-113641>
- Gershman, S. J., Blei, D. M., & Niv, Y. (2010). Context, learning, and extinction. *Psychological Review*, 117(1), 197–209. <https://doi.org/10.1037/a0017808>
- Gillan, C. M., Robbins, T. W., Sahakian, B. J., van den Heuvel, O. A., & van Wingen, G. (2016). The role of habit in compulsivity. *European Neuropsychopharmacology*, 26(5), 828–840. <https://doi.org/10.1016/j.euroneuro.2015.12.033>
- Giordano, N., Iemolo, A., Mancini, M., Cacace, F., De Risi, M., Latagliata, E. C., ... De Leonibus, E. (2018). Motor learning and metaplasticity in striatal neurons: Relevance for Parkinson's disease. *Brain*, 141(2), 505–520. <https://doi.org/10.1093/brain/awx351>
- Gouvêa, T. S., Monteiro, T., Motiwala, A., Soares, S., Machens, C., & Paton, J. J. (2015). Striatal dynamics explain duration judgments. *ELife*, 4, 1–14. <https://doi.org/10.7554/eLife.11386>
- Graybiel, A. M. (1998). The basal ganglia and chunking of action repertoires. *Neurobiology of Learning and Memory*, 70(1–2), 119–136. <https://doi.org/10.1006/nlme.1998.3843>
- Gremel, C. M., & Costa, R. M. (2013). Orbitofrontal and striatal circuits dynamically encode the shift between goal-directed and habitual actions. *Nature Communications*, 4(May), 1–12. <https://doi.org/10.1038/ncomms3264>
- Grillner, S., & Robertson, B. (2015). The basal ganglia downstream control of brainstem motor centres—an evolutionarily conserved strategy. *Current Opinion in Neurobiology*, 33(February), 47–52. <https://doi.org/10.1016/j.conb.2015.01.019>

- Guo, Q., Wang, D., He, X., Feng, Q., Lin, R., Xu, F., ... Luo, M. (2015). Whole-brain mapping of inputs to projection neurons and cholinergic interneurons in the dorsal striatum. *PLoS ONE*, 10(4), 1–15. <https://doi.org/10.1371/journal.pone.0123381>
- 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(6902), 65–70. <https://doi.org/10.1038/nature01221>
- Hammond, L. J., & Weinberg, M. (1984). Signaling unearned reinforcers removes the suppression produced by a zero correlation in an operant paradigm. *Animal Learning & Behavior*, 12(4), 371–377. <http://dx.doi.org/10.3758/BF03199982>
- Hardwick, R. M., Forrence, A. D., Krakauer, J. W., & Haith, A. M. (2017). Skill Acquisition and Habit Formation as Distinct Effects of Practice. *BioRxiv*, 1–35. <https://doi.org/10.1101/201095>
- Hart, G., Bradfield, L. A., & Balleine, B. W. (2018). Prefrontal cortico-striatal disconnection blocks the acquisition of goal-directed action. *The Journal of Neuroscience*, 38(5), 2850–17. <https://doi.org/10.1523/JNEUROSCI.2850-17.2017>
- Hasz, B. M., & Redish, A. D. (2018). Corrigendum: Deliberation and Procedural Automation on a Two-Step Task for Rats. *Frontiers in Integrative Neuroscience*, 12(August), 1–19. <https://doi.org/10.3389/fnint.2018.00040>
- Hikosaka, O., Nakamura, K., & Nakahara, H. (2006). Basal Ganglia Orient Eyes to Reward. *Journal of Neurophysiology*, 95(2), 567–584. <https://doi.org/10.1152/jn.00458.2005>
- Hinds, A. L., Woody, E. Z., van Ameringen, M., Schmidt, L. a., & Szechtman, H. (2012). When too much is not enough: Obsessive-compulsive disorder as a pathology of stopping, rather than starting. *PLoS ONE*, 7(1). <https://doi.org/10.1371/journal.pone.0030586>
- Horvitz, J. C. (2009). Stimulus-response and response-outcome learning mechanisms in the striatum. *Behavioural Brain Research*, 199(1), 129–40. <https://doi.org/10.1016/j.bbr.2008.12.014>
- Howe, M. W., Tierney, P. L., Sandberg, S. G., Phillips, P. E. M., & Graybiel, A. M. (2013). Prolonged dopamine signalling in striatum signals proximity and value of distant rewards. *Nature*, 500(7464), 575–579. <https://doi.org/10.1038/nature12475>
- Howard, C. D., Li, H., Geddes, C. E., & Jin, X. (2017). Dynamic Nigrostriatal Dopamine Biases Action Selection. *Neuron*, 1–15. <https://doi.org/10.1016/j.neuron.2017.02.029>
- Howe, M. W., & Dombeck, D. A. (2016). Rapid signalling in distinct dopaminergic axons during locomotion and reward. *Nature*, 535, 505–510. <https://doi.org/10.1038/nature18942>

- Hull, C. L. (1932). The goal gradient hypothesis and maze learning. *Psychological Review*, 39, 25–43. doi:10.1037/h0072640
- Ilango, A., Kesner, A. J., Keller, K. L., Stuber, G. D., Bonci, A., & Ikemoto, S. (2014). Similar Roles of Substantia Nigra and Ventral Tegmental Dopamine Neurons in Reward and Aversion. *The Journal of Neuroscience*, 34(3), 817–822. <https://doi.org/10.1523/JNEUROSCI.1703-13.2014>
- Jin, X., & Costa, R. M. (2015). Shaping action sequences in basal ganglia circuits. *Current Opinion in Neurobiology*, 33, 188–196. <https://doi.org/10.1016/j.conb.2015.06.011>
- Jin, X., & Costa, R. M. (2010). Start/stop signals emerge in nigrostriatal circuits during sequence learning. *Nature*, 466(7305), 457–462. <https://doi.org/10.1038/nature09263>
- Jin, X., Tecuapetla, F., & Costa, R. M. (2014). Basal ganglia subcircuits distinctively encode the parsing and concatenation of action sequences. *Nature Neuroscience*, 17(3), 423–30. <https://doi.org/10.1038/nn.3632>
- Jog, M. S., Kubota, Y., Connolly, C. I., Hillegaart, V., & Graybiel, A. M. (1999). Building neural representations of habits. *Science*, 286(1999), 1745–1749. <https://doi.org/10.1126/science.286.5445.1745>
- Keiflin, R., Pribut, H. J., Shah, N. B., & Janak, P. H. (2019). Ventral Tegmental Dopamine Neurons Participate in Reward Identity Predictions. *Current Biology*, 1–11. <https://doi.org/10.1016/j.cub.2018.11.050>
- Keiflin, R., Pribut, H. J., & Janak, P. H. (2019, November). Dissociable Contributions of Ventral Tegmental and Substantia Nigra Dopamine Neurons to Reinforcement Learning. Poster presented at Society for Neuroscience, San Diego, CA.
- Khamassi, M., & Humphries, M. D. (2012). Integrating cortico-limbic-basal ganglia architectures for learning model-based and model-free navigation strategies. *Frontiers in Behavioral Neuroscience*, 6(November), 1–19. <https://doi.org/10.3389/fnbeh.2012.00079>
- Killcross, S., & Coutureau, E. (2003). Coordination of actions and habits in the medial prefrontal cortex of rats. *Cerebral Cortex*, 13(4), 400–408. <https://doi.org/10.1093/cercor/13.4.400>
- Killeen, P. R. (1994). Mathematical principles of reinforcement. *Behavioral and Brain Sciences*, 17(1), 105–135. <https://doi.org/10.1017/S0140525X00033628>
- Kim, J., Kim, Y., Nakajima, R., Shin, A., Jeong, M., Park, A. H., ... Kim, D. (2017). Inhibitory Basal Ganglia Inputs Induce Excitatory Motor Signals in the Thalamus. *Neuron*, 95(5), 1181–1196.e8. <https://doi.org/10.1016/j.neuron.2017.08.028>

- Kish, S. J., Shannak, K., & Hornykiewicz, O. (1988). Uneven Pattern of Dopamine Loss in the Striatum of Patients with Idiopathic Parkinson's Disease. *The New England Journal of Medicine*, 318(14), 876–880.
- Klaus, A., Martins, G. J., Paixao, V. B., Zhou, P., Paninski, L., & Costa, R. M. (2017). The Spatiotemporal Organization of the Striatum Encodes Action Space. *Neuron*, 95(5), 1171–1180.e7. <https://doi.org/10.1016/j.neuron.2017.08.015>
- Kravitz, A. V., Tye, L. D., & Kreitzer, A. C. (2012). Distinct roles for direct and indirect pathway striatal neurons in reinforcement. *Nature Neuroscience*, 15(6), 816–818. <https://doi.org/10.1038/nn.3100>
- Kravitz, A. V., Freeze, B. S., Parker, P. R. L., Kay, K., Thwin, M. T., Deisseroth, K., & Kreitzer, A. C. (2010). Regulation of parkinsonian motor behaviours by optogenetic control of basal ganglia circuitry. *Nature*, 466(7306), 622–626. <https://doi.org/10.1038/nature09159>
- Kreitzer, A. C., & Malenka, R. C. (2007). Endocannabinoid-mediated rescue of striatal LTD and motor deficits in Parkinson's disease models. *Nature*, 445(7128), 643–647. <https://doi.org/10.1038/nature05506>
- Kubikova, L., Bosikova, E., Cvikova, M., Lukacova, K., Scharff, C., & Jarvis, E. D. (2014). Basal ganglia function, stuttering, sequencing, and repair in adult songbirds. *Scientific Reports*, 4, 6590. <https://doi.org/10.1038/srep06590>
- Laliv, A. L., Lien, A. D., Roseberry, T. K., Donahue, C. H., & Kreitzer, A. C. (2018). Motor thalamus supports striatum-driven reinforcement. *ELife*. <https://doi.org/10.7554/eLife.34032.001>
- Lashley, K. S. (1951). The Problem of Serial Order in Behavior. In *Cerebral Mechanisms in Behavior: The Hixon Symposium* (pp. 112–146). <https://doi.org/10.1016/j.humov.2007.04.001>
- Lerner, T. N., & Kreitzer, A. C. (2012). RGS4 Is Required for Dopaminergic Control of Striatal LTD and Susceptibility to Parkinsonian Motor Deficits. *Neuron*, 73(2), 347–359. <https://doi.org/10.1016/j.neuron.2011.11.015>
- Lerner, T. N., Shilyansky, C., Davidson, T. J., Evans, K. E., Beier, K. T., Zalocusky, K. A., ... Deisseroth, K. (2015). Intact-Brain Analyses Reveal Distinct Information Carried by SNc Dopamine Subcircuits. *Cell*, 162(3), 635–647. <https://doi.org/10.1016/j.cell.2015.07.014>
- Liang, L., DeLong, M. R., & Papa, S. M. (2008). Inversion of Dopamine Responses in Striatal Medium Spiny Neurons and Involuntary Movements. *The Journal of Neuroscience*, 28(30), 7537–7547. <https://doi.org/10.1523/JNEUROSCI.1176-08.2008>
- Long, M. A., Jin, D. Z., & Fee, M. S. (2010). Support for a synaptic chain model of neuronal sequence generation. *Nature*, 468(7322), 394–399. <https://doi.org/10.1038/nature09514>

- Marsden, C. D. (1980). The enigma of the basal ganglia and movement. *Trends in Neurosciences*, 3(11 C), 284–287. [https://doi.org/10.1016/0166-2236\(80\)90103-4](https://doi.org/10.1016/0166-2236(80)90103-4)
- Martiros, N., Burgess, A. A., & Graybiel, A. M. (2018). Inversely Active Striatal Projection Neurons and Interneurons Selectively Delimit Useful Behavioral Sequences. *Current Biology*, 0(0), 560–573. <https://doi.org/10.1016/j.cub.2018.01.031>
- Matamales, M., Skrbis, Z., Bailey, M. R., Balsam, P. D., Balleine, B. W., Götz, J., & Bertran-Gonzalez, J. (2017). A corticostriatal deficit promotes temporal distortion of automatic action in ageing. *ELife*, 6, 1–24. <https://doi.org/10.7554/eLife.29908>
- Matsumoto, N., Hanakawa, T., Maki, S., Graybiel, A. M., & Kimura, M. (1999). Role of Nigrostriatal Dopamine System in Learning to Perform Sequential Motor Tasks in a Predictive Manner. *Journal of Neurophysiology*, 82(2), 978–998.
- Mattar, M. G., & Daw, N. D. (2018). Prioritized memory access explains planning and hippocampal replay. *Nature Neuroscience*, 21(11), 1609–1617. <https://doi.org/10.1038/s41593-018-0232-z>
- 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. *Behavioral Neuroscience*, 117(4), 760–773. <https://doi.org/10.1037/0735-7044.117.4.760>
- Meck, W. H. (2006). Neuroanatomical localization of an internal clock: A functional link between mesolimbic, nigrostriatal, and mesocortical dopaminergic systems. *Brain Research*, 1109(1), 93–107. <https://doi.org/10.1016/j.brainres.2006.06.031>
- Mello, G. B. M., Soares, S., & Paton, J. J. (2015). A scalable population code for time in the striatum. *Current Biology*, 25(9), 1113–1122. <https://doi.org/10.1016/j.cub.2015.02.036>
- Miller, K. J., Botvinick, M. M., & Brody, C. D. (2017). Dorsal hippocampus contributes to model-based planning. *Nature Neuroscience*, 20(9), 1269–1276. <https://doi.org/10.1038/nn.4613>
- Miyachi, S., Hikosaka, O., Miyashita, K., Kárádi, Z., & Rand, M. K. (1997). Differential roles of monkey striatum in learning of sequential hand movement. *Experimental Brain Research*, 115(1), 1–5. <https://doi.org/10.1007/PL00005669>
- Morgan, M. J. (1974). Resistance to satiation. *Animal Behaviour*, 22, 449–466.
- Morris, R. W., Dezfouli, A., Griffiths, K. R., Le Pelley, M. E., & Balleine, B. W. (2017). The algorithmic neuroanatomy of action-outcome learning. *BioRxiv*, 137851. <https://doi.org/10.1101/137851>
- Nadjar, A., Brotchie, J. M., Guigoni, C., Li, Q., Zhou, S., Wang, G., ... Bezard, E. (2006). Phenotype of Striatofugal Medium Spiny Neurons in Parkinsonian and Dyskinetic

- Nonhuman Primates: A Call for a Reappraisal of the Functional Organization of the Basal Ganglia. *Journal of Neuroscience*, 26(34), 8653–8661. <https://doi.org/10.1523/JNEUROSCI.2582-06.2006>
- Nelson, A. B., & Kreitzer, A. C. (2014). Reassessing Models of Basal Ganglia Function and Dysfunction. *Annual Review of Neuroscience*, 37, 117–35. <https://doi.org/10.1146/annurev-neuro-071013-013916>
- Niv, Y., Daw, N. D., Joel, D., & Dayan, P. (2007). Tonic dopamine: Opportunity costs and the control of response vigor. *Psychopharmacology*, 191(3), 507–520. <https://doi.org/10.1007/s00213-006-0502-4>
- 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(3), 472–479. <https://doi.org/10.1016/j.neuron.2015.12.032>
- O'Hare, J. K., Li, H., Kim, N., Gaidis, E., Ade, K., Beck, J., ... Calakos, N. (2017). Striatal fast-spiking interneurons selectively modulate circuit output and are required for habitual behavior. *eLife*, 6, 1–26. <https://doi.org/10.7554/eLife.26231>
- Okubo, T. S., Mackevicius, E. L., Payne, H. L., Lynch, G. F., & Fee, M. S. (2015). Growth and splitting of neural sequences in songbird vocal development. *Nature*, 528(7582), 352–357. <https://doi.org/10.1038/nature15741>
- Oldenburg, I. A., & Sabatini, B. L. (2015). Antagonistic but Not Symmetric Regulation of Primary Motor Cortex by Basal Ganglia Direct and Indirect Pathways. *Neuron*, 86(5), 1174–1181. <https://doi.org/10.1016/j.neuron.2015.05.008>
- Osakada, F., Kobak, D., Machens, C. K., Watabe-Uchida, M., Tian, J., Uchida, N., ... Cohen, J. Y. (2016). Distributed and Mixed Information in Monosynaptic Inputs to Dopamine Neurons. *Neuron*, 91(6), 1374–1389. <https://doi.org/10.1016/j.neuron.2016.08.018>
- Ostlund, S. B., Winterbauer, N. E., & Balleine, B. W. (2009). Evidence of action sequence chunking in goal-directed instrumental conditioning and its dependence on the dorsomedial prefrontal cortex. *The Journal of Neuroscience*, 29(25), 8280–8287. <https://doi.org/10.1523/JNEUROSCI.1176-09.2009>
- Pan, W., Schmidt, R., Wickens, J. R., & Hyland, B. I. (2005). Dopamine Cells Respond to Predicted Events during Classical Conditioning: Evidence for Eligibility Traces in the Reward-Learning Network. *The Journal of Neuroscience*, 25(26), 6235–6242. <https://doi.org/10.1523/JNEUROSCI.1478-05.2005>
- Panigrahi, B., Martin, K. A., Li, Y., Graves, A. R., Vollmer, A., Olson, L., ... Dudman, J. T. (2015). Dopamine Is Required for the Neural Representation and Control of Movement Vigor. *Cell*, 162(6), 1418–1430. <https://doi.org/10.1016/j.cell.2015.08.014>

- Parker, N. F., Cameron, C. M., Taliaferro, J. P., Lee, J., Choi, J. Y., Davidson, T. J., ... Witten, I. B. (2016). Reward and choice encoding in terminals of midbrain dopamine neurons depends on striatal target. *Nature Neuroscience*, 19(April). <https://doi.org/10.1038/nn.4287>
- Pawlak, V., & Kerr, J. N. D. (2008). Dopamine receptor activation is required for corticostriatal spike-timing-dependent plasticity. *Journal of Neuroscience*, 28(10), 2435–46. <https://doi.org/10.1523/JNEUROSCI.4402-07.2008>
- Paxinos, G., & Watson, C. (2007). The rat brain in stereotaxic coordinates. New York: Academic.
- Perez, O., Aitken, M., Zhukovsky, P., Soto, F. A., Urcelay, G. P., & Dickinson, A. (2016). Human instrumental performance in ratio and interval contingencies: a challenge for associative theory. *The Quarterly Journal of Experimental Psychology*, 0(0), 1–13. <https://doi.org/10.1080/17470218.2016.1265996>
- Redish, a D., Jensen, S., Johnson, A., & Kurth-Nelson, Z. (2007). Reconciling reinforcement learning models with behavioral extinction and renewal: implications for addiction, relapse, and problem gambling. *Psychological Review*, 114(3), 784–805. <https://doi.org/10.1037/0033-295X.114.3.784>
- Rescorla, R. A., & Wagner, A. . (1972). A Theory of Pavlovian Conditioning: Variations in the Effectiveness of Reinforcement and Nonreinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical Conditioning II: Current Theory and Research* (pp. 64–99). New York: Appleton-Century-Crofts.
- Reynolds, J. N., Hyland, B. I., & Wickens, J. R. (2001). A cellular mechanism of reward-related learning. *Nature*, 413(6851), 67–70. <https://doi.org/10.1038/35092560>
- Roseberry, T. K., Bonci, A., Kreitzer, A. C., Lalive, A. L., Lee, A. M., & Wilbrecht, L. (2016). Cell-Type-Specific Control of Brainstem Locomotor Circuits by Basal Ganglia. *Cell*, 164(3), 526–537. <https://doi.org/10.1016/j.cell.2015.12.037>
- Roseberry, T., & Kreitzer, A. (2017). Neural circuitry for behavioural arrest. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 372(1718), 9–12. <https://doi.org/10.1098/rstb.2016.0197>
- Rosenbaum, D. A., Cohen, R. G., Jax, S. A., Weiss, D. J., & van der Wel, R. (2007). The problem of serial order in behavior: Lashley's legacy. *Human Movement Science*, 26(4), 525–554. <https://doi.org/10.1016/j.humov.2007.04.001>
- Rothwell, P. E., Hayton, S. J., Sun, G. L., Fuccillo, M. V., Lim, B. K., & Malenka, R. C. (2015). Input- and Output-Specific Regulation of Serial Order Performance by Corticostriatal Circuits. *Neuron*, 88(2), 345–356. <https://doi.org/10.1016/j.neuron.2015.09.035>

- Rumelhart, D. E., & Norman, D. A. (1982). Simulating a Skilled Typist : A Study of Skilled Cognitive-Motor Performance. *Cognitive Science*, 6, 1–36.
- 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*, 4, 1–19. <https://doi.org/10.7554/eLife.09423>
- Saunders, B. T., Richard, J. M., Margolis, E. B., & Janak, P. H. (2018). Dopamine neurons create Pavlovian conditioned stimuli with circuit-defined motivational properties. *Nature Neuroscience*, 21(8), 1072–1083. <https://doi.org/10.1038/s41593-018-0191-4>
- Schmidt, R., & Berke, J. D. (2017). A Pause-then-Cancel model of stopping: evidence from basal ganglia neurophysiology. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 372, 20160202. <https://doi.org/10.1098/rstb.2016.0202>
- Schmidt, R., Leventhal, D. K., Mallet, N., Chen, F., & Berke, J. D. (2013). Canceling actions involves a race between basal ganglia pathways. *Nature Neuroscience*, 16(8), 1118–24. <https://doi.org/10.1038/nn.3456>
- Schultz, W., Dayan, P., & Montague, P. R. (1997). A Neural Substrate of Prediction and Reward. *Science*, 275(5306), 1593–1599. <https://doi.org/10.1126/science.275.5306.1593>
- Schultz, W. (1998). Predictive Reward Signal of Dopamine Neurons. *Journal of Neurophysiology*, 80(1), 1–27. <https://doi.org/10.1152/jn.1998.80.1.1>
- Schulz, J. M., & Reynolds, J. N. J. (2013). Pause and rebound: sensory control of cholinergic signaling in the striatum. *Trends in Neurosciences*, 36(1), 41–50. <https://doi.org/10.1016/j.tins.2012.09.006>
- Sharpe, M. J., Chang, C. Y., Liu, M. A., Batchelor, H. M., Mueller, L. E., Jones, J. L., ... Schoenbaum, G. (2017). Dopamine transients are sufficient and necessary for acquisition of model-based associations. *Nature Neuroscience*, 2017(April). <https://doi.org/10.1038/nn.4538>
- Sharpe, M. J., Stalnaker, T., Schuck, N. W., Killcross, S., Schoenbaum, G., & Niv, Y. (2019). An Integrated Model of Action Selection : Distinct Modes of Cortical Control of Striatal Decision Making. *Annual Review of Psychology*, 1–24. <https://doi.org/10.1146/annurev-psych-010418-102824>
- Shen, W., Flajolet, M., Greengard, P., & Surmeier, D. J. (2008). Dichotomous dopaminergic control of striatal synaptic plasticity. *Science*, 321(5890), 848–51. <https://doi.org/10.1126/science.1160575>
- Shi, S., Li, Y., & McNally, G. P. (2014). The conditions that promote fear learning: Prediction error and Pavlovian fear conditioning. *Neurobiology of Learning and Memory*, 108, 14–21. <https://doi.org/10.1016/j.nlm.2013.05.002>

- Soares, S., Atallah, B. V., & Paton, J. J. (2016). Midbrain dopamine neurons control judgment of time. *Science*, 354(6317), 1273–1277. <https://doi.org/10.1126/science.aah5234>
- Smith, K. S., Bucci, D. J., Luikart, B. W., & Mahler, S. V. (2016). DREADDs: Use and Application in Behavioral Neuroscience. *Behavioral Neuroscience*, 130(2), 137–155. <https://doi.org/10.1037/bne0000135>
- Smith, K. S., & Graybiel, A. M. (2013). A dual operator view of habitual behavior reflecting cortical and striatal dynamics. *Neuron*, 79, 361–374. <https://doi.org/10.1016/j.neuron.2013.05.038>
- Spence, K. W. (1932). The order of eliminating blinds in maze learning by the rat. *Journal of Comparative Psychology*, 14, 9–27. <https://doi.org/10.1037/h0075997>
- St. Claire-smith, R. (1979). The overshadowing of instrumental conditioning by a stimulus that predicts reinforcement better than the response. *Animal Learning & Behavior*, 7(2), 224–228.
- Stachenfeld, K. L., Botvinick, M. M., & Gershman, S. J. (2017). The hippocampus as a predictive map. *Nature Neuroscience*, 20(11), 1643–1653. <https://doi.org/10.1038/nn.4650>
- Starkweather, C. K., Babayan, B. M., Uchida, N., & Gershman, S. J. (2017). Dopamine reward prediction errors reflect hidden-state inference across time. *Nature Neuroscience*, 2(March), 1–11. <https://doi.org/10.1038/nn.4520>
- Steinberg, E. E., Keiflin, R., Boivin, J. R., Witten, I. B., Deisseroth, K., & Janak, P. H. (2013). A causal link between prediction errors, dopamine neurons and learning. *Nature Neuroscience*, 16(7), 966–973. <https://doi.org/10.1038/nn.3413>
- Suri, R. E., & Schultz, W. (1998). Learning of sequential movements by neural network model with dopamine-like reinforcement signal. *Experimental Brain Research*, 121(3), 350–354. <https://doi.org/10.1007/s002210050467>
- Sutton, R. S., & Barto, A. G. (2018). *Reinforcement Learning: An Introduction* (2nd ed.). Cambridge, MA: MIT Press.
- Sutton, R. S., Precup, D., & Singh, S. (1999). Between MDPs and semi-MDPs: A framework for temporal abstraction in reinforcement learning. *Artificial Intelligence*, 112(1), 181–211. [https://doi.org/10.1016/S0004-3702\(99\)00052-1](https://doi.org/10.1016/S0004-3702(99)00052-1)
- 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(9), 1281–1289. <https://doi.org/10.1038/nn.3188>
- Takahashi, Y. K., Roesch, M. R., Stalnaker, T. a, Haney, R. Z., Calu, D. J., Taylor, A. R., ... Schoenbaum, G. (2009). The orbitofrontal cortex and ventral tegmental area are necessary

- for learning from unexpected outcomes. *Neuron*, 62(2), 269–80. <https://doi.org/10.1016/j.neuron.2009.03.005>
- Takahashi, Y. K., Batchelor, H. M., Liu, B., Khanna, A., Morales, M., & Schoenbaum, G. (2017). Dopamine Neurons Respond to Errors in the Prediction of Sensory Features of Expected Rewards. *Neuron*, 95(6), 1395–1405.e3. <https://doi.org/10.1016/j.neuron.2017.08.025>
- Tanaka, M., Singh, J., Murugan, M., & Mooney, R. (2016). Focal expression of mutant huntingtin in the songbird basal ganglia disrupts cortico-basal ganglia networks and vocal sequences. *Proceedings of the National Academy of Sciences, March 7*, E1720–E1727. <https://doi.org/10.1073/pnas.1523754113>
- Tanji, J. (2001). Sequential Organization of Multiple Movements: Involvement of Cortical Motor Areas. *Annual Review of Neuroscience*, 24, 631–651. <https://doi.org/10.1147/006X010621-063>
- Tecuapetla, F., Jin, X., Lima, S. Q., & Costa, R. M. (2016). Complementary Contributions of Striatal Projection Pathways to Action Initiation and Execution. *Cell*, 166, 1–13. <https://doi.org/10.1016/j.cell.2016.06.032>
- Terrace, H. S. (2005). The simultaneous chain: A new approach to serial learning. *Trends in Cognitive Sciences*, 9(4), 202–210. <https://doi.org/10.1016/j.tics.2005.02.003>
- 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(5), 781–795. <https://doi.org/10.1016/j.neuron.2010.04.036>
- Thorndike, E. L. (1898). Animal intelligence: An experimental study of the associative processes in animals. *The Psychological Review: Monograph Supplements*, 2(4), i-109.
- Thrailkill, E. A., & Bouton, M. E. (2017). Effects of outcome devaluation on instrumental behaviors in a discriminated heterogeneous chain. *Journal of Experimental Psychology: Animal Learning and Cognition*, 43(1), 88–95. <https://doi.org/10.1037/xan0000119>
- Tolman, E. C., & Honzik, C. H. (1930). Degrees of hunger, reward and non-reward, and maze learning in rats. *University of California Publications in Psychology*, 4, 241–256.
- Tritsch, N. X., & Sabatini, B. L. (2012). Dopaminergic Modulation of Synaptic Transmission in Cortex and Striatum. *Neuron*, 76(1), 33–50. <https://doi.org/10.1016/j.neuron.2012.09.023>
- Vicente, A., Galvão-Ferreira, P., Tecuapetla, F., & Costa, R. (2016). Direct and indirect striatal projection pathways reinforce different action strategies. *Current Biology*, 26(7), R267–R269. <https://doi.org/10.1016/j.cub.2016.02.036>

- Von Holst, E. (1973). On The Nature of Order in the Central Nervous System (R.D. Martin, Trans.). *Naturwissenschaften*, 25, 625—31 and 641—647. (Original work published 1937).
- Waelti, P., Dickinson, A., & Schultz, W. (2001). Dopamine responses comply with basic assumptions of formal learning theory. *Nature*, 412, 43–48.
<https://doi.org/10.1080/00028896809342979>
- Wagner, A. R. (1961). Effects of amount and percentage of reinforcement and number of acquisition trials on conditioning and extinction. *Journal of Experimental Psychology*, 62(3), 234–242. <http://dx.doi.org/10.1037/h0042251>
- Wall, N. R., DeLaParra, M., Callaway, E. M., & Kreitzer, A. C. (2013). Differential innervation of direct- and indirect-pathway striatal projection neurons. *Neuron*, 79(2), 347–360.
<https://doi.org/10.1016/j.neuron.2013.05.014>
- Walsh, M. M., & Anderson, J. R. (2011). Learning from delayed feedback: neural responses in temporal credit assignment. *Cognitive, Affective & Behavioral Neuroscience*, 11(2), 131–43. <https://doi.org/10.3758/s13415-011-0027-0>
- Wassum, K. M., Ostlund, S. B., & Maidment, N. T. (2012). Phasic mesolimbic dopamine signaling precedes and predicts performance of a self-initiated action sequence task. *Biological Psychiatry*, 71(10), 846–854. <https://doi.org/10.1016/j.biopsych.2011.12.019>
- Wickens, J. R., Begg, A. J., & Arbuthnott, G. W. (1996). Dopamine reverses the depression of rat corticostriatal synapses which normally follows high-frequency stimulation of cortex in vitro. *Neuroscience*, 70(1), 1–5. [https://doi.org/10.1016/0306-4522\(95\)00436-M](https://doi.org/10.1016/0306-4522(95)00436-M)
- Wickens, J. R., Budd, C. S., Hyland, B. I., & Arbuthnott, G. W. (2007). Striatal contributions to reward and decision making: Making sense of regional variations in a reiterated processing matrix. *Annals of the New York Academy of Sciences*, 1104, 192–212.
<https://doi.org/10.1196/annals.1390.016>
- Wickens, J. R., Reynolds, J. N. J., & Hyland, B. I. (2003). Neural mechanisms of reward-related motor learning. *Current Opinion in Neurobiology*, 13(6), 685–690.
<https://doi.org/10.1016/j.conb.2003.10.013>
- Williams, B.A (1994). Conditioned reinforcement: Neglected or outmoded explanatory construct? *Psychonomic Bulletin & Review*, 1(4), 457–475.
- Williams, B. A. (1999). Associative competition in operant conditioning: Blocking the response-reinforcer association. *Psychonomic Bulletin & Review*, 6, 618–623.
- Wilson, R. C., Takahashi, Y. K., Schoenbaum, G., & Niv, Y. (2014). Orbitofrontal cortex as a cognitive map of task space. *Neuron*, 81(2), 267–79.
<https://doi.org/10.1016/j.neuron.2013.11.005>

- Witten, I. B., Steinberg, E. E., Lee, S. Y., Davidson, T. J., Zalocusky, K. A., Brodsky, M., ... Deisseroth, K. (2011). Recombinase-driver rat lines: Tools, techniques, and optogenetic application to dopamine-mediated reinforcement. *Neuron*, 72(5), 721–733. <https://doi.org/10.1016/j.neuron.2011.10.028>
- Yagishita, S., Hayashi-Takagi, A., Ellis-Davies, G. C. R., Urakubo, H., Ishii, S., & Kasai, H. (2014). A critical time window for dopamine actions on the structural plasticity of dendritic spines. *Science*, 345(6204), 1616–1620. <https://doi.org/10.1126/science.1255514>
- Yin, H. H. (2010). The Sensorimotor Striatum Is Necessary for Serial Order Learning. *The Journal of Neuroscience*, 30(44), 14719–14723. <https://doi.org/10.1523/JNEUROSCI.3989-10.2010>
- Yin, H. H., Knowlton, B. J., & Balleine, B. W. (2004). Lesions of dorsolateral striatum preserve outcome expectancy but disrupt habit formation in instrumental learning. *European Journal of Neuroscience*, 19, 181–189. <https://doi.org/10.1111/j.1460-9568.2004.03095.x>
- Yin, H. H., Knowlton, B. J., & Balleine, B. W. (2005a). Blockade of NMDA receptors in the dorsomedial striatum prevents action-outcome learning in instrumental conditioning. *European Journal of Neuroscience*, 22(2), 505–512. <https://doi.org/10.1111/j.1460-9568.2005.04219.x>
- Yin, H. H., Ostlund, S. B., Knowlton, B. J., & Balleine, B. W. (2005b). The role of the dorsomedial striatum in instrumental conditioning. *European Journal of Neuroscience*, 22(2), 513–523. <https://doi.org/10.1111/j.1460-9568.2005.04218.x>
- Yttri, E. A., & Dudman, J. T. (2016). Opponent and bidirectional control of movement velocity in the basal ganglia. *Nature*, 533(7603), 402–406. <https://doi.org/10.1038/nature17639>

Figure 1. (A) An idealized depiction of how the ordinal position of an action in a sequence determines the rate of acquisition (left) and extinction (right). Actions that are more proximal to a rewarding outcome are learned about sooner than more distal actions. (B) Left: A set of actions are executed to earn a rewarding outcome. Following training, the outcome is devalued outside of the sequence, after which sensitivity to outcome devaluation is probed. Right: An idealized depiction of how the ordinal position of an action in a sequence determines its sensitivity to

outcome devaluation. Actions more proximal to the outcome are more sensitive to outcome devaluation, and thus more goal-directed.

Figure 2. A sagittal view of the rodent basal ganglia. The direct and indirect pathways are represented by blue and red, respectively. Black connections are common to both pathways. The yellow arrow represents nigrostriatal dopamine. Connections terminating with a circle are excitatory projections, while connections terminating with a vertical line are inhibitory. GPe, globus pallidus external segment; EP, entopeduncular nucleus; StN, subthalamic nucleus; SNr, substantia nigra pars reticulata; SNc, substantia nigra pars compacta. Figure inspired by Nelson & Krietzner (2014) and adapted from Paxinos & Watson (2007).

Figure 3. Behavioral paradigm for studying action sequence learning based on interval timing.

Actions are separately trained such that each time an action is performed, an outcome is obtained after a fixed interval (top left). Each action is separated from the outcome by a unique interval. Then, subjects are required to learn to perform those actions in a specific order that either maintains the prior temporal relationships between action and outcome ('Maintained') or is independent of those relationships ('Random'). The prediction is that the Maintained group will learn the sequence at a faster rate compared to the Random group, and a group given no prior training will learn at a relatively intermediate rate (right). Lesions or inactivations of the dorsal striatum (bottom left) are predicted to erase the advantage conferred by maintained temporal relationships between actions and outcome (right, red arrows).