The University Of Sheffield.

The Impact of Striatal Neuropeptides and Topography on Action Sequence Selection

The University Of Sheffield

Department of Psychology

Adaptive Behaviour Research Group

Author:
David Buxton

Supervisors:
Prof. Kevin Gurney
Prof. Paul Overton
Dr. Enrico Bracci

A thesis submitted in partial fulfilment of the requirements for the degree of Doctor of Philosophy

October 2018

To my mother Karen,
whom I miss, and love

Acknowledgements

I would like to extend my deep and heartfelt gratitude to my supervisors, whose advice and encouragement has helped me avoid innumerable pitfalls and dead-ends over the past four years. My thanks to Professor Kevin Gurney for always being available to rescue me from being lost amongst the weeds, and for many insightful discussions on the way back. My thanks also to Professor Paul Overton and Dr. Enrico Bracci, whose knowledge and guidance has kept me on course and prevented several modelling flights of fancy.

I am indebted to my parents, who have always supported my endeavours in and out of university, at home and abroad, through good times and bad; and to my grandparents, for their wise life lessons and for making the pursuit of my studies possible.

I am grateful for the many friends who have comforted me with companionship, refreshed me with drinks, or inspired me with their own struggles and academic journeys. My thanks especially to Alex, Carola, James, Jon, and Nicola for their frequent willingness to let me bend their ear about some trifle or other.

Finally, I am deeply thankful for my dove Michelle, who has shared my stress and despair, and suffered through many practice presentations, yet has inexplicably endured and responded with mugs of tea, patience, and love ♡

Abstract

Many common behaviours are a sequence of several actions. As action sequences are learned their activation often becomes habitual, allowing smooth, rapid, and semi-automatic execution; learning and performing action sequences is central to normal motor function.

The striatum is the primary input nucleus for the basal ganglia and receives glutamatergic cortical afferents. These afferents innervate localised populations of medium spiny neurons (MSNs) and may encode 'action requests'. Striatal interactions ensure that only non-conflicting, high salience requests are selected, but the mechanisms enabling clean, rapid switching between sequential actions are poorly understood.

Substance P (SP) and enkephalin are neuropeptides co-released with GABA by MSNs preferentially expressing D_1 or D_2 dopamine receptors respectively. SP facilitates subsequent glutamatergic inputs to target MSNs while enkephalin has an inhibitory effect. We construct models of these glutamatergic effects and integrate them into a basal ganglia model to demonstrate that diffuse neuropeptide connectivity enhances action selection. For action sequences with an ordinal structure, patterning SP connectivity to reflect this ordering enhances the selection of correctly–ordered actions and suppresses disordered selection. We also show that selectively pruning SP connections allows context–sensitive inhibition of specific undesirable requests that otherwise interfere with action group selection.

We then construct a striatal microcircuit model with physical topography and show that inputs to this model generate oscillations in MSN spiking. Input salience and active neuronal density have differentiable impacts on oscillation amplitude and frequency, but the presence of oscillations has little effect on the mean MSN firing rate or action selection.

Our model suggests that neuropeptide interactions enhance the contrast between selected and rejected action requests, and that patterned SP connectivity enhances the selection of ordered sequences. Our model further suggests that striatal topography does not directly impact action selection, but that evoked oscillations may represent an additional form of population coding that could bind together semantically related MSN groups.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

The mammalian brain is a fascinating and complex machine capable of solving a wide range of problems, from basic tasks such as finding food and shelter to more evolutionarily recent issues such as navigating social groups. Some problems have forced the development of new brain structures and cognitive systems to enable fundamentally new skills, but many have been solved by adapting and expanding existing abilities for new situations. And in all cases, an animal's behaviour depends on the structure and connectivity of its neural systems, including neurons, glia, neurotransmitters, and other messenger chemicals.

High-level behaviours are often too complex for us to understand their entire neurophysiology in detail, but exploration of the underlying processes can provide significant insight into salient features. By examining the operation of a specific function in detail we may be able to better understand how it has been adapted for more complex tasks, and one function in particular is of great importance for the behaviour of all animals.

1.1 Actions and Sequences

1.1.1 Action Selection

Action selection is arguably the most fundamental problem for animal life: what to do next? How should the body's limited resources be utilised when presented with nearly limitless behavioural possibilities in a complex and changing environment? How can competing or contradictory goals be achieved with speed and accuracy? A working definition of key terms will prove useful.

Definition of Terms

Searle (1979) defines an *intentional action* as "an intention in action together with the bodily movement... which is caused by it." This broad definition encompasses everything

from tapping a finger to playing a piano concerto, so for our purposes an 'action' refers to what Searle calls a *basic action*: any intentional action that can be completed without recourse to other intentional actions. A finger tap cannot be sensibly broken down into smaller steps and is thus a basic action, while playing a concerto requires the co-ordination of many complex intermediate actions and therefore is not.

Importantly, this definition allows different agents to possess different sets of basic actions according to their individual abilities, and implies that it may be possible to increase the set of basic actions available to an agent.

Actions are undertaken in the pursuit of goals, which we may simply define as an agent's internal representation of a desired state of the world. An agent is likely to have multiple goals simultaneously, often resulting from biological needs; the states of hunger or fatigue will promote satiety or restfulness as short-term goals, which will encourage the actions of eating or sleeping to achieve them.

Achieving every goal therefore requires the execution of multiple distinct and potentially contradictory actions; one cannot eat and sleep simultaneously. An agent must meet all of its bodily needs in order to survive, yet it can only perform one meaningful action with any body part at a time. And for the best outcome, the agent's entire body should usually be directed towards performing a single action at a time.

It is therefore necessary for an agent to be able to select which action to perform at any given moment to effectively achieve its goals.

Requirements for Selection

To explore the process of action selection and assess the performance of selection mechanisms we must first explicate the parameters of successful selection. Prescott et al. (1999) propose that any effective selection mechanism must, at a minimum:

Exclude Incompatible Alternatives

The primary feature of effective selection is that only a single action from a set of incompatible options is selected at any one time. As a corollary to this, multiple actions may be selected simultaneously so long as they are not mutually incompatible.

Prefer Salience

McFarland (1989) proposes that given several otherwise indistinguishable inputs to a selection mechanism, the one of highest salience should be preferred, allowing the selection mechanism to be ignorant of the semantic content of inputs.

Select Cleanly

A winning action should be fully selected as soon as possible after emerging as the strongest candidate. No action should be partially selected and the selection strength should not be proportional to the input salience.

Prevent Distortion

Conversely, losing actions should be fully inhibited and have no influence over the mechanism's output. Selected inputs should face no interference from non-selected alternatives; in general, there is no benefit to dithering, indecision, or attempting to perform multiple incompatible actions simultaneously, and these outcomes should always be avoided.

Persist

McFarland (1989) also suggests that once selected, an action should persist with a lower salience input than required for initial selection. This helps to prevent dithering and allows actions to be completed when multiple inputs of similar salience are present.

However, many goals are impossible to fulfil with only a single action, such as acquiring territory or building a nest. Achieving such goals will often require the completion of several intermediate sub-goals (e.g. fighting rivals or collecting materials), and achieving each of these may involve sequences of multiple actions.

1.1.2 Action Sequences

Action sequences are a subset of the more general trait of cognitive sequencing, described by Savalia et al. (2016) as "the ability to perceive, represent and execute a set of actions that follow a particular order". Cognitive sequencing is necessary for many forms of advanced cognition such as planning, problem solving, and general skill learning, and action sequencing is fundamental to behaviours such as singing for birds, grooming for rats, and everyday activities such as tying shoelaces or buttoning shirts for humans.

Sequences exhibit several advantageous features that set them apart from a series of unrelated actions. Actions within a learned sequence are executed more rapidly and with smoother transitions than unrelated actions (Benecke et al., 1987), and once initiated sequences flow to completion with little or no conscious oversight (Graybiel, 2008) and suffer less interference from distractions (Verwey et al., 2010).

Action sequences are intrinsically hierarchical (Botvinick et al., 2009); a sequence is composed of multiple sub-actions and may also be part of a broader sur-action. This hierarchical organisation enables a wide range of higher cognitive processes (Balleine et al., 2015; Savalia et al., 2016) and underlies both goal-directed and habitual behaviours (Dezfouli and Balleine, 2013).

Goal-directed and Habitual Behaviours

All instrumental behaviours (including action sequences) may be described as either goal-directed or habitual depending primarily on whether the behaviour occurs in pursuit of

a specific outcome or in response to a particular stimulus (Balleine and O'Doherty, 2010; Dickinson, 1985). Colloquially, habits are often associated with troublesome vices, but many sequencing benefits depend on this response automatisation.

Habituation is the result of a necessary tradeoff between the accuracy of goal-directed behaviour and the speed of habitual behaviour (Keramati et al., 2011). Because there is a cognitive and temporal cost to analysing a problem and deciding on a solution, a clear incentive exists to automatise responses when possible.

To achieve this, individual actions within a sequence become 'chunked' together through repeated performance (Graybiel, 1998; Verwey, 2001) to form units that can be utilised without separately selecting every action. Chunking is computationally efficient and allows for rapid selection with low cognitive overhead (Ramkumar et al., 2016; Solopchuk et al., 2016), making the ability to automatise action sequences highly valuable. Indeed, a declining ability to learn and execute action sequences is a primary symptom of conditions such as Huntington's disease or Parkinson's disease (Agostino et al., 1992; Benecke et al., 1987).

Action chunks potentially constitute individually manipulable higher-order atomic actions (Verwey, 2001), but their increased selection speed comes at the cost of reduced accuracy. Because each chunk is selected and executed as a unit the efficacy of individual actions within a chunk cannot be assessed (Dezfouli et al., 2014), and environmental changes may result in habitual behaviours that are inefficient or counterproductive.

The formation of action chunks is necessary for automatising responses and habituating behaviours, but the process by which it occurs is not fully understood. To investigate this further we must first discuss the area of the brain responsible for action selection in general.

1.2 The Basal Ganglia

The basal ganglia (BG) are a collection of subcortical neural structures generally accepted to serve as the main vertebrate action selection mechanism (Berns and Sejnowski, 1996; Doya, 1999; Prescott, 2007; Redgrave et al., 1999), and are therefore heavily involved with related functions such as reinforcement learning, the formation of habits, and the initiation of movement (Lanciego et al., 2012). It is likely that all jawed vertebrates (Grillner and Robertson, 2016; Medina and Reiner, 1995) and potentially some jawless vertebrates (Reiner et al., 1998; Stephenson-Jones et al., 2011) possess at least rudimentary BG structures, and most mammalian BG structures and connections have homologues in these earliest examples (Reiner, 2010).

1.2.1 Structures

Core BG nuclei comprise the striatum, the subthalamic nucleus (STN), the globus pallidus (GP) or pallidum, and the substantia nigra (SN). Most of these structures may be further

subdivided based on their connectivity, functionality, or cytochemistry. Major striatal subregions include dorsomedial (DMS), dorsolateral (DLS), and ventral; pallidum incorporates external (GPe) and internal (GPi) segments in primates, or the entopeduncular nucleus in rodents and some other groups (Redgrave, 2007); substantia nigra is divided into pars compacta (SNc) and pars reticulata (SNr). Though not a part of the basal ganglia proper, many descriptions and models of the region also include portions of cortex and thalamus as their interactions with BG are key to its function.

1.2.2 Functional Architecture

Many aspects of BG functionality may be inferred from its structure and connectivity, illustrated in **Figure 1.1**. The regions linked by patterned input and output projections are suggestive of the primary function of the BG as a whole, and the arrangement of connections within and between structures allows for a deeper understanding of how that functionality is achieved.

Inputs

The striatum and STN are the main input nuclei for the BG, receiving excitatory inputs from thalamus and nearly all regions of cortex. Sensorimotor projections form a somatotopic map on their targets in striatum (Brown et al., 1998; Verstynen et al., 2012) and STN (Romanelli et al., 2004, 2005), and remaining inputs to both striatum and STN are also topographically organised (François et al., 1991; Nakano et al., 2000; Voorn et al., 2004).

This topographic organisation segregates the striatum into partially overlapping regions that receive inputs from separate functional domains (Draganski et al., 2008; Hunnicutt et al., 2016), creating what McGeorge and Faull (1989) describe as a "functional gradient" within the structure represented by coloured bands in **Figure 1.1**. Within each domain, afferents from reciprocally–connected cortical regions frequently converge onto localised striatal targets (Shipp, 2017; Yeterian and Van Hoesen, 1978), and specific limb movements or single body parts may be represented by striatal activity in units as small as a single neuron (Coffey et al., 2016; West et al., 1990).

These patterened, convergent corticostriatal projections may encode 'requests' or 'bids' for action (Prescott et al., 2006) that integrate command and contextual information about each request (Rueda-Orozco and Robbe, 2015; Shipp, 2017). Information about each request may be encoded in both the spatial distribution and overall activity of target populations (Koechlin and Burnod, 1996), allowing all relevant inputs to co-operate and compete for selection against other requests using a "common input currency" of salience (Redgrave et al., 1999); similarly convergent inputs from thalamus (McFarland and Haber, 2000) may encode significant sensory and attentional information that modulates this selection (Matsumoto et al., 2001; Minamimoto and Kimura, 2002).

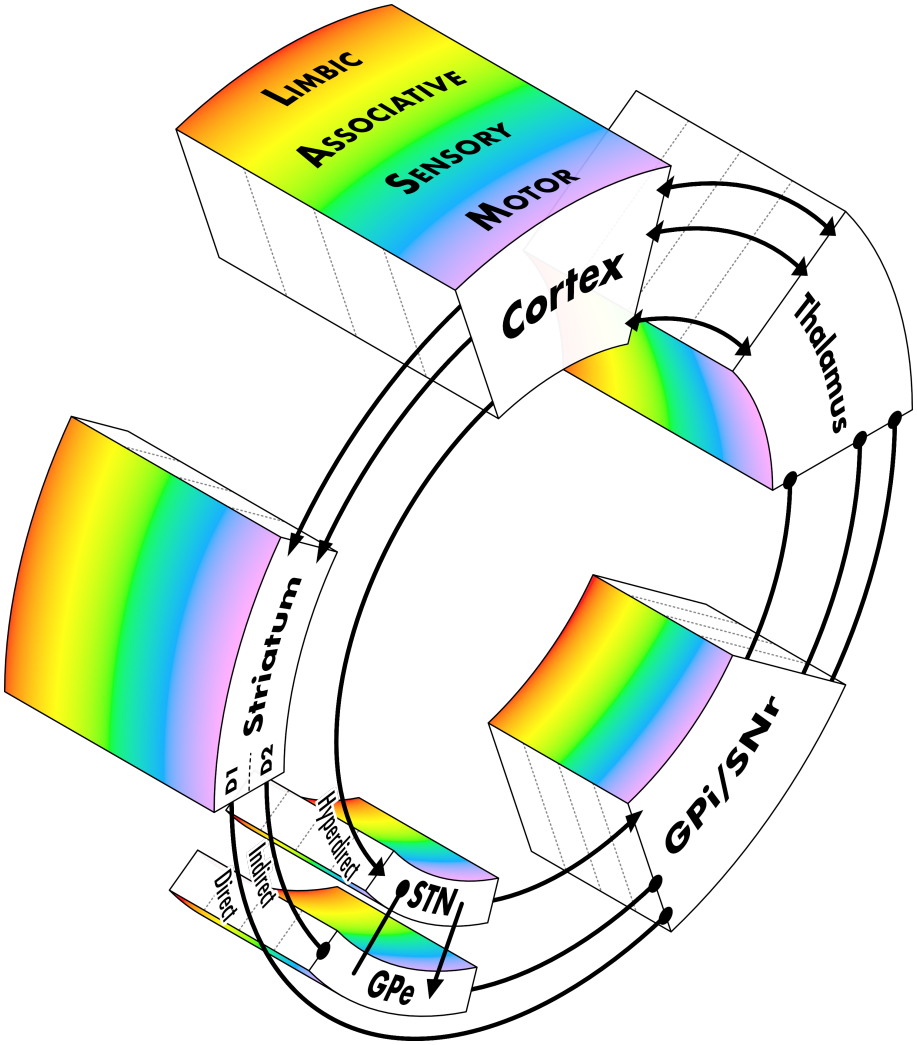

Figure 1.1: Principal connectivity of the basal ganglia–thalamocortical loop. Functional domains are mapped onto segregated but overlapping regions within each structure, and connections between structures preserve this separation while compressing the informational content. Inputs reach GPi/SNr via direct, indirect, and hyperdirect pathways and disinhibit corticothalamic feedback loops. Colours indicate functional domains, and connection terminators ► and ● indicate excitatory and inhibitory projections respectively. Diagram adapted from Redgrave (2007) and Shipp (2017).

Outputs

GPi and SNr are the main BG output nuclei in primates and project tonic inhibition to thalamus, which exerts inhibitory control over cortex. Their similar outputs, cytochemistry, and complementary roles in selection mean they are often grouped together as a single functional unit, though GPi outputs are more topographically segregated and correlate with sensorimotor activity while SNr outputs are more integrated and target associative regions (Kaneda et al., 2002; Parent and Hazrati, 1995b).

GPi and SNr outputs preserve the topographic organisation of their inputs (Mana and

Chevalier, 2001; Mengual et al., 1999; Middleton and Strick, 2002) and are organised into discrete information streams or 'channels' (Hoover and Strick, 1993) that relate to distinct aspects of behaviour or cognition (Middleton and Strick, 2000). Inhibited output activity in these channels may therefore represent the successful selection of a particular action request — a hypothesis supported by improvements in Parkinsonian patients following targeted pallidal lesioning (Kishore et al., 2000; Laitinen et al., 1992).

Tonic inhibition of thalamic targets allows the BG to control the flow of information to motor and cognitive systems (Deniau et al., 2007; Freeze et al., 2013), so targeted disinhibition of thalamus grants selected action requests control over the required systems. GPi/SNr disinhibition of thalamus forms a centre-surround pattern (Mink and Thach, 1993; Nambu et al., 2002) that ensures a high activation contrast and clean selection.

Loops

The principal connections between basal ganglia structures form segregated loops that largely preserve the topographic organisation of its inputs and segregation of functional domains (Middleton and Strick, 2000; Parent and Hazrati, 1995a). These loops partially converge as they transit the BG (Draganski et al., 2008; Yelnik, 2002), though their informational content is likely compressed rather than lost as it moves to structures with fewer neurons (Bar-Gad et al., 2003; Oorschot, 1996).

Selection processes are further segregated according to their secondary functional aspects; for example, separate loops may be responsible for predicting rewards at different time scales (Tanaka et al., 2004) and for comparing action outcomes to expectations (Stephenson-Jones et al., 2016). Habitual and goal-directed behaviours may also be processed in separate loops (Kim and Hikosaka, 2015), enabling the long-term storage of learned abilities (Hikosaka et al., 2017). The topographic organisation of information within the basal ganglia therefore extends beyond somatotopic body maps or separation of functional domains, but includes meta-informational aspects of selection.

This informational segregation allows the basal ganglia to perform parallel selection operations across multiple functional domains simultaneously without interference (Redgrave et al., 2011), and exaptation of this organisational feature has allowed the basal ganglia to grow from a mere selector of motor programs to a generalised selection hub additionally processing associative, limbic and cognitive information (Stephenson-Jones et al., 2011).

The integration of information throughout the basal ganglia is equally important. Learning in particular relies on integrating feedback information from previously selected actions with ongoing selection processes. This includes identifying events caused by previous actions (Redgrave and Gurney, 2006; Redgrave et al., 2008) and reward thus obtained, so that an agent may identify which of its actions lead to beneficial outcomes. Information from multiple cortical regions must also be integrated for the basal ganglia to successfully

co-ordinate complex activities or those comprising multiple actions (Bednark et al., 2015; Gobel et al., 2011; Lehéricy et al., 2006).

Informational integration occurs as a result of convergence between loops, but also due to connectivity within basal ganglia structures or inputs from external sources. The striatum (Hunnicutt et al., 2016) and thalamus (Haber and Calzavara, 2009) are notable loci of these latter forms of integration, and thalamic connections in particular form part of several secondary loops that modulate basal ganglia activity. Inputs from subcortical structures such as the superior colliculus (McHaffie et al., 2005) integrate information about the body's current physical state into the selection process, and return links from GPi/SNr complete these subcortical loops.

Cortex and thalamus also exhibit reciprocal excitatory connections forming a positive feedback loop (Haber and McFarland, 2001), which could ensure actions are always selected at maximum strength and may also represent a form of working memory that allows actions to persist once selected (Chambers et al., 2005).

Pathways

Connectivity between basal ganglia populations can also be described in terms of three distinct activity pathways. The *direct* and *indirect* pathways (Alexander and Crutcher, 1990) comprise projections from striatal medium spiny neurons (MSNs) directly to GPi/SNr or indirectly via GPe, as shown in **Figure 1.1**. Activity in these pathways has classically been associated with the promotion or inhibition of movement, with elevated activity levels in each corresponding to hyper- and hypo-kinetic states respectively (Albin et al., 1989; DeLong, 1990). This relationship is reinforced by evidence from optogenetic stimulation of rats that shows increased motion and reduced freezing in response to activation of direct–pathway MSNs, and vice-versa for stimulation of indirect–pathway MSNs (Kravitz et al., 2010; Sippy et al., 2015).

Several researchers have noted that movement initiation is marked by a transient increase in activity in both direct– and indirect–pathway MSNs (Cui et al., 2013; Isomura et al., 2013; Jin and Costa, 2010). This appears to directly contradict the promotion / inhibition paradigm; however, it has been suggested that co-activation of indirect–pathway MSNs may help to inhibit competing motor programs, thereby enhancing the execution of the selected action (Freeze et al., 2013; Sano et al., 2013). In support of this, it has been shown that ablation of indirect–pathway MSNs results in hyperactive, spontaneous movement potentially resulting from an inability to inhibit undesired actions (Durieux et al., 2009; Sano et al., 2003). The co-activation of direct– and indirect–pathway MSNs may therefore help prevent the selection of multiple incompatible action requests.

It has also been suggested that rather than simply promoting or inhibiting movement, activity in direct or indirect pathways could signal persistent reinforcement or transient

punishment respectively (Kravitz et al., 2012). According to this hypothesis, activity in the direct pathway identifies and encourages behaviour that results in rewarding outcomes, while activity in the indirect pathway discourages behaviour associated with aversive stimuli (Hikida et al., 2010; O'Hare et al., 2016).

Cortical afferents to striatum innervate direct– and indirect–pathway MSNs equally (Doig et al., 2010) but differentially (Wall et al., 2013); sensory and limbic afferents preferentially target the direct pathway while motor afferents preferentially target the indirect pathway (Wall et al., 2013). However, the functional implications of this are still speculative.

A third *hyperdirect* pathway has also been described that bypasses the striatum entirely, comprising projections from cortex to STN that transmit a fast-acting inhibitory signal (Nambu et al., 2002). This may help to specify the timing of selection (Frank, 2006; Nambu et al., 2002) and rapidly cancel selection already in progress when necessary (Aron and Poldrack, 2006; Schmidt et al., 2013).

More generally, it has been proposed that increased STN activation raises the evidence threshold required to make a decision thus resulting in slower, more accurate choices (Frank, 2006; Mansfield et al., 2011), while increased striatal activation lowers the evidence threshold resulting in faster but less accurate choices (Forstmann et al., 2008; Mansfield et al., 2011). It is therefore likely that interactions between all three pathways are necessary to dynamically optimise the speed / accuracy tradeoff inherent to efficient selection (Bogacz et al., 2010; Tewari et al., 2016).

1.2.3 Dopamine

Dopamine is a key modulatory neurotransmitter in the basal ganglia of vital importance for its normal function; though not of direct relevance to the present study a brief overview will be useful.

Dopamine is produced in the SNc and ventral tegmental area (VTA) (Björklund and Dunnett, 2007) and projected to other basal ganglia structures, especially the striatum (Bolam et al., 2000). It is released in tonic and phasic modes (Grace, 1991), both of which have been implicated in multiple functions at various time scales (Schultz, 2007).

All cell types in the striatal microcircuit (Section 1.3.3) express dopamine receptors, which can be classed as either D_1-like (comprising D_1 and D_5 receptors) or D_2-like (comprising D_2, D_3, and D_4 receptors) (Seeman and Van Tol, 1994), with broadly excitatory and inhibitory postsynaptic effects respectively. This bimodal action allows dopamine to exert a strong and complex modulatory influence on striatal function.

In particular, dopamine modulates MSN activity by acting on the D_1- and D_2-type receptors preferentially expressed by direct– and indirect–pathway neurons respectively (Gerfen and Surmeier, 2011). Activation of D_1 receptors enhances neuronal excitability by increasing

L-type Ca^{2+} currents (Hernández-López et al., 1997) and reducing somatic K^+ currents (Kitai and Surmeier, 1993), while activation of D_2 receptors decreases Ca^{2+} currents (Higley and Sabatini, 2010) and increases K^+ currents (Kitai and Surmeier, 1993) to reduce neural excitability.

Dopamine also modulates the activity of striatal interneurons; fast-spiking interneurons are depolarised through activation of the D_5 receptors they express and are further excited by activation of D_2 receptors on GABAergic presynaptic terminals which reduces inhibitory inputs (Bracci et al., 2002; Centonze et al., 2003). Tonically-active interneurons express D_2 and D_5 receptors and thus individual neurons experience both inhibitory and excitatory dopaminergic effects (Kreitzer, 2009).

These opposing effects allow dopamine to bias striatal activity, and balanced levels of tonic dopamine are necessary for normal striatal function. Heightened levels of dopamine enhance the excitability of direct–pathway MSNs, causing hyperactivity and habituation difficulties (Zhuang et al., 2001), while depressed dopamine levels cause preferential activation of the indirect pathway which contributes to the cardinal symptoms of Parkinson's disease (Mehler-Wex et al., 2006).

In contrast to the role of tonic dopamine, the phasic release of dopamine is commonly considered to signal reward prediction error (Nakahara et al., 2004; Schultz and Dickinson, 2000), though careful analysis of the timing of afferent spikes to dopaminergic neurons (Dommett et al., 2005) suggests that the phasic dopamine burst may instead signal that an unexpected sensory event was self-caused (Redgrave et al., 2008). The phasic release of dopamine has been strongly implicated in learning (Smith-Roe and Kelley, 2000) and motivation (Bromberg-Martin et al., 2010), and it has been proposed that this targeted, transient modulation of striatal excitability may also assist selection by enhancing the contrast of cortical afferents (Nicola et al., 2004).

. . .

The striatum is thus heavily implicated in the integration and organisation of relevant action information, appropriate selection or inhibition of actions, and learning and adaptation in response to reward or environmental changes. To explore how neuropeptides might influence this process we must first understand more about the composition and structure of the striatum itself.

1.3 The Striatum

The striatum is the basal ganglia's primary input structure, comprising the caudate nucleus and putamen (generally referred to as dorsal striatum) and the nucleus accumbens (or ventral striatum) (Gerfen and Wilson, 1996). It takes its name from the striated appearance

caused by the large number of fibres passing through from cortex (Steno, 1669), though under a microscope it has a broadly uniform appearance and no clear structural organisation (Kreitzer, 2009).

1.3.1 Neural Populations

Medium Spiny Neurons

The medium spiny neuron (MSN) is the most common neural type in the striatum, comprising over 95% of the striatal population in cats (Kemp, 1968) and over 97% in rats (Rymar et al., 2004) (though potentially only 80–85% in primates (Graveland and Difiglia, 1985; Wu and Parent, 2000)), with the remainder consisting of several types of interneuron. MSNs are GABAergic and principally synapse onto other MSNs (Wilson and Groves, 1980), though overall connectivity is likely sparse (Humphries et al., 2010; Tunstall et al., 2002). They express AMPA and NMDA glutamate rececptors (Nicola et al., 2000) and are commonly placed into one of two categories based on their relative expression of dopamine receptor subtypes and downstream projection targets within the basal ganglia (Smith et al., 1998). D_1 (or striatonigral) MSNs preferentially express D_1-type dopamine receptors, co-release the neuropeptides substance P and dynorphin with GABA (Reiner and Anderson, 1990), and project to internal regions of globus pallidus and SNr. D_2 (or striatopallidal) MSNs preferentially express D_2-type dopamine receptors, co-release enkephalin with GABA (Reiner and Anderson, 1990), and project to external regions of globus pallidus. These categorisations are not absolute; MSNs project minor axonal arbors to the 'opposite' downstream structure (Parent et al., 1995) and many MSNs express both D_1 and D_2 receptors (Aizman et al., 2000; Surmeier et al., 1996), though this dual expression may be largely limited to neurons in patch regions (Biezonski et al., 2015).

At rest, MSNs exhibit a hyperpolarised membrane potential (Kita et al., 1985). In this 'Down' state, depolarisation and spiking are suppressed by an inwardly–rectifying potassium current (Wilson and Kawaguchi, 1996). A significant and sustained glutamatergic influx blocks the potassium current (Blackwell et al., 2003) and shifts the MSN into an 'Up' state in which a much lower additional glutamatergic influx readily causes spiking. The 'Up' state is dependent on continued glutamatergic influx, and MSNs therefore exhibit bimodality but not bistability in their spiking behaviour (Humphries et al., 2009a; Kasanetz et al., 2006).

This bimodality is differentiated in D_1 and D_2 MSNs. D_2 MSNs fire at higher rates in response to current injection (Kreitzer and Malenka, 2007), though it is not correct to say that they are generally more excitable than D_1 MSNs. Rather, D_1 MSNs require greater stimulation to bring them into an 'Up' state, but are then more prone to spiking. Conversely, D_2 MSNs enter the 'Up' state in response to lower levels of input but subsequently require a greater number of synchronous inputs to fire (Kreitzer, 2009).

Interneurons

The striatum contains at least three varieties of interneuron that modulate striatal activity but do not project outside the structure: fast-spiking, low-threshold, and tonically active. Of these, only fast-spiking interneurons are of direct relevance to the current study and so the others will be discussed only briefly.

Fast Spiking

Fast-spiking (or parvalbumin–positive) interneurons (FSIs) are GABAergic and comprise approximately 1% of the rat striatal population (Humphries et al., 2010; Luk and Sadikot, 2001). They have a lower input resistance and a more rapid spike than other striatal cells (Kawaguchi, 1993), and are preferentially found in dorsolateral striatum (Bennett and Bolam, 1994). FSIs receive input from widespread cortical regions (Lapper et al., 1992; Parthasarathy and Graybiel, 1997) and thalamus (Sidibé and Smith, 1999), and exert a strong inhibitory influence on the large number of MSNs to which they synapse (Koós and Tepper, 1999). They also form synapses with other FSIs, as well as electrical gap junctions (Kita et al., 1990) that do not appear to cause correlated firing (Berke, 2008) but may allow FSI networks to preferentially respond to correlated inputs (Hjorth et al., 2009).

Low Threshold Spiking

Low-threshold spiking (or somatostatin–positive) interneurons (LTSes) exhibit a high input resistance and a relatively depolarised resting potential (Kawaguchi, 1993). They also receive input from cortex and thalamus (Sidibé and Smith, 1999) and form weak synapses with MSNs, FSIs and TANs (Do et al., 2013; Gittis et al., 2010). In addition to somatostatin, LTS interneurons release neuropeptide Y and nitric oxide synthase (Kreitzer, 2009), thought to play a role in long-term plasticity within striatum (Calabresi et al., 1999).

Tonically Active

Tonically active (or cholinergic) interneurons (TANs) comprise less than 1% of striatal neurons (Rymar et al., 2004) and are physically very large, with a somatic diameter potentially greater than 40 µm (Tepper and Bolam, 2004). TANs receive sparse excitatory input from cortex (Thomas et al., 2000) and thalamus (Lapper and Bolam, 1992), and their low-frequency tonic firing (Kreitzer, 2009) is paused in response to certain stimuli (Aosaki et al., 1994). Their ability to exert inhibitory control over striatal MSNs (Pakhotin and Bracci, 2007) and to modulate dopamine transmission (Threlfell and Cragg, 2011) make them of particular interest for studies of learning.

Neuropeptides

D_1 and D_2 MSNs co-release neuropeptides with GABA, the effects and functions of which are not fully known. Substance P (SP) is co-released from D_1 MSNs, but is also found throughout the body and has been linked to a wide variety of physiological and neurological functions (Muñoz and Coveñas, 2014). Within the central nervous system, SP is implicated in pain (Muñoz and Coveñas, 2014), neuroinflammation (Thornton and Vink, 2012), and a range of behavioural effects.

Intraventricular administration of SP causes increased respiratory movements in cats and rabbits (von Euler and Pernow, 1954, 1956) and excessive grooming in rats (Van Wimersma Greidanus and Maigret, 1988), while SP injections directly to the rat substantia nigra cause circling movements (James and Starr, 1977) and increased rearing and grooming (Kelley and Iversen, 1979). Similarly, SP agonist injections to the ventral pallidum increase motor activity (Napier et al., 1995).

Conversely, blocking SP's action decreases amphetamine–induced behaviour in rats (Gonzalez-Nicolini and McGinty, 2002), while mice lacking the NK_1 receptors on which SP acts exhibit behavioural traits comparable to ADHD (Porter et al., 2015; Yan et al., 2009, 2010) and lose the rewarding aspect of opiates (Murtra et al., 2000).

Within the striatum, SP has been found to have direct and indirect excitatory effects on target MSNs (Blomeley and Bracci, 2008; Blomeley et al., 2009), acting principally on presynaptic NK_1 receptors. The identification of an excitatory intra-striatal signal is of particular interest as it implies a co-operative function for what has been thought to be a purely competitive inhibitory network. Given the striatum's established role in action selection and SP's notable effects on motor output, this discovery further reinforces the case for exploring its role in action selection.

Another neuropeptide, enkephalin, is co-released with GABA from D_2 MSNs and has also been linked to behavioural changes. Injection of an enkephalin analog into the VTA stimulates feeding behaviour (Cador et al., 1986), and may either stimulate or suppress behavioural responding (Kelley et al., 1989), locomotion, and rearing (Kalivas et al., 1983) depending on the dose and time elapsed since administration.

Within the striatum, enkephalin has inhibitory effects on target MSNs (Blomeley and Bracci, 2011), acting principally on presynaptic μ-opioid receptors. It has been proposed that enkephalin plays a role in limiting excessive MSN excitation (Steiner and Gerfen, 1998) and in balancing the activation of direct and indirect pathways (Presti and Lewis, 2005). It has also been suggested that the motor dysfunctions associated with Borna disease are caused by heightened levels of striatal enkephalin (Solbrig et al., 2002). Thus, enkephalin's role in action selection also merits further investigation.

The impact of these neuropeptides on action selection is not yet known. However, their widespread presynaptic modulation of glutamatergic inputs to striatum suggests a potential

ability to dynamically regulate action requests. We hypothesise that the modification of cortical input salience via neuropeptide interactions represents a neurological basis for the chunking of sequential actions.

1.3.2 Organisational Features

In addition to the D_1 / D_2 MSN distinction, the striatum exhibits several other important organisational features of tangential interest to the present study that will therefore only be briefly discussed.

Patch / Matrix Organisation

The striatum is interspersed with a complex arrangement of tubular striosomes or 'patches' (Gerfen, 1984) whose connectivity and neuropeptide concentrations differ from the surrounding 'matrix'. Patches express high levels of μ-opioid receptors (Herkenham and Pert, 1981), SP and enkephalin (Bolam et al., 1988; Wang et al., 2007), and reduced levels of acetylcholine (Graybiel and Ragsdale, 1978). However, within patch regions enkephalin acts primarily on δ-opioid receptors on intra-striatal connections (Banghart et al., 2015), suggesting enkephalin's function in patches is differentiated from matrix regions. Patches also preferentially receive inputs from limbic cortical regions and project along the indirect pathway to SNc rather than SNr (Gerfen, 1984), and it has been proposed that this connectivity is optimised for high-conflict or emotional decision-making (Friedman et al., 2015).

Conversely, matrix neurons receive input preferentially from sensory and motor cortices (Donoghue and Herkenham, 1986) and appear to be functionally organised into 'matrisome' clusters (Flaherty and Graybiel, 1993) that project to localised regions of globus pallidus. However, it is not clear whether matrisomal clusters correspond to individual BG output channels (Shipp, 2017).

A further subset of 'exo-patch' neurons have been identified in matrix regions with similar neurochemistry and connectivity to patch neurons (Smith et al., 2016), further confusing the functional roles of these compartments. The significance and function of patch / matrix segregation is still very much under debate.

Dorsal / Ventral Regions

The striatum has several significant structural and functional subdivisions. The principal distinction is between dorsal and ventral regions; dorsal striatum incorporates the caudate and putamen while nucleus accumbens forms the ventral striatum (Gerfen and Wilson, 1996), though these delineations vary slightly between species and researchers.

Dorsal striatum is often further subdivided into dorsomedial (DMS) and dorsolateral (DLS) regions, which preferentially receive input from associative and sensorimotor cortices respectively (McGeorge and Faull, 1989). DMS and DLS have been implicated in goal-directed and habitual behaviours respectively (Devan et al., 2011), and evidence suggests that the transition from goal-directed to habitual behaviours involves a shift from processing in DMS to DLS (Alloway et al., 2017; Redgrave et al., 2010; Yin et al., 2009)

Ventral striatum comprises core and shell divisions (Záborszky et al., 1985) and receives input primarily from limbic cortical regions (Brog et al., 1993). Ventral striatum receives dopamine innervation from VTA rather than SNc (Fields et al., 2007), is generally implicated in motivational and emotional selection operations (Cardinal et al., 2002; Kelley, 2004), and may be important for integrating the emotive and affective value of actions into the selection process (Knutson et al., 2009; Mannella et al., 2013).

1.3.3 Functional Architecture

Connectivity

Inputs from cortex and thalamus are topographically organised throughout striatum (Section 1.2.2) and converge on target MSNs (Flaherty and Graybiel, 1993; Takada et al., 1998). However, cortical projections to target FSIs are widespread (Ramanathan et al., 2002) and FSI activity is uncorrelated with task-related behaviours (Berke, 2008). As FSIs are capable of inhibiting large MSN populations with little input (Section 1.3.1) this corticostriatal connectivity may represent an on–centre, off–surround pattern that could support selection via activation of localised MSN populations representing specific action requests or acting as "local controllers" (Graybiel and Grafton, 2015) for sequences. D_1 and D_2 MSNs also form asymmetric connections within the striatum (Taverna et al., 2008) which may help it perform as a threshold device during selection (Bahuguna et al., 2015).

It was previously thought that the abundance of inhibitory MSNs in the striatum implied the existence of mutually inhibitory networks that performed selection with a 'winner–take–all' (WTA) operation (Beiser and Houk, 1998; Groves, 1983). However, WTA relies on strong reciprocal inhibitory connections that are rare in striatum (Jaeger et al., 1994; Tunstall et al., 2002). Although it has been estimated that MSNs may each receive inputs from over 700 partner MSNs (Humphries et al., 2010) striatal connectivity is nonetheless sparse (Planert et al., 2010; Taverna et al., 2008) and recent research has placed greater emphasis on the feedforward inhibitory role of striatal interneurons (Tepper et al., 2008) and the function of stiatal microcircuits.

Striatal Microcircuits

Analysis and modelling of systems at the level of microcircuits — defined by Grillner et al. (2005) as a "minimal number of interacting neurons that can collectively produce a functional output" — is increasingly useful for understanding the functionality of complex interconnected neural populations. Microcircuit descriptions of the striatum have proven useful for understanding the formation of cell assemblies (Carrillo-Reid et al., 2008; Humphries et al., 2009b), Parkinsonian impairments (Jáidar et al., 2010), and other disorders (Gittis and Kreitzer, 2012).

Bolam et al. (2006) describe a canonical striatal microcircuit that incorporates the key connectivity features outlined above and is able to perform selection via a combination of convergent corticostriatal inputs and both feedforward (FSI → MSN) and feedback (MSN → MSN) GABAergic inhibition. It also incorporates the key modulatory effects of dopamine and acetylcholine (**Figure 1.2**).

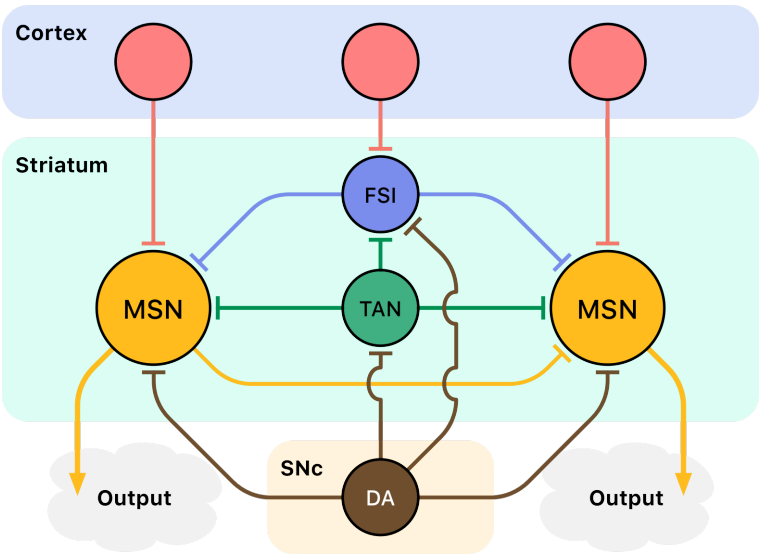

Figure 1.2: Simplified schematic of the canonical striatal microcircuit, omitting axonal and dendritic detail. Two interconnected MSNs receive input from cortex, fast-spiking and tonically active interneurons, and provide output to other structures. FSIs also receive input from cortex and TANs. Dopaminergic neurons in SNc innervate all striatal microcircuit neurons. Diagram adapted from Bolam et al. (2006).

This microcircuit encompasses the key connectivity and modulatory features discussed so far and could therefore be considered to represent a functional striatal unit. As the microcircuit is able to select, it is a potentially useful tool for exploring the impact of neuropeptides on selection.

1.4 Computational Models

Computational models are a uniquely versatile scientific tool for synthesising and explaining previous observations, for rapidly iterating experiments to explore relevant hypotheses, and for generating predictions or suggestions about future data collection (Epstein, 2008). This is particularly useful in neuroscience, where the quantity and complexity of available data often precludes any useful implicit or informal modelling and where highly dynamical systems often produce unexpected and unpredictable results. A computational model is therefore an ideal environment for drawing together currently available data about striatal processing and testing our hypotheses about neuropeptide contributions to selection. Neural models have been created at all levels of description, so we will need to utilise a model that simulates basal ganglia function at an appropriate level of detail.

The basal ganglia has been extensively modelled for over two decades, so before creating our own model we should explore existing options to determine if any of them are suitable for our purposes. Albert Einstein is paraphrased as having said that "everything should be as simple as possible, but not simpler" (Calaprice, 2010), and to that end we should look for a model that contains those components necessary for testing our hypotheses and omits those that are extraneous.

Substance P and enkephalin modulate the glutamatergic EPSPs of target neurons (Section 1.3.1), so our model must capture the membrane potential and synaptic currents of individual MSNs in addition to their spiking activity. MSNs receive significant inhibition from FSIs, and both MSNs and FSIs receive excitatory input from cortex (Section 1.2.2), so these principal microcircuit features should be included.

However, to explore the impact of striatal neuropeptides on selection resulting from basal ganglia output we require a model of the entire basal ganglia. Although spiking BG models exist (Chersi et al., 2013; Humphries et al., 2006), our BG model will principally serve as an embedding architecture for evaluating the results of striatal computation and does not need to be modelled in such detail. We may therefore need to employ two separate models at differing levels of description and integrate them into a unified whole.

In Section 1.2.2 we presented evidence that inputs to basal ganglia may represent 'bids for action' that feed into parallel loops and may become selected as the result of thalamic disinhibition. Therefore our model must support multiple parallel loops. Secondly, although we are not specifically investigating the roles of action pathways in selection we are studying the separable effects of neuropeptides released from direct– and indirect–pathway MSNs, thus requiring the separation of striatum and its downstream targets. Thirdly, to explore the behaviour of the entire basal ganglia loop our model should include key external structures such as cortex and thalamus, utilising the positive feedback loop between cortex and thalamus as a form of working memory to sustain selection.

1.4.1 Basal Ganglia Models

There are numerous basal ganglia models, and several detailed reviews compare their objectives, design, and performance (Beiser et al., 1997; Gillies and Arbuthnott, 2000; Helie et al., 2013; Schroll and Hamker, 2013). Many of these models are immediately unsuitable for our purposes; some include only a subset of structures or connections, and many are constructed to explore a specific functional aspect such as categorisation or learning, imposing unwanted structural requirements on inputs or processing. Our BG model is intended to enhance the biological plausibility and aid in the evaluation of our striatal model, and we therefore wish to make few additional assumptions about the underlying computations.

Berns and Sejnowski (1998) describe one potential BG model that includes the main BG structures, thalamus, and the connections between them. It is also designed to explore sequence production and makes no significant assumptions about the structure of inputs. The model does lack a cortical population and is therefore not a closed loop, but a more significant problem is the utilisation of GPe–STN oscillations as a form of working memory. Experimental and modelling data suggest that these oscillations arise from dopamine depletion (Magill et al., 2001) and are unlikely to have a specific role in selection (Humphries et al., 2006), making this model unsuitable for our research.

Gurney et al. (2001a,b) introduce a model that proposes the basal ganglia perform action selection by means of input signal selection, using the 'common currency' of input salience. This model does not impose any additional structure on inputs, supports the separation of action channels into multiple parallel loops, and was expanded by Humphries and Gurney (2002) to include a complete basal ganglia–thalamocortical loop. As this includes all the key structures and features needed to support our striatal microcircuit model, we will adapt this model for our work.

1.4.2 Striatal Models

Many early striatal models focused on its presumed role in pattern recognition (Beiser et al., 1997) and were often abstracted from the underlying biology. An exception is Wickens et al. (1995) who describe a biophysical MSN model used to explore connectivity changes in Huntington's disease. Their model includes conductance–based inhibitory synapses and estimates of MSN–MSN connectivity, but lumps together all excitatory inputs and omits other connectivity and modulatory influences.

Gruber et al. (2003) detail another biologically–inspired MSN model that simulates a greater number of ion currents and includes D_1 receptor modulation. However, this model focuses on a single neuron and lacks D_2 receptor modulation, and would therefore require a significant amount of effort to expand into a functional microcircuit model. Conversely, Gurney and Overton (2004) describe a striatal model that incorporates both MSNs and FSIs

and is able to perform selection but consists of leaky integrator neurons lacking necessary membrane dynamics.

Moyer et al. (2007) outline the most detailed model of MSN behaviour to date, constructing a multi-compartment model that includes a wealth of intrinsic and synaptic currents and dopaminergic modulation of D_1 and D_2 receptors. It is likely that this level of detail would make simulations of large populations infeasible.

However, Izhikevich (2003) describes a computationally simple and flexible canonical neuron model that captures the biological plausibility of Hodgkin and Huxley (1952) and for which Humphries and Gurney (2007) describe an efficient solution method. Humphries et al. (2009a) use these developments to create a spiking MSN model that accurately reproduces the behaviour of the detailed Moyer et al. model in a computationally tractable form.

Striatal interneuron models are rare, though Kotaleski et al. (2006) describe an FSI model including intrinsic and synaptic currents at a level of detail that would be suitable. However, Humphries et al. (2009b) expand on their previous MSN model to construct a microcircuit including FSIs, obviating the need to integrate the Kotaleski et al. model. This model was subsequently expanded still further to include detailed connectivity estimates based on 3D simulations of axonal and dendritic arbors (Humphries et al., 2010). Although lacking cholinergic modulation, this model represents the most complete and biologically plausible striatal microcircuit of which we are aware and will form the basis of our study.

1.5 Research Questions

To summarise:

- The basal ganglia perform action selection and functions such as habit learning
- During habit learning component actions are chunked together into units
- The neurological basis of chunking is unclear
- The striatum receives cortical inputs that may represent action requests
- MSNs release neuropeptides SP and enkephalin that modulate cortical inputs
- The function of these neuropeptides is not yet known
- *Hypothesis:* One function of striatal neuropeptides may be to facilitate chunking

This overarching hypothesis invites the formulation of specific research questions:

1. How do striatal neuropeptides facilitate chunking?
2. Under what conditions do neuropeptides facilitate chunking?
3. What are the differentiable effects of substance P and enkephalin on selection?
4. To what extent do striatal inputs, connectivity, and neuropeptide modulation impact sequence execution?
5. Are sequence chunks stored in the striatum?
6. How do other striatal features impact sequence selection?

To answer these questions, we will use computational models of the relevant neural structures as a biologically plausible testbed for experimentation.

1.5.1 Thesis Outline

Chapter 2: Neuron and Neuropeptide Models

We instantiate the MSN model of Humphries et al. (2009a) in our modelling environment, and we create and validate novel phenomenological models of substance P and enkephalin's glutamatergic EPSP modulation.

Chapter 3: Neuropeptides and Sequence Selection

We instantiate the Humphries et al. (2009b) FSI model and the Humphries and Gurney (2002) basal ganglia–thalamocortical loop model in our modelling environment. We recreate the Humphries et al. (2009b) striatal microcircuit model (incorporating our neuropeptide model) and integrate it with the BG loop model to create a novel hybrid model. We test and evaluate the impact of striatal neuropeptides on the selection of various action groups.

Chapter 4: Striatal Topography and Action Selection

We expand the striatal microcircuit model to simulate 1 mm^3 of striatal tissue, following the connectivity algorithm outlined in Humphries et al. (2010). We assess the impact of biologically grounded (rather than merely statistically accurate) connectivity on selection and explore other changes introduced by physical topography.

Chapter 5: General Discussion

We discuss the validity and limitations of the current approach, summarise our contributions to the research literature and suggest potential avenues for future work.

Chapter 2

Neuron and Neuropeptide Models

2.1 Introduction

The neuropeptides substance P and enkephalin have been shown to have excitatory (Blomeley et al., 2009) and inhibitory (Blomeley and Bracci, 2011) effects respectively on glutamatergic EPSPs in MSNs. The impact of this modulation on selection is not yet known, but we hypothesise it enables the preferential selection of cortical inputs that represent a learned action sequence. Testing our hypothesis will require a neuropeptide model that captures this glutamatergic modulation, and as one does not currently exist it will be necessary to create one.

To simulate EPSP modulation it will be necessary to utilise a neuron model that captures the membrane potential, synaptic currents, and spiking activity of individual neurons. As we are interested only in the effects of glutamatergic EPSP facilitation or inhibition and not the underlying biological mechanisms, our neuropeptide model will be purely phenomenological. Finally, because we intend to employ the neuropeptide model at scale in a simulated striatal microcircuit consisting of several thousand neurons, computational efficiency will be a primary concern in development.

2.2 Model Construction

Unless otherwise specified all models in this and later sections are instantiated in the SpineCreator environment (Cope et al., 2017) using the SpineML syntax (Richmond et al., 2014) and executed using the BRAHMS simulation engine (Mitchinson et al., 2010) with a forward Euler solver and a 0.1 ms timestep (Humphries and Gurney, 2007).

2.2.1 Medium Spiny Neuron

The spiking MSN model uses the canonical neuron model of Izhikevich (2007) with updates from Humphries et al. (2009a) that capture the effects of dopamine modulation. In this model, v is the neuron membrane potential and u is the dominant ion channel contribution (which for MSNs is the slow A-type potassium current (Nisenbaum et al., 1994)), which becomes a recovery variable:

$$C\dot{v} = k(v - v_r)(v - v_t) - u + I \tag{2.1}$$

$$\dot{u} = a[b(v - v_r) - u] \tag{2.2}$$

with a reset condition:

$$\text{if } v > v_{\text{peak}} \text{ then } v \leftarrow c\,,\ u \leftarrow u + d$$

In this description, C is capacitance, v_r and v_t are the resting and threshold potentials respectively, I is a current source, and c is the reset potential. Dimensionless parameters a, b, d, and k further tune the neuron model (Izhikevich, 2007); a is a recovery time constant, d describes net spike–activated currents affecting post-spike behaviour, and parameters b and k are derived from the neuron's rheobase and input resistance. MSN model parameters are unchanged from the values in Humphries et al. (2009a) and are listed in Table A.7.

Dopamine Modulation

Dopamine modulation of MSNs is modelled with the process described in Humphries et al. (2009a). The proportion of active dopamine receptors is represented by ϕ_1 and ϕ_2 in the interval [0, 1] for D_1 and D_2 MSNs respectively, and these values are used to modify MSN parameters. For D_1 MSNs, two parameters are updated to be dependent on the activation of D_1 receptors:

$$v_r \leftarrow v_r(1 + K\phi_1) \tag{2.3}$$

$$d \leftarrow d(1 - L\phi_1) \tag{2.4}$$

Updates to these two parameters model the enhanced hyperpolarising influence of inwardly–rectifying potassum currents and the lowered activation threshold for depolarising L-type Ca^{2+} currents respectively (Hernández-López et al., 1997). For model tuning, we use $\phi_1 = \phi_2 = 0.8$ per Humphries et al. (2009a), but for neuropeptide model validation and later simulations we use $\phi_1 = \phi_2 = 0.3$ as in Tomkins et al. (2014).

Dopamine modulation of D_2 MSNs is captured by altering the value for k to model the inhibitory effect on the slow potassium current caused by activation of D_2 receptors (Moyer et al., 2007):

$$k \leftarrow k(1 - \alpha\phi_2) \tag{2.5}$$

For an in-depth explanation of the rationale for these modifications we refer the reader to Humphries et al. (2009a).

Synaptic inputs

Synaptic input to MSNs is modelled by:

$$I = I_{\text{ampa}} + I_{\text{gaba}} + B(v)I_{\text{nmda}} \tag{2.6}$$

Each input z (where z is one of AMPA, GABA or NMDA) is further modelled by:

$$I_z = \bar{g}_z h_z (E_z - v) \tag{2.7}$$

with $\bar{g}$ representing maximum conductance and E_z the reversal potential of the synaptic current.

The single exponential model of post-synaptic currents h_z is modified with the inclusion of a saturation effect from Tomkins et al. (2012):

$$\dot{h}_z = \frac{-h_z}{\tau_z} \quad \text{and} \quad h_z(t) \leftarrow h_z(t) + \left[1 - \frac{h_z(t)}{\omega_z}\right] S_z(t) \tag{2.8}$$

Here, τ_z is the synaptic time constant, ω_z is the maximum number of synapses or receptor groups and $S_z(t)$ is the number of presynaptic spikes of type z arriving at all the neuron's receptors at time t. As only NMDA currents are susceptible to saturation in normal conditions (Clements et al., 1992) the AMPA and GABA saturation value ω_z is set high enough to have negligible effect.

In Equation 2.6 the $B(v)$ term models the voltage–dependent magnesium plug in NMDA receptors, which per Jahr and Stevens (1990) is given by:

$$B(v) = \frac{1}{1 + \frac{[\text{Mg}^{2+}]_0}{3.57} \exp\left(-0.062v\right)} \tag{2.9}$$

where $[\text{Mg}^{2+}]_0$ represents the equilibrium concentration of magnesium ions.

The synaptic input equations are also modified to account for dopaminergic modulation. D_1 modulation of NMDA–evoked EPSPs is captured with:

$$I^{\text{D1}}_{\text{nmda}} = I_{\text{nmda}}(1 + \beta_1 \phi_1) \tag{2.10}$$

and D_2 modulation of AMPA-evoked EPSPs with:

$$I^{\text{D2}}_{\text{ampa}} = I_{\text{ampa}}(1 - \beta_2 \phi_2) \tag{2.11}$$

where β_1 and β_2 are scaling coefficients that determine the relationship between the occupancy of dopamine receptors and the effect magnitude. Values for dopamine modulation and synaptic equations may be found in Table A.9.

2.2.2 Neuropeptides

Neuropeptide action is simulated in two stages. The amount of neuropeptide released in response to a given level of MSN activity is calculated using a simple sum of exponentials, and this value is converted into a facilitation or inhibition effect multiplier with a tuned response curve.

Thus, for a single spike–induced neuropeptide release event at time t_i the amplitude $a_p^i(t)$ of neuropeptide p induced by this event is given by:

$$a_p^i(t) = S_p \left[\exp\left(\frac{-(t-t_i)}{\tau_p^f} \right) - \exp\left(\frac{-(t-t_i)}{\tau_p^r} \right) \right] \tag{2.12}$$

where p is either SP or enkephalin. τ_p^r and τ_p^f represent neuropeptide release rise and fall time constants respectively, and S_p is the number of incoming spikes causing release of neuropeptide p. Multiple events over a period of time combine to form a net amplitude:

$$A_p(t) = \sum_i a_p^i(t) \tag{2.13}$$

The net amplitude $A_p(t)$ of neuropeptide release determines the resulting modulatory effect $N_p(t)$, which is normalised using the Weibull cumulative distribution function and a scaling factor β_p:

$$N_p(t) = \beta_p \left[1 - \exp\left(-\frac{A_p(t)}{\lambda_p} \right)^{\kappa_p} \right] \tag{2.14}$$

The neuropeptide modulatory effect is scaled with the Weibull function as it allows for greater control over slope parameters compared to a standard sigmoid function, which is necessary to achieve the required tuning of (and differentiation between) both neuropeptide models. This effect is appended to the synaptic input equation (Equation 2.7) giving a final form for glutamate input to MSNs:

$$I_z^{\text{ms}} = \bar{g}_z h_z (E_z - v) \Big[1 + N_{\text{sp}}(t - \tau_{\text{sp}}^d) \Big] \Big[1 - N_{\text{enk}}(t - \tau_{\text{enk}}^d) \Big] \tag{2.15}$$

where z is either AMPA or NMDA. The multiplicative interaction between neuropeptides implicitly assumes that a single input may be simultaneously modulated by both SP and enkephalin, which we consider to be a reasonable assumption given the available data. Additionally, reformulating the equation so that NP interactions are additive has little effect on the input current modifications.

As the model is purely phenomenological, the delay between MSN activity and the onset of

neuropeptide effects (Blomeley et al., 2009; Blomeley and Bracci, 2011) is captured using a fixed time offset τ_p^d. Values for all neuropeptide parameters are listed in Table A.10.

2.3 Model Validation

2.3.1 Medium Spiny Neuron

We reconstruct the MSN model exactly as described in Humphries et al. (2009a) and therefore forego repetition of the full suite of validation tests performed in that paper. We instead recreate the current–frequency (f–I) and input frequency–output frequency (f–f) response curves to ensure no errors are introduced during model transcription.

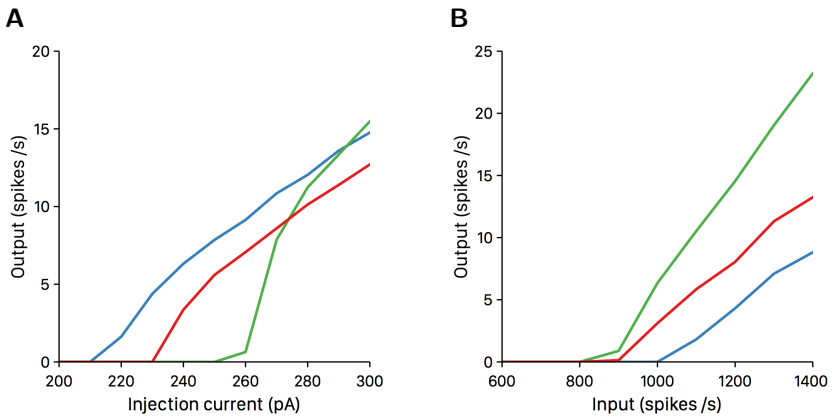

Figure 2.1: Current–frequency (f–I) and input frequency–output freqneucy (f–f) response curves for the Humphries et al. (2009a) MSN models instantiated in SpineCreator. **(A)** f–I and **(B)** f–f response curves for baseline MSN without dopamine (red), D_1 MSN (green), and D_2 MSN (blue) match those from Humphries et al. (2009a).

As in Humphries et al. (2009a), we construct and calibrate three versions of the MSN model to ensure baseline model accuracy and correct dopamine modulation. **Figure 2.1** shows MSN model response to injection current and spiking inputs. As in Humphries et al. (2009a), D_1 receptor activation inhibits spiking in response to low levels of current injection but causes increased activation in response to greater current injection, and an enhanced response to all levels of spiking input. Conversely, D_2 receptor activation causes greater activation in response to all levels of injection current but a reduced response to spiking input. As expected, these results match Humphries et al. (2009a), confirming the correct transcription of the reduced MSN model.

2.3.2 Neuropeptides

Substance P

Substance P is known to have both direct (Blomeley and Bracci, 2008) and indirect (Blomeley et al., 2009) effects on target MSNs; only the indirect effects are modelled here as it is unlikely that neuropeptides mediate direct communication between MSNs (Blomeley et al., 2009).

Neurophysiological recordings show that SP has a presynaptic facilitatory effect on subsequent glutamate inputs to MSNs it targets; in a pair of MSNs A and B where A projects to B, a burst of five spikes over 50 ms in MSN A elicits on average a 14% increase in glutamatergic EPSP amplitude in MSN B 100 ms after the first spike (Blomeley et al., 2009). No facilitation is seen at 50 ms after the first spike, and only residual facilitation is seen after 250 ms.

When antidromic spikes are evoked in MSNs, glutamatergic facilitation due to SP release is ∼40% after 250 ms and ∼22% after 500 ms (Blomeley et al., 2009). Bath application of SP increases the amplitude of glutamatergic EPSPs by 47% on average (Blomeley and Bracci, 2008). These data formed the primary fitness criteria for the SP model; **Figures 2.2A,B** show a comparison of neurophysiological data and model performance for SP, confirming that the phenomenological model captures the facilitatory effect.

Enkephalin

Enkephalin has a similar but inhibitory presynaptic effect on glutamatergic inputs to MSNs. In a similar paired–recording experiment, a burst of five spikes in MSN A elicits on average a 17.1% inhibition of glutamatergic EPSP amplitude in MSN B 500 ms after the first spike (Blomeley and Bracci, 2011). No inhibition is seen at 250 ms after the first spike, and minimal inhibition is seen at 1000 ms.

Evocation of antidromic spikes results in an average inhibition of 29.6% after 500 ms, becoming undetectable after 2 s (Blomeley and Bracci, 2011). No data were available on inhibitory effects as the result of bath application. These data formed the primary fitness criteria for the enkephalin model; **Figures 2.2D,E** show a comparison of neurophysiological data and model performance for enkephalin, confirming that the phenomenological model captures the inhibitory effect.

Calibration Procedure

Calibration of each neuropeptide model thus required the fixing of six variables; the NP release rise, fall, and onset delay time constants τ_r, τ_f, and τ_d, tuned response curve values λ and κ, and an overall effect multiplier β. The β value for each neuropeptide was fixed at

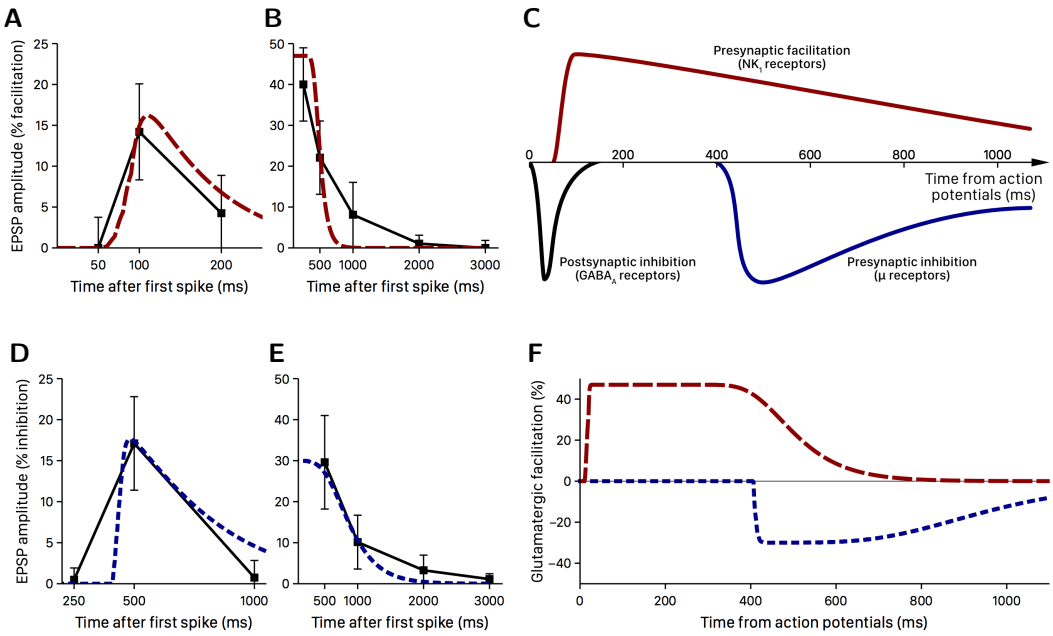

Figure 2.2: Neuropeptide models validated against neurophysiological data. (A) Facilitation of glutamatergic EPSPs by SP after release induced by five spikes from a neighbouring MSN *in vitro* (black) and *in silico* (red). **(B)** Facilitation of glutamatergic EPSPs by SP after release induced by antidromic stimulation of MSNs *in vitro* (black) and *in silico* (red). **(C)** Qualitative estimate of the timecourse of MSN interactions after a burst of spikes causing a release of GABA, SP (acting on NK_1 receptors) or enkephalin (acting on μ-opioid receptors) from Blomeley and Bracci (2011). **(D)** Inhibition of glutamatergic EPSPs by enkephalin after release induced by five spikes from a neighbouring MSN *in vitro* (black) and *in silico* (blue). **(E)** Inhibition of glutamatergic EPSPs by enkephalin after release induced by antidromic stimulation of MSNs *in vitro* (black) and *in silico* (blue). **(F)** Timecourse of the effects of SP (red) and enkephalin (blue) on glutamatergic EPSPs after a burst of spikes from model MSNs.

the maximum experimentally observed effect for that neuropeptide and the τ_d onset delay time constant was set to the minimum observed delay between neuropeptide release and onset of glutamate modulation in Blomeley et al. (2009) or Blomeley and Bracci (2011). Preliminary manual tuning established reasonable values for τ_r, leaving three free variables for each neuropeptide.

Two experiments were constructed to mimic the paired–recording experiments from Blomeley et al. (2009) and Blomeley and Bracci (2011). A single model MSN was provided with five spikes from another model MSN over a 50 ms window, with the resulting effect on subsequent glutamatergic inputs compared to physiological data at 50 ms, 100 ms, and 200 ms after the first spike for SP and 250 ms, 500 ms, and 1000 ms after the first spike for enkephalin. Our model cannot produce antidromic spikes, so to approximate their evocation a single model MSN was provided with five spikes from each of ten separate model MSNs over a 50 ms window. The impact on subsequent glutamatergic inputs was compared to

physiological data at 250 ms and 500 ms after the first spike for SP, and 500 ms and 2 s after the first spike for enkephalin.

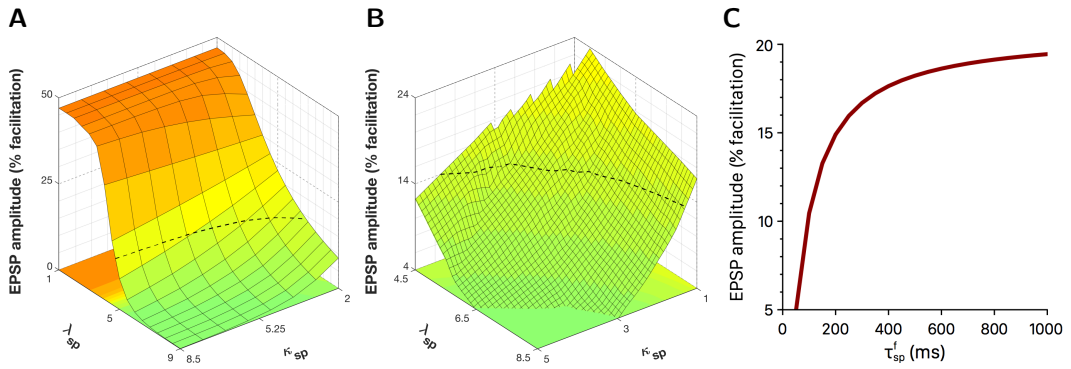

Figure 2.3: Search space for SP modulatory values λ_{sp}, κ_{sp}, and τ^f_{sp}. EPSP facilitation at 100 ms after innervation by five MSN spikes shows the relative effects of changes to each value. Dashed line indicates desired 14% EPSP facilitation. Initial testing at low resolution **(A)** provided an estimated range for each value, and detailed exploration of this narrower search space region **(B)** revealed potential solutions. **(C)** Impact of SP fall time constant τ^f_{sp} on EPSP facilitation with final values of $\lambda_{sp} = 5.5$ and $\kappa_{sp} = 2.5$.

Calibration for each neuropeptide occurred in two stages; first, both experiments were repeated using a wide range of τ_f, λ, and κ values at low resolution to obtain a rough estimate of values that produced results in line with the data. This estimate was then used to generate a narrower range of values that was tested at high resolution, and the set of values that produced EPSP modulation with the lowest mean absolute error across all time intervals in both experiments were selected for inclusion in the model. **Figure 2.3** illustrates the SP calibration search space for the first paired–recording experiment only; EPSP modulation is maximally impacted by changes to each value where $\lambda \approx 5$, $\kappa \approx 2$, and $\tau_f \approx 200$ms as shown by the steep slope gradient in these regions (**Figures 2.3A,C**), so it is perhaps unsurprising that this search space region (**Figure 2.3B**) provided the final values for both neuropeptide models, listed in Table A.10.

Figure 2.2C shows a qualitative estimate of the timecourse of neuropeptide action adapted from Blomeley and Bracci (2011), and **Figure 2.2F** shows the glutamatergic effects of the model neuropeptides after calibration.

2.4 Discussion

We have recreated a previously validated model of striatal MSNs incorporating dopamine modulation and used this to create a novel model of neuropeptide–mediated glutamatergic EPSP modulation. Using a two-phase simulation that calculates neuropeptide release amplitude and facilitation separately enables us to accurately capture the effects of two different neuropeptides over both short and long timescales with modifications to a single

model. Importantly, the model remains computationally simple and is thus suitable for inclusion in large-scale spiking striatal models.

This simplicity does introduce some minor inaccuracies. Glutamate inhibition from enkephalin in response to low levels of input is slightly overstated at ∼1 s, falling slightly outside the margin of error of observations (**Figure 2.2D**). Similarly, glutamate facilitation from substance P in response to high levels of input falls off more rapidly at ∼500 ms than in observations (**Figure 2.2B**). However, we consider it unlikely that these minor deviations will cause meaningful inaccuracies when the model is employed in more complex simulations.

Our model necessarily makes some assumptions regarding currently unknown factors. We assume that the amount of neuropeptide released scales linearly with spiking activity and GABA release, and that the maximum observed effects in a lab setting are appropriate limits for the *in silico* effect. We also do not attempt to account for the temperature difference between *in vitro* observations and living tissue and the potentially slower timecourse of lab-based results.

We further assume that all neurons release identical levels of neuropeptide in response to a given input, and that all neurons obtain identical levels of glutamate facilitation or inhibition in response to neuropeptide release. Experimental data and biological plausibility suggests that both these assumptions are false; however, simulating an accurate distribution of release and reaction values is of dubious merit, would greatly increase the computational complexity, and given the data currently available is unlikely to increase the model's overall accuracy.

Our model is purely phenomenological, and therefore captures the effects of glutamate modulation but makes no attempt to simulate or explain the biological processes underlying dynamics such as the ∼400 ms delay prior to the onset of enkephalin's action. We therefore implicitly assume that such delays are constant in all circumstances; future neurophysiological studies into these dynamics may invite an updated version of the model with greater biological accuracy.

Chapter 3

Neuropeptides and Sequence Selection

Chapter 3 is based on research previously published by Buxton et al. (2017), reproduction and adaptation of which is permitted under the Creative Commons Attribution License (CC BY).

3.1 Introduction

We propose that the action and interaction of striatal neuropeptides plays a key role in encoding action sequences within corticostriatal networks by allowing successive actions in a learned sequence to be preferentially selected over comparable non-sequence actions. This would assist with rapid, smooth transitions between sequential actions while still allowing exceptionally salient action requests to interrupt sequences mid-execution.

We test this hypothesis using a spiking model of a GABAergic striatal microcircuit that incorporates the MSN and neuropeptide models from Chapter 2, additional striatal components, and inputs representing cortical action requests. We use several different neuropeptide connectivity configurations to explore how the presence or lack of neuropeptide signalling between striatal regions representing different action requests influences the selection and inhibition of those requests.

We embed the striatal microcircuit model in a model of the basal ganglia–thalamocortical loop so the results and implications of striatal computation may be assessed at the level of motor cortex output, providing a closer correspondence to behaviour than analysis of MSN spiking. This also allows us to incorporate feedback into the model by constructing a cortico–basal ganglia–thalamic loop that represents a complete sensorimotor loop and utilises a corticothalamic positive feedback loop to represent a form of working memory.

3.2 Striatal Microcircuit Model

We recreate the GABAergic striatal microcircuit model described in Humphries et al. (2009b), utilising the MSN and neuropeptide models described in Section 2.2 and a fast-spiking interneuron model.

3.2.1 Fast-Spiking Interneuron

The dopamine–modulated FSI model of Humphries et al. (2009b) builds on the canonical Izhikevich (2007) model by extending Equation 2.1 to:

$$C\dot{v}_{\text{fs}} = k[v_{\text{fs}} - v_r(1 - \eta\phi_1)](v_{\text{fs}} - v_t) - u_{\text{fs}} + I \tag{3.1}$$

and implementing a nonlinear term for u:

$$\dot{u}_{\text{fs}} = \begin{cases} -au_{\text{fs}} & \text{if } v_{\text{fs}} < v_b \\ -a[b(v_{\text{fs}} - v_b)^3 - u_{\text{fs}}] & \text{if } v_{\text{fs}} \geq v_b \end{cases} \tag{3.2}$$

The nonlinear u term in Equation 3.2 allows the FSI model to display Type 2 dynamics, and the resting potential v_r is increased with the $(1 - \eta\phi_1)$ term to model the effects of modulation by D_1 dopamine receptors. FSIs do not express D_2 receptors on their membranes (Centonze et al., 2003) so these are not modelled. Intrinsic properties for the FSI model are unchanged from Humphries et al. (2009b) and are listed in Table A.8.

NMDA receptors are rare on FSIs (Blackwell et al., 2003), so only synaptic input to AMPA or GABA receptors is modelled. However, FSIs also express dendrodendritic gap junctions (Koós and Tepper, 1999), giving total current contributions of:

$$I = I_{\text{ampa}} + I_{\text{gaba}} + I_{\text{gap}} \tag{3.3}$$

D_2 receptor modulation of GABA input to FSIs is captured by:

$$I_{\text{gaba}}^{\text{fs}} = I_{\text{gaba}}(1 - \epsilon_2\phi_2) \tag{3.4}$$

Gap junctions between FSIs i and j are modelled as compartments with voltage v_{ij}^*, with dynamics of:

$$\tau\dot{v}_{ij}^* = (v_i - v_{ij}^*) + (v_j - v_{ij}^*) \tag{3.5}$$

where τ is a voltage decay time constant and v_i and v_j are membrane potentials for the

FSI pair. The current introduced by the FSI pair is:

$$I^*_{\text{gap}}(i) = g(v^*_{ij} - v_i) \quad \text{and} \quad I^*_{\text{gap}}(j) = g(v^*_{ij} - v_j) \tag{3.6}$$

with g as the effective conductance of the gap junction. Total gap junction input I_{gap} to any FSI is thus the sum of all contributions I^*_{gap}. Synaptic properties for the FSI model are listed in Table A.9.

3.2.2 Model Connectivity

The striatal microcircuit model is composed of 6,000 model MSNs divided into two equal groups of 3,000 D_1 and D_2 MSNs. These are both further subdivided into six 'action channels' $c_1 \ldots c_6$ of 500 neurons each, representing the striatal targets of six distinct cortical action requests. The MSN population is complemented by an additional 60 FSIs — 1% of the MSN population (Humphries et al., 2010; Luk and Sadikot, 2001) — that are all innervated by each action request.

For MSN → MSN, FSI → MSN, and FSI → FSI connections an exhaustive all–to–all list of connections between the two populations is probabilistically culled according to the expected number of connections for each type (Table 3.1). Culling is entirely independent of the action channel represented by a given neuron. Each neuron therefore makes contact with and receives inputs from a statistically accurate number of partners, but the model lacks any topography and is similar to the *random* model from Tomkins et al. (2014). All 500 D_1 or D_2 MSNs in channel c_n also project to the single neuron in the basal ganglia–thalamocortical loop model (Section 3.3) representing channel c_n in GPi/SNr or GPe respectively.

Table 3.1: Expected number of striatal contacts, from Humphries et al. (2010)

Connection type	Contacts
MSNs → 1 MSN	728 ± 25.7
FSIs → 1 MSN	30.6 ± 5.39
1 FSI → MSNs	3017 ± 45.1
FSIs → 1 FSI	12.8 ± 3.37
FSI gap junctions	0.65 ± 0.81

Neuropeptide projections co-exist with MSN → MSN GABA connections targeting both D_1 and D_2 MSNs (Blomeley et al., 2009; Blomeley and Bracci, 2011; Yung et al., 1996) and never appear on their own, though GABA connections without an associated neuropeptide projection are permitted. Several neuropeptide connectivity configurations (described in Section 3.5.1) determine which GABA connections co-release a neuropeptide.

3.3 Basal Ganglia–Thalamocortical Loop Model

Although our research focuses on the impact of neuropeptides on the striatal response to cortical inputs, directly analysing changes in MSN activity risks misinterpretation of the data or the unintentional imposition of bias. We therefore embed the striatal microcircuit in a rate-coded model of the basal ganglia–thalamocortical loop to create a more biologically plausible selection system that allows the results of striatal computation to be quantitatively analysed at the level of motor cortex output, providing a closer correspondence to behaviour.

3.3.1 Model Connectivity

The basal ganglia–thalamocortical loop (BG loop) model is largely unchanged from the *thalamocortical loop* (TC) model described in Humphries and Gurney (2002). Five populations represent major structures within the basal ganglia–thalamocortical complex, each composed of six separate leaky integrators representing a distinct 'action channel' $c_1 \ldots c_6$ as in the striatal microcircuit model.

With the exception of STN, BG loop populations are connected with one–to–one links that preserve the channel-based architecture of the basal ganglia (Table A.13). The combination of diffuse excitatory STN and focused inhibitory striatal projections to GPe and GPi/SNr models an off–centre, on–surround pattern of activation (Mink and Thach, 1993; Nambu et al., 2002) that allows selective disinhibition of the motor cortex (MCtx) action channel corresponding to the selected action request. GPi/SNr neurons therefore govern the flow of information to motor cortex via ventrolateral thalamus (VLT) (Romanelli et al., 2005) and their inhibition predicts motor activity (Deniau et al., 2007; Freeze et al., 2013).

3.3.2 Neural Dynamics

Neural dynamics and connectivity for the BG loop model are unchanged from Humphries and Gurney (2002), with a few exceptions:

1. The connection weight $w_{\text{sc–mc}}$ from sensory cortex to motor cortex is reduced to 0.5 from 1 in order to emphasise the role of GPi/SNr in disinhibiting VLT and thus promoting the selected action request.

2. The connection weight $w_{\text{vlt–mc}}$ from VLT to motor cortex is increased to 1.05 from 1 in order to allow for a stable MCtx–VLT feedback loop while GPi/SNr output remains below 0.05, the selection threshold in Humphries and Gurney (2002).

3. Leaky integrator populations D1 & D2 representing striatum are not used beyond initial validation and are replaced with the striatal microcircuit model.

The activation a of each leaky integrator population is dependent on afferent input u and constant decay factor k:

$$\dot{a} = k(a - u) + u \tag{3.7}$$

Output y at time t is representative of mean firing rate and is bound between 0–1, where it is governed by a piecewise linear output function with threshold θ:

$$y(t) = F(a(t), \theta) = \begin{cases} 0 & \text{if } a(t) \leq \theta \\ a(t) - \theta & \text{if } \theta < a(t) < 1 - \theta \\ 1 & \text{if } a(t) \geq 1 - \theta \end{cases} \tag{3.8}$$

External driving input to the model from rate–converted Poisson source sc_i is provided to channel i of MCtx and STN, with input to all other populations provided by the preceding population's output. Dynamics for net input u_i and output y_i for channel i of each population are as follows:

$$\begin{aligned}
\text{D1:} \quad u_i^{\mathrm{d1}'} &= (w_{\mathrm{d1}} y_i^{\mathrm{sc}} + w_{\mathrm{d1}} y_i^{\mathrm{mc}})(1 + \chi_{\mathrm{d1}'}) \\
y_i^{\mathrm{d1}'} &= F\left(a_i^{\mathrm{d1}'}, 0.2\right) \\
\text{D2:} \quad u_i^{\mathrm{d2}'} &= (w_{\mathrm{d2}} y_i^{\mathrm{sc}} + w_{\mathrm{d2}} y_i^{\mathrm{mc}})(1 - \chi_{\mathrm{d2}'}) \\
y_i^{\mathrm{d2}'} &= F\left(a_i^{\mathrm{d2}'}, 0.2\right) \\
\text{STN:} \quad u_i^{\mathrm{stn}} &= w_{\mathrm{stn}} y_i^{\mathrm{sc}} + w_{\mathrm{stn}} y_i^{\mathrm{mc}} + w_{\mathrm{gp}} y_i^{\mathrm{gp}} \\
y_i^{\mathrm{stn}} &= F\left(a_i^{\mathrm{stn}}, -0.25\right) \\
\text{GPe:} \quad u_i^{\mathrm{gp}} &= w_{\mathrm{gp}} \sum_j^n y_j^{\mathrm{stn}} - y_i^{\mathrm{d2}} \\
y_i^{\mathrm{gp}} &= F\left(a_i^{\mathrm{gp}}, -0.2\right) \\
\text{GPi/SNr:} \quad u_i^{\mathrm{snr}} &= w_{\mathrm{snr}} \sum_j^n y_j^{\mathrm{stn}} - y_i^{\mathrm{d1}} - w_{\mathrm{gp}} y_i^{\mathrm{gp}} \\
y_i^{\mathrm{snr}} &= F\left(a_i^{\mathrm{snr}}, -0.2\right) \\
\text{VLT:} \quad u_i^{\mathrm{vlt}} &= w_{\mathrm{mc}} y_i^{\mathrm{mc}} + w_{\mathrm{snr}} y_i^{\mathrm{snr}} \\
y_i^{\mathrm{vlt}} &= F\left(a_i^{\mathrm{vlt}}, 0\right) \\
\text{MCtx:} \quad u_i^{\mathrm{mc}} &= w_{\mathrm{mc}} y_i^{\mathrm{sc}} + w_{\mathrm{vlt}} y_i^{\mathrm{vlt}} \\
y_i^{\mathrm{mc}} &= F\left(a_i^{\mathrm{mc}}, 0\right)
\end{aligned}$$

Weights and properties for the BG loop model are listed in Table A.16.

3.4 Hybrid Model Integration

The integration of rate-coded and spiking populations into a single model necessitates the creation of neural interconnects to translate activity rate to spiking output and vice-versa.

3.4.1 Rate–to–Spike Conversion

Activity rate output from motor cortex in the basal ganglia–thalamocortical loop model is converted to spike trains and projected to each MSN and FSI in the striatal model. Conversion is achieved by assigning a Poisson spike generator and a random number generator P to every striatal projection from motor cortex, and generating a spike every timestep on any connection where the activity rate y_{mc} is greater than P up to a maximum possible firing rate r_{max}. The rate–to–spike conversion is thus achieved by:

$$\text{Emit spike if } y_{\text{mc}} r_{\text{max}} \tau_{\text{bg}} > P \tag{3.9}$$

where τ_{bg} is the timestep value for the overall simulation. Rate–to–spike properties are listed in Table A.9.

3.4.2 Spike–to–Rate Conversion

Converting spike train outputs into a normalised activity rate is necessary for the connections between model input and BG loop populations, and for projections from striatal MSNs to GPe and GPi/SNr. An instantaneous measurement of spiking output is insufficient to generate a continuous activity rate, so a sum of exponentials captures a dynamic rate r of spiking in sensory cortex or striatum which is converted to an activity rate y in the range 0–1. The dynamic rate of spiking activity is thus captured by:

$$r_s(t) = \sum_i S_s \left[\exp\left(\frac{-(t - t_i)}{\tau_s^f} \right) - \exp\left(\frac{-(t - t_i)}{\tau_s^r} \right) \right] \tag{3.10}$$

where s is sensory cortex, D_1 MSN, or D_2 MSN, S_s is the number of spikes arriving from population s, and τ_s^r and τ_s^f govern the duration over which spiking activity should be averaged. The dynamic rate r_s is normalised with the Weibull cumulative distribution function to provide a final activity rate y_s:

$$y_s(t) = 1 - \exp\left(-\frac{r_s(t)}{\lambda_s} \right)^{\kappa_s} \tag{3.11}$$

Values for spike–to–rate properties are listed in Table A.16.

3.4.3 Inputs

The striatal microcircuit and BG loop models both receive input from populations of Poisson spike generators representing action requests from sensory cortex[1]. Each spike source population comprises 500 separate Poisson spike generators, collectively defined as representing sensory cortex activity corresponding to a single action request. Each spike generator $\mathrm{sc}_c^i, 1 \leq i \leq 500$ has a one–to–one connection with a single D_1 MSN and a single D_2 MSN, and motor cortex leaky integrator mc_c projects to D_1 and D_2 MSNs $\mathrm{ms}_c^i, 1 \leq i \leq 500$. All 500 spike generators in each action channel also project to the single leaky integrator representing that channel in STN and MCtx within the BG loop model. The first 60 spike generators in each channel also have a one–to–one connection to a single FSI and each motor cortex neuron projects to all 60 FSIs. Each striatal neuron therefore receives the same number of afferent connections from sensory and motor cortices; **Figure 3.1** illustrates the connectivity between sensory cortex spike generators, striatal neurons, and motor cortex leaky integrators.

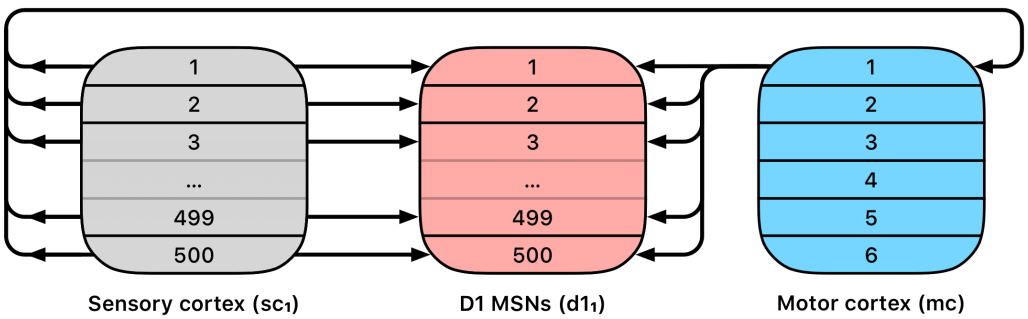

Figure 3.1: Hybrid model input connectivity. Partial input connectivity for channel c_1 showing one–to–one connections from sensory cortex (grey) to striatum (red), convergent connections from sensory cortex to motor cortex (blue), and divergent connections from motor cortex to striatum. Inputs to D_2 MSNs and STN match those to D_1 MSNs and motor cortex respectively.

All six action channels are therefore uniquely represented in the striatal model by a distinct population of 1,000 MSNs split evenly into D_1 and D_2 subtypes, while each FSI receives input from all channels. This corticostriatal connectivity reflects the convergence of cortical afferents from functionally related cortical regions on target MSNs (Flaherty and Graybiel, 1993; Takada et al., 1998) and widespread input to target FSIs (Berke, 2008; Ramanathan et al., 2002) that provide distributed inhibition of MSNs. This represents an on–centre, off–surround pattern of corticostriatal connectivity that could support selection via activation of specific MSN populations.

Figure 3.2 illustrates population–level connectivity of the complete hybrid model.

[1]We refer to inputs as originating from sensory cortex so that the model may represent a complete sensorimotor loop; however, inputs are entirely abstracted and could plausibly originate from any non-motor cortical source that provides the striatum with patterned inputs.

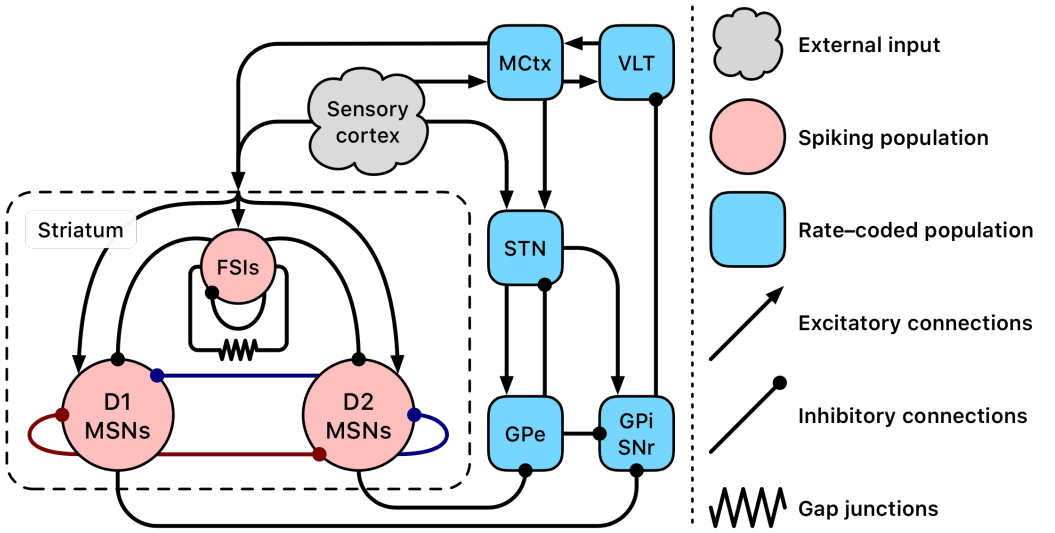

Figure 3.2: Complete connectivity of the hybrid basal ganglia–thalamocortical loop model. Coloured connections indicate co-release of substance P (red) or enkephalin (blue) with GABA.

3.5 Action Groups and Selection Metrics

To explore the effects of striatal neuropeptides on selection, several different action selection scenarios are simulated. To maintain consistency with Humphries and Gurney (2002) all simulations are conducted with a six channel model. Action channels $c_1 \ldots c_4$ represent distinct actions within an action group, and channel c_5 represents a generic action that marks the end of every action group. Channel c_6 is a null channel that receives no external input, except for Section 3.6.4 where it is utilised as an intrusive 'distractor' action.

To assess model performance, we define three types of action groups:

Action Series

The term *action series* refers to any group of action requests that occur one after the other but have no preferred semantic order; for example, taking a sip of tea; putting on glasses; scratching the nose. The specific order in which these actions occur is not important, but it is important that selection of more than one does not occur simultaneously. For the purposes of what follows, the four actions $1 \ldots 4$ in any order comprise a valid action series.

Action Sequence

An *action sequence* refers to a specific group of action requests that must occur in a predefined semantic order; for example, raising the foot from the accelerator, moving it across to the brake, and pressing it down on the brake. These actions must occur one after the other and in a specific order. For the purposes of what follows, only the four actions $1 \ldots 4$ in the order $1 \rightarrow 2 \rightarrow 3 \rightarrow 4$ comprise a valid action sequence.

Action Clique

An *action clique* refers to a group of action requests that may or may not occur in a predefined order but must exclude other specific actions from occurring alongside actions within the clique. For example, putting a teabag in a mug, pouring milk into a mug, pouring water into a mug, but *not* putting instant coffee into a mug. The specific order in which tea is made is unimportant, but it is important to not make coffee at the same time. For the purposes of what follows, the four actions 1…4 in any order coupled with the exclusion of action 6 comprise a valid action clique.

Each action group thus consists of several distinct action requests, and each action request consists of two phases of activity. The onset of each action request is marked by a transient burst of activity from Poisson generators representing an action request from sensory cortex, followed by a quiet 'gap' period during which the model receives no external stimulus but may sustain selection via feedback between motor cortex and thalamus (Chambers et al., 2005; Haber and McFarland, 2001). The transient burst and the subsequent gap together comprise the 'valid' selection period for that action request, which ends at the onset of the next action request. This input scheme is comparable to phasic activity in macaque prefrontal cortex corresponding to saccades during a learned sequence (Fujii and Graybiel, 2003). **Figure 3.3** illustrates these input features and shows example rate outputs from selected populations.

To ensure consistent initial activation of an action group, the first action request in a trial is always active from 100–400 ms at 2,000 spikes/s, and the post-transient gap is 200 ms for all action requests. The input duration and salience of the remaining three action requests in a group varies *between* trials, but each uses the same input salience and duration *within* a trial. Channel c_5 is always active at 2,000 spikes/s following the gap period after the last action request in the group to mark the end of the trial. Every action group is therefore initiated by an input of standard strength, and the remaining action requests within a group all have identical input duration, input salience, and valid length.

In the results that follow we use suprathreshold activity in motor cortex as a quantitative measure of selection. An action channel is determined to be selected when the motor cortex activity rate for that channel is above a threshold of $\theta = 0.95$. A score of 1 is assigned for every simulation timestep of channel c_n validity where that channel is selected, and a score of -1 is assigned for every timestep outside of the valid period where the channel is selected. A score of -1 is also assigned for each timestep any channel is selected at the same time as another, even if this occurs during a valid period. Subthreshold activity in motor cortex at any time scores 0.

The total score is averaged across the duration of the entire action group presentation to give a final overall selection score between -1 and 1. Channel c_5 is not considered part of any action group and so selection of this channel is not scored.

Figure 3.3 shows a time series of activity rate outputs from key populations in response

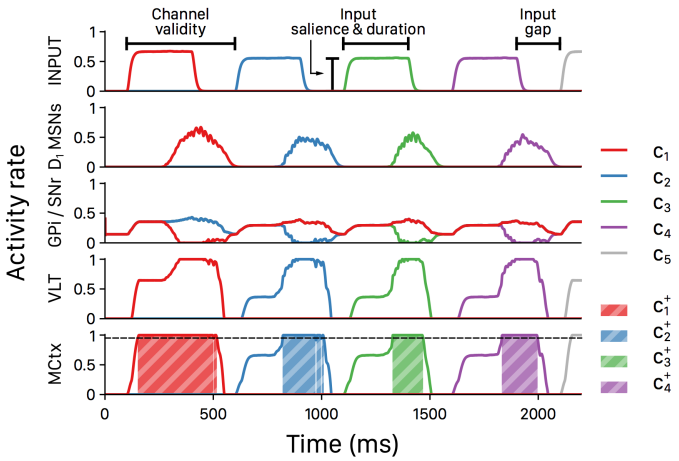

Figure 3.3: Activity rates of key model populations in channels $c_1 \ldots c_5$ in response to a typical action group presentation. Input features and MCtx selection threshold $\theta = 0.95$ are indicated; selection of channel *n* resulting in a positive selection score is highlighted as c_n^+.

to presentation of a typical action series with input duration 300 ms and input salience 1,600 spikes/s; the selection score for this presentation is 0.4256. Selection scores from multiple such presentations are aggregated to provide a summary of the overall selection performance of four neuropeptide connectivity configurations.

3.5.1 Neuropeptide Connectivity Configurations

We define four neuropeptide connectivity configurations, described in Table 3.2. There are no other differences between model configurations, and in all cases the sole effect of neuropeptide release is the facilitation or inhibition of subsequent glutamatergic inputs to the postsynaptic neuron.

Table 3.2: Striatal microcircuit model neuropeptide connectivity configurations

Configuration	Substance P co-released from	Enkephalin co-released from
Control	—	—
Diffuse	All D_1 MSN projections	
Unidirectional	Only D_1 MSNs projections $c_n \rightarrow c_{n+1}$ (for $n < 4$)	All D_2 MSN projections
Pruned	All D_1 MSN projections except $c_1 \rightarrow c_6$	

3.5.2 Potential Neuropeptide Benefits

To assess the impact of neuropeptides on the selection of actions and sequences, we must define potential benefits or enhancements to selection. We consider five potential selection benefits that may occur as the result of neuropeptide modulation:

Temporal locking
: If action A's input downslope overlaps with action B's input upslope, ensure that switching always takes place at either action A offset or action B onset.

Input enhancement
: Ensure that actions forming part of a sequence are selected even when their cortical inputs have reduced salience and/or duration when compared with non-sequential actions.

Clean transitions
: Ensure that a sharp, clean transition between actions takes place even when action onsets and offsets are gradual.

Rapid transitions
: Ensure that transitions between actions occur more rapidly than they would between equivalent non-sequential actions and that periods where no action is selected are reduced.

Distraction reduction
: Ensure that otherwise salient non-sequential actions are not selected and do not interrupt the execution of the action sequence.

The presence of neuropeptides is unlikely to provide all or even most of these benefits, but any realised benefits will fall into one of these categories.

3.6 Simulation Results

3.6.1 Hybrid Model Validation

To confirm that the hybrid model performs in line with the model from Humphries and Gurney (2002) we recreated an experiment from that paper exploring the model's response to a transient change in input strength. **Figure 3.4** shows the activity rate of all neural populations in the Humphries and Gurney TC model compared to the hybrid model. Both models show similar activity rates and overall response to external input, suggesting that the conversion between spiking and rate output is suitably tuned and that the model is behaving in line with expectations. However, structural changes in the hybrid model give rise to some notable differences; MSN membrane dynamics cause a delay between the onset of input activity and striatal output in the hybrid model, and internal connectivity within the striatum that is not present in the Humphries and Gurney model prevents simultaneous striatal activity in multiple channels. This also allows the sustained selection of a single action request during the transient event, a feature which required the inclusion of an additional population representing the thalamic reticular nucleus in Humphries and Gurney (2002).

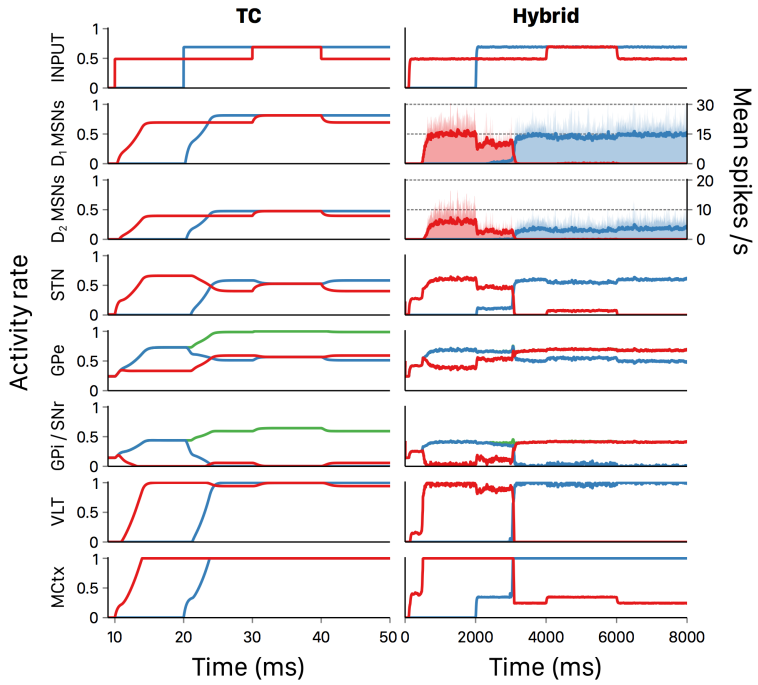

Figure 3.4: Validation of the hybrid BG model against the TC model. Comparison of activity rates of channels c_1 (red), c_2 (blue), and c_3 (green) in response to a transient input event confirms that population–level behaviour of the hybrid model matches the Humphries and Gurney (2002) rate-coded TC model. Mean spikes/s for striatal populations averaged over a 15 ms rolling window.

3.6.2 Diffuse Neuropeptides Enhance Action Series Selection

In order to explore the ability of the model neuropeptides to influence selection of an action series we presented both control and diffuse configurations of the model with a four–action series $1 \rightarrow 2 \rightarrow 3 \rightarrow 4$ of varying input duration and salience. Using a control configuration where the model striatum included no neuropeptide projections, the selection score for each trial remained close to zero until input salience rose to at least 1,350 spikes/s for a duration of 500 ms (**Figure 3.5A**). This rise in selection score corresponded to successful selection of the entire series, which occurred at lower durations as the input salience increased thereby showing that both input features have an impact on selection of an action request and implying that they are to some degree interchangeable. The mean selection score for all action series presentations using the control configuration was 0.3273.

Using the diffuse configuration resulted in an increase of the mean selection score to 0.3788 (**Figure 3.5B**); however, this did not correspond to the series as a whole being reliably selected at lower input salience or duration values. Instead, action requests that were successfully selected using the control configuration were selected for a longer duration, and some action requests that were just below the threshold for selection in the control configuration were able to become selected.

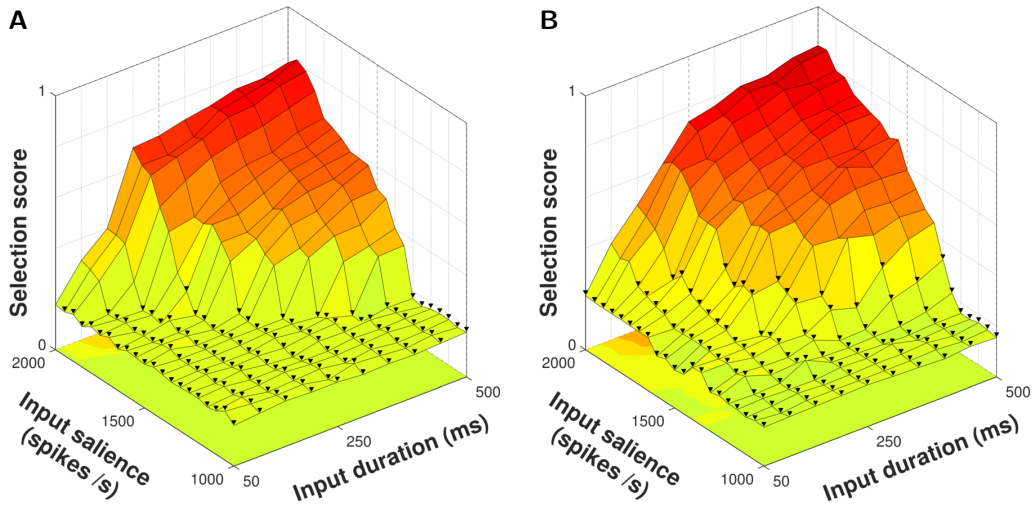

Figure 3.5: Selection scores following presentation of a four–action series to control and diffuse neuropeptide configurations. Trials with fewer than four actions selected indicated with ▼. **(A)** The control configuration shows a clear delineation between trials where all action requests are selected and those where some are not. **(B)** The diffuse configuration facilitates selection but does not enable selection of the entire action series at lower input values.

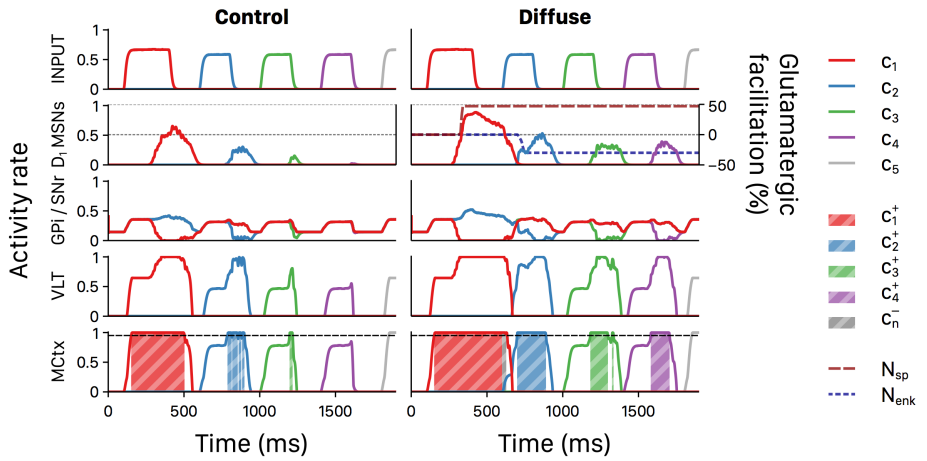

Figure 3.6: Diffuse neuropeptide connectivity facilitates selection of above– or near–threshold action requests. Activity rates of key populations in channels $c_1 \ldots c_5$ following presentation of an action series with input duration 200 ms and input salience 1,700 spikes/s. MCtx selection threshold $\theta = 0.95$ is indicated and selection of channel n resulting in a positive selection score is highlighted as c_n^+. Selection of any channel resulting in a negative score is highlighted as c_n^-. D_1 MSN activity is overlaid with neuropeptide glutamatergic effects N_{sp} and N_{enk}.

Figure 3.6 illustrates this facilitation with a comparison of selected rate outputs following presentation of a single action series with input duration 200 ms and salience 1,700 spikes/s to both configurations. Under the control configuration the selection of channel c_2 was intermittent, channel c_3 was selected very briefly, and channel c_4 was entirely unselected. The selection score for this trial was 0.2491. Using the diffuse configuration, the selection of

channels c_2 and c_3 was enhanced and activity in channel c_4 was raised above the threshold for selection. The extended selection of channel c_1 also resulted in a brief period in which selection was sustained into the onset of channel c_2, and the onset of enkephalin–based inhibition caused a brief dip in striatal activity during the transient burst of activity to channel c_2. As a result of the improved selection under this configuration the score for this trial was 0.4953 and the entire series was successfully selected.

Striatal output under the control configuration in **Figure 3.6** appears to progressively decay in response to subsequent inputs despite lacking the inter-channel enkephalin connections necessary to facilitate this. However, repeat presentations (not shown) revealed this relationship to be illusory, caused by the use of a specific input duration and salience combination that is on the threshold of what is sufficient for selection. Even with identical inputs each channel's response varies slightly due to differences in intra-striatal connectivity and random input spike clustering, and with borderline inputs these variations are sufficient to cause transient selection. The control configuration's responses in **Figure 3.6** happen to occur in order of decreasing selection performance, giving the appearance of a decaying relationship where none exists. This trial is represented by the single raised ▼ in **Figure 3.5A**, showing that partial selection is unique to this trial and that under the control configuration all other input combinations result in either selection or non-selection of all action requests.

Because all D_1 MSNs release SP irrespective of their targets in the diffuse configuration, facilitation of an action request using this configuration arises from both within– and between–channel effects. Once an action request is selected, SP released from that action channel feeds back into the same channel and encourages the sustained selection of that channel, as evidenced by the longer periods of selection for channels $c_1 \ldots c_3$. In addition, selection of any request releases SP that projects to every other channel and facilitates subsequent action requests, shown by the successful selection of channel c_4 and the more rapid selection of channel c_2.

Thus, diffuse neuropeptide connectivity improved the selection of action requests already above threshold and in some cases was able to raise previously subthreshold action requests to a suprathreshold level. With only a few exceptions, however, this did not allow for entire action series to be selected given inputs of lower salience.

3.6.3 Unidirectional Substance P Enhances Action Sequence Selection

The undirected SP projections in the diffuse configuration make it inherently unable to preferentially facilitate one action request over another, and therefore unable to distinguish between ordered and disordered presentations of a semantically ordered action sequence. To investigate the performance of a unidirectional SP configuration in this regard we presented an action sequence in both ordered $(1 \rightarrow 2 \rightarrow 3 \rightarrow 4)$ and disordered $(4 \rightarrow 3 \rightarrow 2 \rightarrow 1)$ states to models using control, diffuse and unidirectional configurations. Because an action sequence is semantically ordered, a higher selection score corresponds to greater selection

of the ordered presentation and greater inhibition of the disordered presentation.

Figure 3.7 shows selection scores for ordered and disordered sequence presentations to control, diffuse, and unidirectional configurations. Because the first two configurations are necessarily sequence–agnostic their response to both presentations is identical to an arbitrarily–ordered series as in **Figure 3.5**; the difference in selection scores between ordered and disordered presentations is thus due to scoring the disordered presentation according to the correct semantic order rather than a change in model performance.

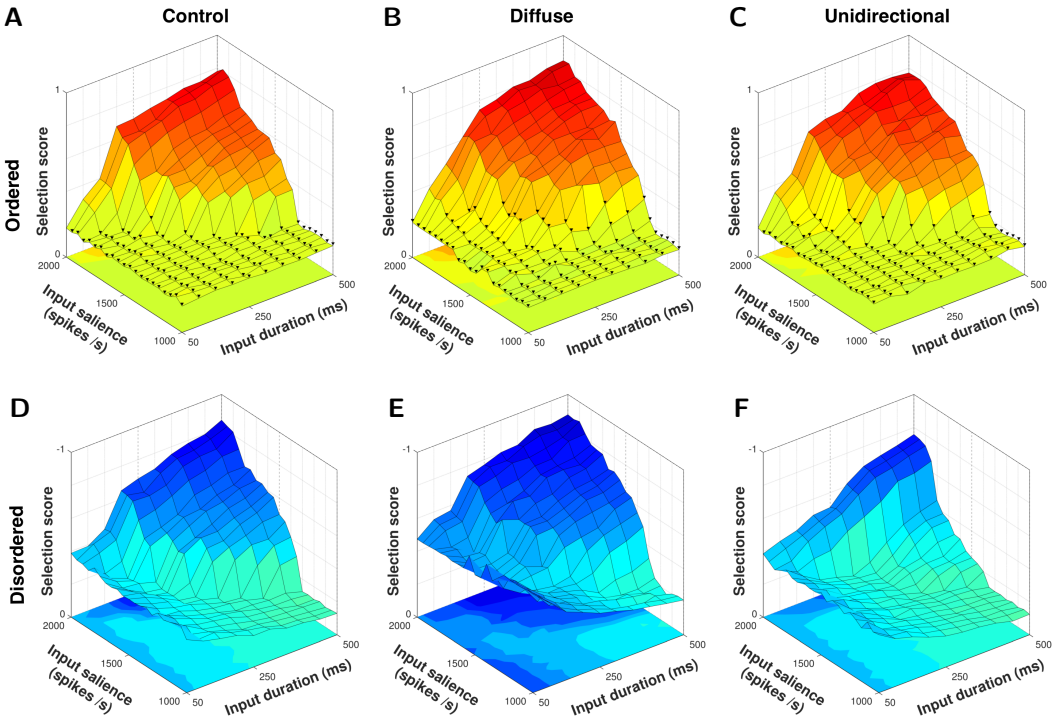

Figure 3.7: Selection scores following ordered and disordered presentations of a four–action sequence to multiple neuropeptide configurations. Ordered presentations with fewer than four actions selected indicated with ▼ and disordered presentations shown with inverted Z–axis. Presentation of an ordered sequence to the diffuse configuration **(B)** results in facilitation of selection compared to control **(A)**. However, because these configurations are sequence–agnostic, disordered presentations to control **(D)** and diffuse **(E)** configurations result in erroneous selection across a wide range of input values. Compared to control, the unidirectional configuration both facilitates selection of the ordered presentation **(C)** and inhibits selection of the disordered presentation **(F)**.

The unidirectional configuration conferred a slight advantage to selection of an ordered sequence presentation over control, with a mean selection score of 0.3542 (**Figure 3.7C**) compared to 0.3273 for control (**Figure 3.7A**). This was lower than the mean selection score of 0.3788 for the diffuse configuration (**Figure 3.7B**), implying that the difference in neuropeptide projection architecture between these configurations impacted selection facilitation. In the unidirectional configuration, only SP projections of the form $c_n \rightarrow c_{n+1}$ are permitted, with the result that facilitation of an action request can only occur if the previous request in the sequence is selected. Thus, within–channel feedback facilitation

cannot occur in the unidirectional configuration, resulting in a lower mean selection score than the diffuse configuration in response to an ordered sequence presentation.

However, the distributed between–channel facilitation present in the diffuse configuration caused erroneous selection in response to disordered presentations (**Figure 3.7E**). The inability to inhibit (or prevent facilitation of) undesired action requests resulted in a mean selection score of -0.5005 for the diffuse configuration compared to -0.3654 for control (**Figure 3.7D**). Conversely, the unidirectional configuration was able to actively inhibit disordered requests, resulting in a higher mean selection score of -0.2913 (**Figure 3.7F**).

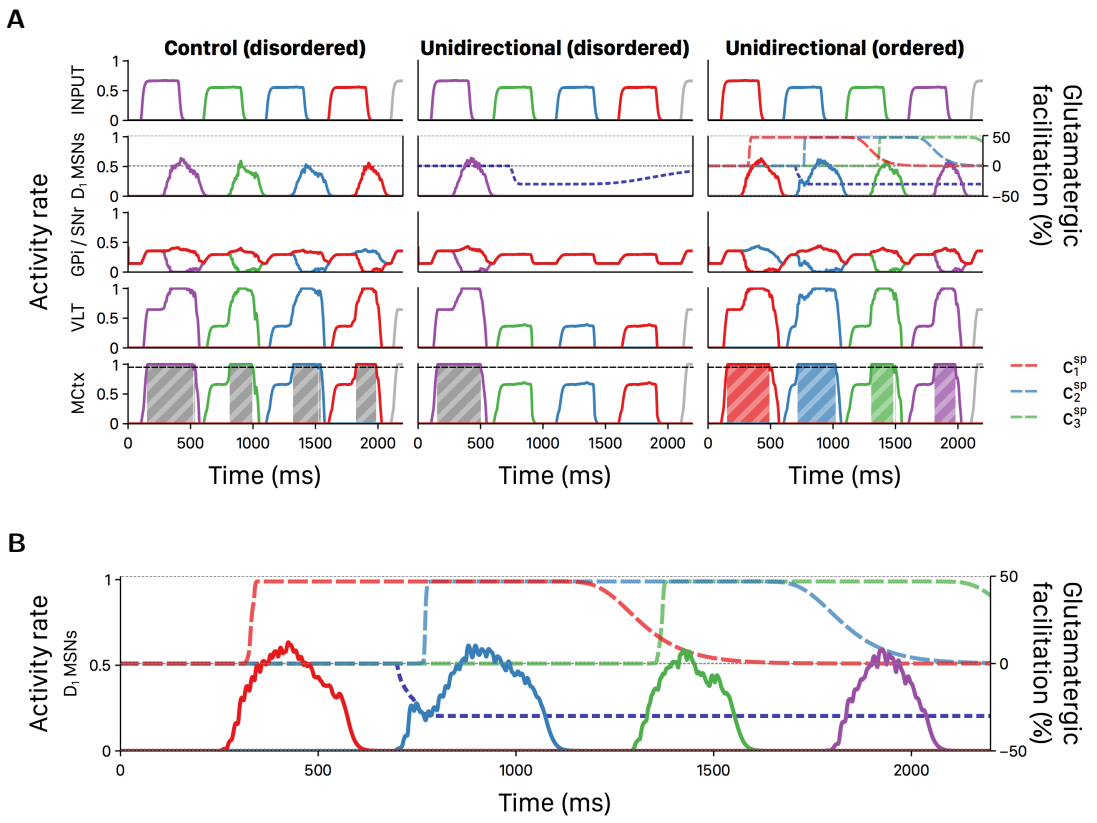

Figure 3.8: Active inhibition of disordered action requests requires enkephalin and is counteracted by substance P. (A) Activity rates of key populations in channels $c_1 \ldots c_5$ following presentation of an ordered or disordered action sequence with input duration 300 ms and input salience 1,600 spikes/s. Interpretation of graphical elements as in **Figure 3.6**, with the addition of channel-specific SP effects c_n^{sp}. **(B)** Enhanced reproduction of D_1 MSN activity in the ordered presentation, showing per-channel SP facilitation and global enkephalin inhibition of inputs.

Figure 3.8A illustrates how between–channel SP and diffuse enkephalin interact in the unidirectional configuration to enhance selection and inhibition of ordered and disordered action requests respectively. In the control configuration, each individual action request was salient enough to become (erroneously) selected without additional facilitation. However, in the disordered presentation to the unidirectional configuration the selection of the first

action request released sufficient enkephalin to inhibit selection of the remaining requests; enkephalin's long timecourse enabled the inhibition of multiple successive action requests in the absence of SP-ergic facilitation. The ordered presentation to the unidirectional configuration shows that the presence of between–channel SP was sufficient to counteract enkephalinergic inhibition and enabled ordered action requests to be selected at least as, or even more strongly than in the control configuration; **Figure 3.8B** shows an expanded view of D_1 MSN activity during the ordered trial to illustrate the period during which selection of subsequent inputs is facilitated.

The presence of distributed enkephalin projections therefore makes the selection of subsequent action requests contingent on the presence of SP; inhibition of other requests becomes a default counteracted by SP influx. Not only was selection of an ordered sequence improved by the unidirectional facilitation of action requests, but the semantic ordering of the sequence was protected by the inhibition of disordered action requests.

3.6.4 Substance P Pruning Enhances Action Clique Separation

We have shown that the diffuse configuration confers the strongest advantage to selection of unordered action requests by utilising both within– and between–channel SP projections to facilitate action requests, and that the unidirectional configuration allows active inhibition of disordered requests by limiting facilitation to ordered between–channel projections. We sought to explore if selective inhibition of undesired requests could allow for further structuring of action groups. We therefore presented all four model configurations with action requests in the order $1 \rightarrow 6 \rightarrow 2 \rightarrow 3 \rightarrow 4$, where action channels $c_1 \ldots c_4$ formed an action clique and channel c_6 represented an undesired distractor action.

In exploring the model's response to an action clique, the input duration and salience of each action request in the clique was fixed at 300 ms and 1,600 spikes/s respectively, and only the salience and duration of the non-clique distractor was varied. **Figure 3.9** shows clique and non-clique selection scores for each model configuration.

In line with results from the disordered sequence presentation (**Figure 3.7F**), the unidirectional configuration was able to inhibit the distractor request (**Figure 3.9F**) better than the control condition (**Figure 3.9E**) due to the presence of diffuse enkephalin that was not counteracted by a directed SP projection. Mean distractor selection score for the unidirectional configuration was -0.0986, compared to -0.2090 for control. However, the ability of the unidirectional configuration to facilitate clique action requests dropped to levels below that of control as the duration of the distractor rose above $\sim$250 ms (**Figures 3.9A,B**). Mean clique selection score for the unidirectional configuration was 0.1919, compared to 0.2943 for control. Closer examination of the neuropeptide dynamics (not shown) revealed that the longer timecourse of enkephalin coupled with the lack of within–channel feedback facilitation resulted in between–channel SP facilitation decaying before the onset of channel c_2, resulting in effective inhibition of the remaining clique action requests.

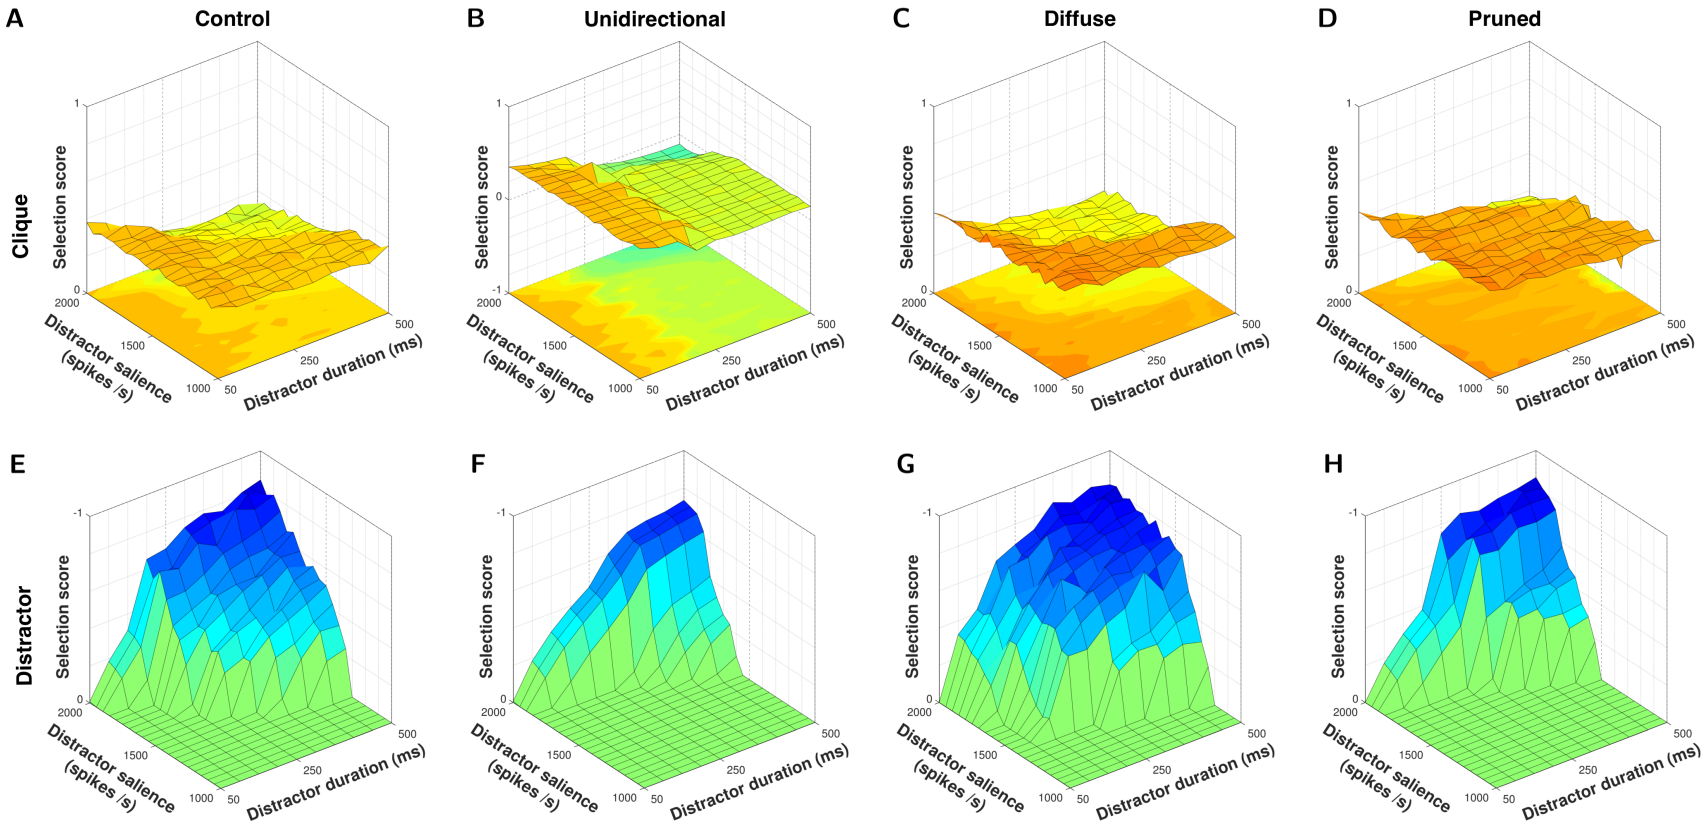

Figure 3.9: Selection scores for clique and distractor actions for all model configurations. Distractor scores shown with inverted Z–axis. The control configuration is moderately able to select clique action requests **(A)** but is unable to inhibit selection of the distractor **(E)**. The unidirectional configuration is able to actively inhibit the distractor **(F)** but is less able to select clique action requests when distractor duration is high **(B)**. The diffuse configuration facilitates selection of clique action requests **(C)** but also facilitates selection of the distractor **(G)**. The pruned configuration is able to selectively inhibit the distractor **(H)** and facilitate selection of the clique across a wider range of distractor input values **(D)**.

It should also be noted that the presented results represent the action clique scenario most favourable to the unidirectional configuration, as the clique actions are presented in an order that matches the unidirectional configuration's semantic preference. Other valid action clique presentations would result in the inhibition of some or all actions and a much poorer mean selection score.

In line with results from the action series presentation (**Figure 3.5B**), the diffuse configuration was able to select clique action requests better than control (**Figure 3.9C**) due to the combination of within– and between–channel facilitation, resulting in a mean selection score of 0.3415. However, the diffuse SP projections caused equivalent facilitation of the distractor request (**Figure 3.9G**) resulting in a mean distractor selection score of -0.3112 and thus no effective separation of clique from non-clique requests.

The pruned configuration was the only neuropeptide projection scheme that performed better than control at both clique selection and distractor inhibition (**Figures 3.9D,H**). Mean clique selection score for the pruned configuration was 0.3718, and mean distractor selection score was -0.1314. The removal of SP projections from channel $c_1 \rightarrow c_6$ inhibited the non-clique action request and allowed within–channel SP feedback in channel c_1 to sustain selection of that channel until the delayed onset of channel c_2, resulting in strong selection of the clique even with a high duration distractor. The inclusion of feedback SP projections within channel c_6 allowed for strong selection of the distractor request when it was salient enough to overcome the diffuse inhibition, therefore resulting in a greater contrast between inhibition and selection of non-clique action requests.

Selective removal of channel–specific SP projections therefore allowed for inhibition of individual action requests that enhanced both the separation of distinct action cliques and facilitation of selection within a given clique. Furthermore, because this was dependent on the removal of specific SP projections, the resulting inhibition is context–sensitive and potentially allows for the inhibition of distinct distractors in different situations with only minimal changes to the SP network.

3.6.5 Impact of Gap Duration on Selection Performance

All the results described so far have emerged from manipulations to the duration, salience, and ordering of inputs to channels $c_2 \ldots c_4$ without varying the inter-channel 'gap' period. As sustained selection throughout this period is a key requirement for good model performance we sought to explore the impact of altering the gap duration.

We provided control, diffuse, and unidirectional configurations with repeat presentations of an input series with duration 300 ms and salience 1,600 spikes/s while modifying only the gap duration. **Figure 3.10** shows the selection scores for each configuration for gap durations ranging from 0 ms to 500 ms.

All three configurations exhibit a similar overall pattern in their response; selection scores

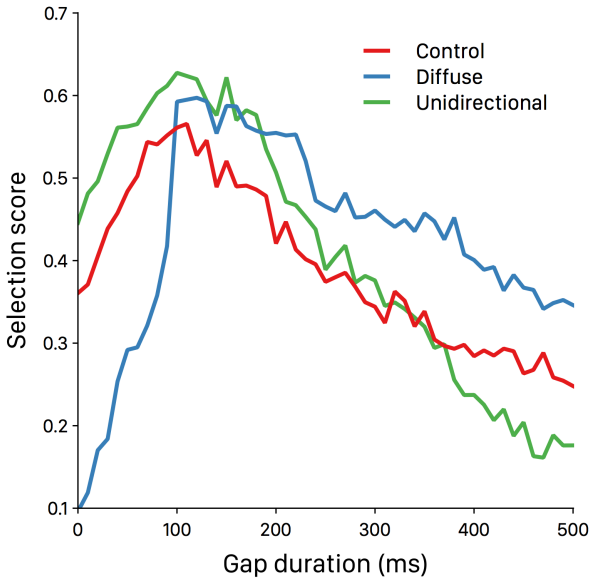

Figure 3.10: Selection scores for multiple neuropeptide configurations for gap durations up to 500 ms. Presentation of an action series with input duration 300 ms, salience 1,600 spikes/s, and varying gap duration demonstrates the relative strengths and weaknesses of within– and between–channel SP facilitation.

rise as the gap duration increases from 0 ms to ∼100 ms, peak at approximately 100 ms, and then decline with further increases to the gap duration. However, the diffuse configuration performs significantly worse than both control and unidirectional configurations at very low gap durations, but rapidly improves with increased gap duration to reach its peak performance along with both other configurations at ∼100 ms. Unlike the control and unidirectional configurations, the diffuse configuration is able to maintain near–peak performance at gap durations up to ∼220 ms, before declining with increasing gap durations at approximately the same rate as the control configuration.

Conversely, performance under the unidirectional configuration is higher than both other configurations at gap durations under 100 ms, but declines more rapidly as gap durations increase past ∼150 ms. This highlights a fundamental tradeoff between sustained selection and rapid transitions within action series; at very low gap durations there is little need to sustain selection outside of the phasic input period, and within–channel SP feedback in this scenario is detrimental as it promotes sustained selection of the current channel over rapid transition to the next channel. Furthermore, at such low gap durations selection transitions are inhibited not just by enkephalinergic but also GABAergic inhibition, and it is this latter source of inhibition that acts on the control configuration to impede its performance at gap durations below 100 ms.

However, at progressively higher gap durations the need for sustained selection during the gap period becomes increasingly important, and the lack of within–channel SP to counteract enkephalinergic inhibition causes the unidirectional configuration's performance

to rapidly fall below that of control. For gap durations between ∼100 ms and ∼190 ms diffuse and unidirectional configurations perform nearly identically, suggesting that the competing requirements of sustained selection and rapid transitions are well balanced with these gap values. Performance in all configurations decays with increasing gap duration, but the diffuse configuration consistently outperforms control at gap durations up to at least 500 ms.

We may therefore remain confident that the results pertaining to diffuse and pruned configurations described in Sections 3.6.2 and 3.6.4 respectively will hold for gap durations above 200 ms up to at least 500 ms. However, the results described in Section 3.6.3 are more reliant on a narrow range of gap values, and at gap durations above ∼300 ms unidirectional SP may no longer be able to counteract the inhibitory effects of enkephalin. Additionally, the dramatically altered performance of all configurations under 100 ms suggests that another mechanism may be required to explain sequence transitions with very low gap durations.

3.7 Discussion

We have integrated our novel neuropeptide models into a hybrid basal ganglia model based on previous work (Humphries and Gurney, 2002; Humphries et al., 2010). By exploring several neuropeptide connectivity configurations and presenting the model with groups of action requests we have shown that inclusion of these neuropeptides can improve the model's ability to both select and reject action requests, and that pruning the SP network can improve the model's ability to selectively inhibit actions that are disordered or undesired. In addition, we have shown that the interaction of SP and enkephalin is of key importance to the appropriate facilitation and inhibition of action requests.

3.7.1 Neuropeptide Action and Interaction

Presenting the model with multiple groups of action requests allowed us to examine the effects of SP and enkephalin within and between MSN populations representing each channel. Presentation of an action series with no preferred semantic order revealed that SP facilitates action requests in two distinct ways. Firstly, SP acts within an action channel to promote its continued selection by facilitating glutamatergic inputs from MCtx regions representing the currently active channel. This leads to sustained inhibition of GPi/SNr activity and thereby promotes continued activity in the MCtx–VLT feedback loop representing the current channel. Secondly, SP acts between action channels by facilitating glutamatergic inputs from cortical regions representing other action requests, raising their effective salience and potentially enabling the selection of otherwise subthreshold requests (**Figure 3.6**). The combination of these two effects allows the diffuse release of striatal neuropeptides to facilitate selection of action requests within an unordered series.

The presentation of a disordered action sequence showed that this indiscriminate facilitation can be a detriment to the inhibition of undesired action requests (**Figure 3.7E**). When SP release is restricted to MSN connections targeting MSNs representing the next sequential action, enkephalin can actively inhibit disordered requests and prevent the erroneous selection of a disordered sequence (**Figure 3.7F**). This enhanced inhibitory ability is not present in the absence of enkephalin (**Figure 3.7D**).

We also showed that between–channel SP facilitation is sufficient to counteract this inhibition and allow the selection of a correctly–ordered sequence with greater strength than if no neuropeptides were present (**Figure 3.8**). Therefore, a transition between two actions A and B within the active timecourse of enkephalin is contingent on either the presence of an SP projection $A \rightarrow B$ or an extremely salient request for action B.

Presentation of an action clique with an additional distractor request showed that by selectively pruning SP connections, the ability of the model to both select an action clique and inhibit distractions is improved with minimal changes to the SP network (**Figure 3.9**). Although diffusely–connected SP and enkephalin facilitate action requests and raise their effective salience above control levels (**Figure 3.6**), the more compelling result may be the inhibition of otherwise salient action requests that occurs when SP projections between MSN populations are removed (**Figure 3.8A**).

In summary, diffuse enkephalin release causes broad inhibition of incoming action requests counteracted by targeted SP release that facilitates the initial and sustained selection of specific requests. This agrees with results from serial selection tasks in rats; increased activation of D_1 receptors reduced selection accuracy and increased premature responses (Agnoli et al., 2013), which could result from enhanced firing of D_1 MSNs and concomitant greater SP release. Similarly, increased activation of D_2 receptors increased both premature and perseverative responses (Agnoli et al., 2013) which could result from reduced firing of D_2 MSNs and accordingly lower enkephalin release. Equivalent but opposing effects were observed with DA receptor antagonists (Domenger and Schwarting, 2006).

Neuropeptide projections between MSN populations may support striatal connectivity supporting ordered action sequences or distinct action cliques that also enhances the contrast between selected and inhibited requests, similar to the 'unsharp mask' model of striatal processing proposed by Stocco and Lebiere (2014). In their model, lateral inhibitory connections within MSN populations form a 'blurred' representation of cortical input signals that is subtracted from the original signal to form a dynamically enhanced output representation. Neuropeptides may complement this process by enabling MSNs to dynamically enhance input signals as well.

3.7.2 Substance P and Learning

The ability of the unidirectional and pruned configurations to selectively inhibit an otherwise salient action request in favour of an alternative suggests that sequence learning and habituation may involve plasticity of striatal SP connections. Plasticity need not directly promote the selection of related action requests but could still raise their relative salience by reducing the quantity or strength of SP connections targeting MSNs representing potential distractions, thereby inhibiting undesired requests.

Though we are as yet unaware of any direct evidence for SP plasticity within the striatum, SP has been previously implicated in the facilitation of learning (Hasenöhrl et al., 2000; Huston and Hasenöhrl, 1995) and in affecting motivational aspects of reward (Murtra et al., 2000). Skill learning has also been shown to result in a relative increase in D_2 MSN activity (Yin et al., 2009) that could plausibly correlate with refinement of SP connectivity. Furthermore, a pattern of preferentially within–channel and within–sequence SP projections combined with diffuse enkephalin projections is consistent with closed– and open–loop reverberations resulting from stimulation of D_1 or D_2 dopamine receptors respectively as reported by Carrillo-Reid et al. (2011).

We have explored only a basic example of selective inhibition; as it is context–sensitive and channel–specific, more complex modifications could allow for groups of MSNs to represent multiple overlapping cliques simultaneously. Patterned SP connectivity may allow the striatum to incorporate probabilistic links between component actions that could represent a neurological basis for chunking (Graybiel, 1998) and higher–order hierarchical groups (Balleine et al., 2015).

MSN populations thus delineated by patterned SP connections may act as "local controllers" (Graybiel and Grafton, 2015) of sequence chunks that facilitate the efficient selection of ordered cortical inputs; indeed, striatal activity during motor chunking has been shown to correspond to the concatenation of sequence elements while cortical activity correlates with their segmentation (Wymbs et al., 2012). It would therefore not be correct to say that sequences are 'stored' within the striatum. While the striatum is heavily involved in the chunking of cortical inputs (Graybiel, 2008; Jin et al., 2014), it also relies on cortical inputs to provide additional information about the order, timing, duration, and salience of individual actions. Sequence execution is therefore a result of the co-ordinated interaction of both structures.

We may posit a role for SP as a modulator of dopamine in the formation of sequence chunks. It has recently been reported that SP may modulate dopamine transmission differently according to neuronal location within the striosomal–matrix axis (Brimblecombe and Cragg, 2015) and movement chunking itself has been shown to be dependent on dopamine (Tremblay et al., 2010). The question of whether SP projections in striatum are plastic as part of sequence learning — and if so, how this plasticity is mediated — should prove enlightening.

3.7.3 Clinical Implications

The impact of striatal neuropeptides on action selection has several potential clinical implications. Huntington's disease causes the degeneration of both D_1 and D_2 MSNs (Turjanski et al., 1995; Weeks et al., 1996) but preferentially impacts D_2–expressing neurons (Mitchell et al., 1999; Richfield et al., 1995). Huntington's disease specifically impairs the learning of motor sequences (Willingham and Koroshetz, 1993), and one plausible explanation is that an enkephalin deficit resulting from the degeneration of D_2 MSNs impairs the ability to selectively facilitate sequential action requests and thus to integrate semantic relationships into the structure of striatum.

Huntington's disease also has well-established impacts on cognitive function that often manifest long before the first indications of motor problems (Duff et al., 2010; Paulsen et al., 2001). Because the basal ganglia also process signals from limbic and associative functional territories, we may reasonably speculate that striatal neuropeptides play a similar selective role across these cognitive domains. Enkephalinergic degeneration damaging the ability to inhibit undesired action requests could therefore also potentially explain the increase in impulsivity and risk–taking behaviours seen in Huntington's disease (El Massioui et al., 2016; Kalkhoven et al., 2014).

Conversely, the increased apathy seen in Huntington's disease seems contrary to expectations from a loss of inhibitory signalling. Heightened apathy is correlated to decline in both cognitive and motor function (Baudic et al., 2006; Naarding et al., 2009; Thompson et al., 2002), and it has been suggested that this apathy is a manifestation of an overall reduction in "drive and motivation" (Hamilton et al., 2003) resulting from these problems. Some preliminary success in slowing cognitive decline in Huntington's disease has been seen following cholinergic interventions (de Tommaso et al., 2004; Morton et al., 2005), suggesting that the role of striatal interneurons and neurotransmitters absent from the current model may be important for a fuller understanding of this condition.

The ADHD–like symptoms of mice lacking functional NK_1 receptors (Porter et al., 2015; Yan et al., 2009) may also be partially explained by a lack of within–channel facilitation to sustain selection of a given action. An inability to sustain selection without external stimuli could cause changes in corticostriatal activity that may reflect "an increased need for or reliance on vigilance or sustained visual attention" reported in children with ADHD (Durston et al., 2003). Definitive conclusions about the mechanisms underlying specific pathologies is beyond the scope of our research, however.

The interaction of neuropeptides clearly allows for an additional layer of computational complexity within the striatum. The facilitatory effects of SP imply an unexplored co-operative role for MSNs that warrants further investigation, especially in regards to the development of patterned SP connectivity. Exploring the computational role of omitted striatal interneurons should also prove illuminating, as SP has been found to modulate the

activity of cholinergic interneurons (Govindaiah et al., 2010) and the responses of tonically active interneurons are strongly correlated with the likelihood of a behavioural response to a stimulus (Blazquez et al., 2002). Future additions to models of neuropeptide interactions will doubtless reveal new striatal functions.

Chapter 4

Striatal Topography and Action Selection

4.1 Introduction

In Chapter 3 we embedded a striatal microcircuit model of 6,000 MSNs in a rate-coded basal ganglia–thalamocortical loop model to probe the effects of striatal neuropeptides on action selection. The relatively low number of neurons in the striatal model meant that while each neuron received a statistically appropriate number of connections, they had a biologically unrealistic distribution and the connectivity had no topography. The impact of this on the modelling results is uncertain.

In this chapter we enhance the biological plausibility of the striatal microcircuit model by expanding it to 84,900 MSNs and placing the neurons in a virtual space representing 1 mm^3 of tissue. Biologically grounded connectivity based on estimations of dendritic and axonal arborisation is established using an algorithm developed by Humphries et al. (2010) and we perform input manipulation based on the network's new physical characteristics.

We integrate the expanded striatal microcircuit model with the BG loop model described in Section 3.3 and attempt to verify the primary results from Section 3.6, before exploring differences introduced with the expanded model.

4.2 Model Construction

Though most aspects of the striatal microcircuit model remain identical to Section 3.2, changes to neuron connectivity must be addressed. To distinguish the new hybrid model from its predecessor, we will henceforth refer to the model utilising the 84,900 MSN microcircuit with physical topography as the *physical* model, and the model utilising the 6,000 MSN microcircuit with statistically–derived connectivity as the *statistical* model.

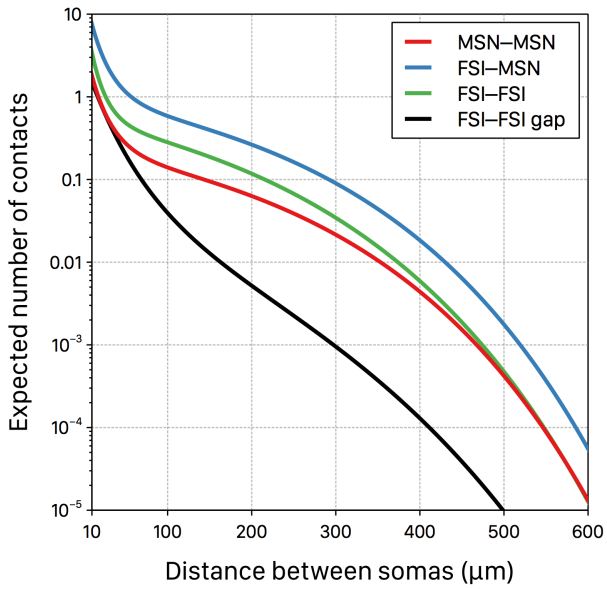

Figure 4.1: Expected number of contacts between neurons in the physical model as a function of distance. Adapted from Humphries et al. (2010).

4.2.1 Neuron Populations and Connectivity

In Section 3.2 we described the construction of the statistical striatal microcircuit model, in which 6,000 MSNs and 60 FSIs are connected at random until the number of neural connections is equivalent to what would be expected given a complete dendritic arbor simulation (Humphries et al., 2010). The physical microcircuit model builds on this by placing each neuron at specific co-ordinates in a three-dimensional virtual space representing 1 mm^3 of striatal tissue and increasing the number of neurons to match a biologically plausible neuronal density.

Per Oorschot (1996) we place 84,900 MSNs in our virtual striatal region, assigning each one random co-ordinates with the stipulation that a minimum distance of 10 μm is maintained between any two neurons (Humphries et al., 2010). Any neuron assigned co-ordinates within 10 μm of another is given new co-ordinates until this minimum distance is met.

For every MSN that is placed there is a 1% chance of generating and placing an FSI, thus maintaining an appropriate MSN:FSI ratio (Humphries et al., 2010; Luk and Sadikot, 2001). FSIs are subject to the same minimum distance requirements as MSNs.

Once all neurons are placed, connections between neurons are generated. We calculate the expected number of synaptic contacts for every pair of neurons based on their type and distance (**Figure 4.1**) following the algorithm from Humphries et al. (2010), and the total number of each type of connection is verified to ensure it is within expected bounds.

Simulating a virtual 1 mm^3 cube of striatal tissue thus allows for more realistic connectivity, albeit with some compromises. Most significantly, neurons towards the edge of the virtual

space will necessarily connect to fewer partners than those towards the centre, potentially resulting in 'edge effects' and differentiated behaviour as the result of this diminished innervation. However, the large size of our simulated region should minimise the impact of any such effects on output behaviour.

Once all intra-striatal connections are made, MSNs are split into two equal populations of 42,450 D_1 and D_2 neurons and assigned to an action channel. For validation of the statistical model (Section 4.3.1) MSNs are further subdivided at random into six equal populations representing action channels $c_1 \ldots c_6$, giving 7,075 neurons of each type in each channel. For the input segregation experiments in Sections 4.3.2 and 4.3.3, MSNs are divided into two populations representing action channels c_1 and c_2 with 10% of the total population left as background neurons receiving no input.

We define two subtypes of the two-channel physical model:

Uniform

D_1 and D_2 MSNs are each split at random into two equal populations representing c_1 and c_2. The physical location of an MSN is unrelated to its assigned channel. (MSNs in the six-channel physical model are assigned at random and are thus equivalent to the uniform model.)

Segregated

D_1 and D_2 MSNs are split into two populations representing c_1 and c_2 based on their x–axis co-ordinate; neurons where $x < 500$ μm are assigned to c_1 and neurons where $x \geq 500$ μm are assigned to c_2. As channel assignation is based on physical location, channels may have a slightly unequal number of neurons.

4.2.2 Inputs and Spike / Rate Conversion

Inputs to the striatal microcircuit and BG loop models follow the scheme detailed in Section 3.4.3 with an increased number of Poisson spike generators to account for the greater number of neurons. As with the statistical model, each MSN receives a single Poisson spike input from its assigned channel and each FSI receives a single Poisson spike input from every channel. An additional 'background' channel additionally provides background noise to each MSN and FSI.

Rate–to–spike conversion is unchanged from Section 3.4.1, and spike–to–rate conversion has only minor modifications. As spiking inputs to the BG loop model are summed before conversion to rate output, we take only the first 500 spike trains for each connection to be converted. This is due to technical limitations on the number of spikes that SpineCreator may receive at any population in a single timestep, and also to avoid retuning the spike–to–rate conversion for each population. The physical model therefore implicitly assumes that the first 500 striatal spike trains in each action channel are representative of the activity

of the entire channel, and that no intra-channel variation in activity exists that would impact the downstream performance of the model. The spiking outputs of all MSNs are still recorded and available for analysis, however.

4.3 Simulation Results

4.3.1 Striatal Topography Does Not Affect Sequence Selection

To explore the impact of physical striatal topography on the neuropeptide modulation results we repeated several key simulations from Section 3.6. Due to the number of neurons in the physical model and the compute times involved in each trial it was not possible to recreate the surface plots showing all combinations of input salience and duration, so we instead highlight specific examples of each case.

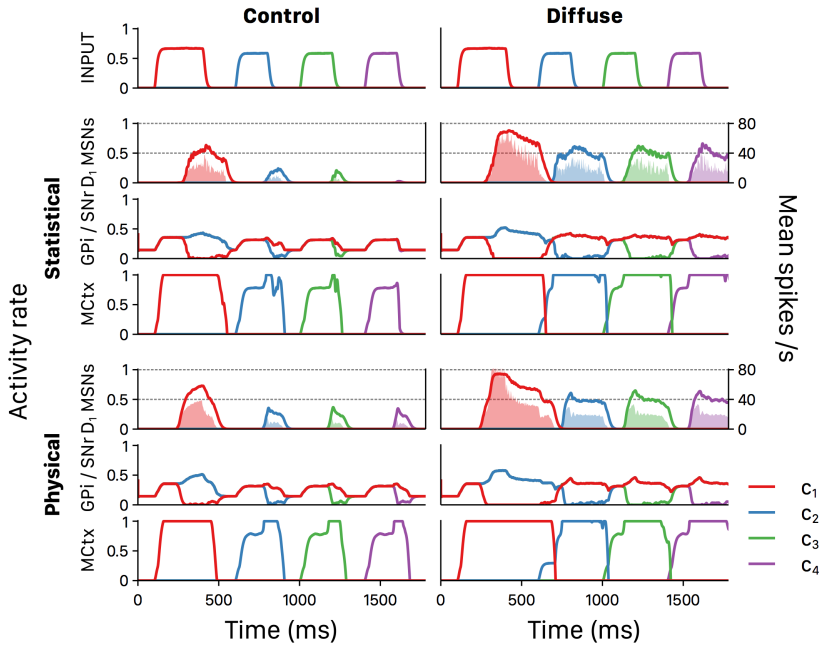

Figure 4.2: Diffuse neuropeptide connectivity in both statistical and physical models facilitates the selection of above– or near–threshold action requests. Activity rates of key populations in channels $c_1 \ldots c_4$ following presentation of an action series with input duration 200 ms and input salience 1,700 spikes/s. Mean spikes/s for striatal populations averaged over a 15 ms rolling window.

Figure 4.2 shows a comparison of rate outputs for selected populations from statistical and physical models following presentation of an unordered action series with input duration 200 ms and input salience 1,700 spikes/s. Using the control configuration in the statistical model, input to channels $c_2 \ldots c_4$ is just below the threshold for sustained selection; channels c_2 and c_3 are selected only briefly and channel c_4 is entirely unselected. In the physical model, the same level of input elicits a slightly stronger striatal response and thus motor

cortex activity in these channels is just able to reach levels required for selection.

Under the diffuse configuration both models are able to strongly select the entire action series, though the increased striatal response in the physical model results in MSN spiking activity in c_1 briefly reaching biologically implausible levels of 80 spikes/s and selection of c_1 extending past the valid period into the onset of c_2. The only other notable difference is the decreased variability in MSN spiking output in the physical model; this is addressed in Section 4.3.2. Aside from these minor variations the main result of Section 3.6.2 is accurately recreated in the physical model, and the presence of diffuse SP and enkephalin facilitates the selection of an unordered action series in both cases.

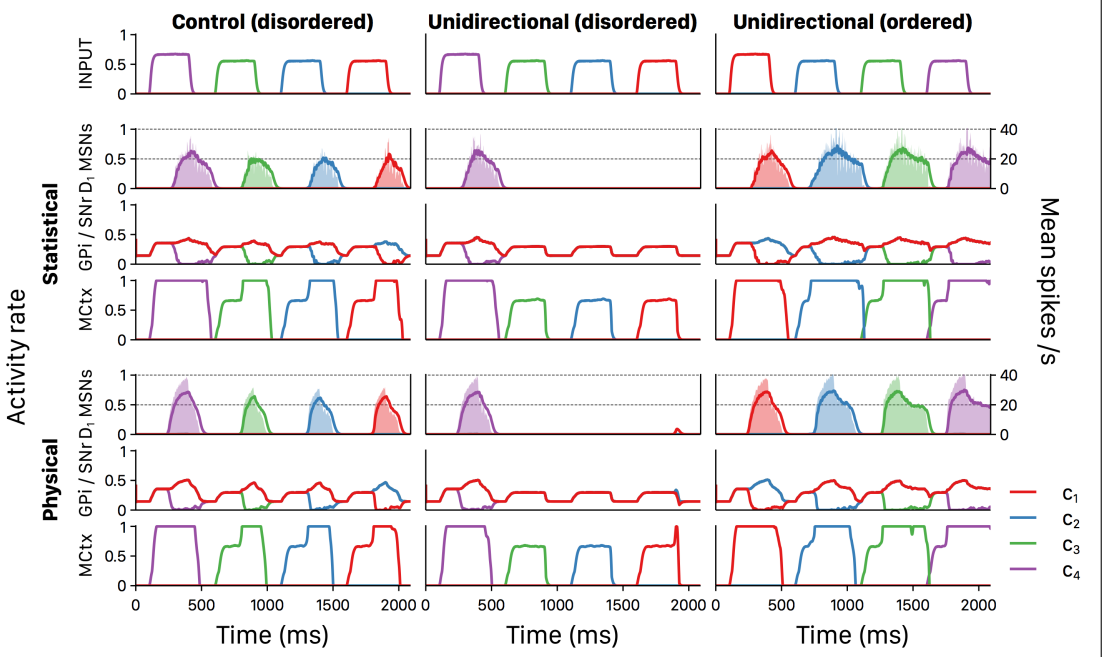

Figure 4.3: Enkephalin–mediated inhibition of disordered action requests is counteracted by targeted substance P in both statistical and physical models. Activity rates of key populations in channels $c_1 \ldots c_4$ following presentation of an ordered or disordered action sequence with input duration 300 ms and input salience 1,600 spikes/s. Striatal mean spikes/s averaged over a 15 ms rolling window.

Figure 4.3 shows a comparison of rate outputs from the statistical and physical models in response to disordered and ordered sequence presentations with input duration 300 ms and input salience 1,600 spikes/s. The physical model exhibits an almost exact recreation of the results from the statistical model, including selection of all disordered channels in the control configuration, inhibition of disordered channels $c_2 \ldots c_4$ in the unidirectional configuration (with the exception of an extremely brief transient spike in MCtx activity at the very end of c_4), and enhanced selection of all ordered channels in the unidirectional configuration. The physical model again shows a slightly enhanced response over the statistical model in response to the same inputs and reduced variability in MSN spiking. The presence of unidirectional SP and diffuse enkephalin therefore allows both the statistical and physical

models to enhance the selection of ordered requests and inhibit disordered actions in a sequence.

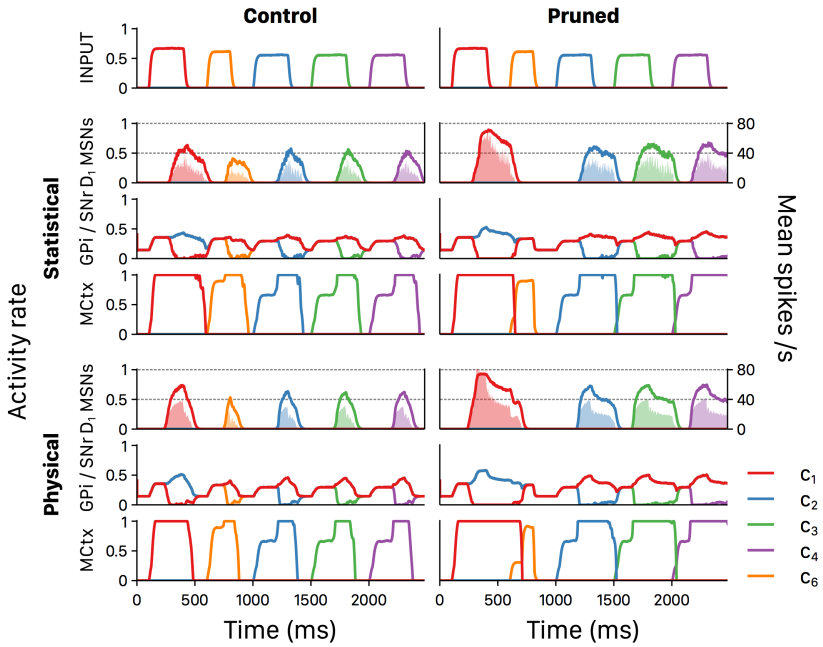

Figure 4.4: Selective substance P pruning enables targeted inhibition of a distracting action request in both statistical and physical models. Activity rates of key populations in channels $c_1 \ldots c_4$, c_6 following presentation of an action clique with input duration 300 ms and input salience 1,600 spikes/s, interrupted by a distractor with input duration 200 ms and input salience 1,800 spikes/s. Striatal mean spikes/s averaged over a 15 ms rolling window.

Figure 4.4 compares rate outputs in the statistical and physical models in response to presentation of an action clique with input duration 300 ms and input salience 1,600 spikes/s that is interrupted by a distractor with input duration 200 ms and input salience 1,800 spikes/s. Both models again exhibit a nearly identical response under both neuropeptide configurations; the high-salience distractor is selected by both models under the control configuration but is successfully inhibited in striatum (and thus kept just below the MCtx selection threshold) under the pruned configuration. Following distractor offset, both models are able to select channels $c_2 \ldots c_4$ for a longer period of time than under the control configuration, resulting in enhanced selection of the action clique even after inhibition of the distractor. Thus, patterned SP connectivity allows for selective inhibition of specific distractor actions in both the statistical and physical models.

4.3.2 Microcircuit Activation Causes Oscillatory Spiking

To explore the impact of physically segregating striatal action channels, we recreated the two-channel transient experiment from Section 3.6.1 and assigned MSNs to each action channel at random (uniform model) or according to physical location (segregated model).

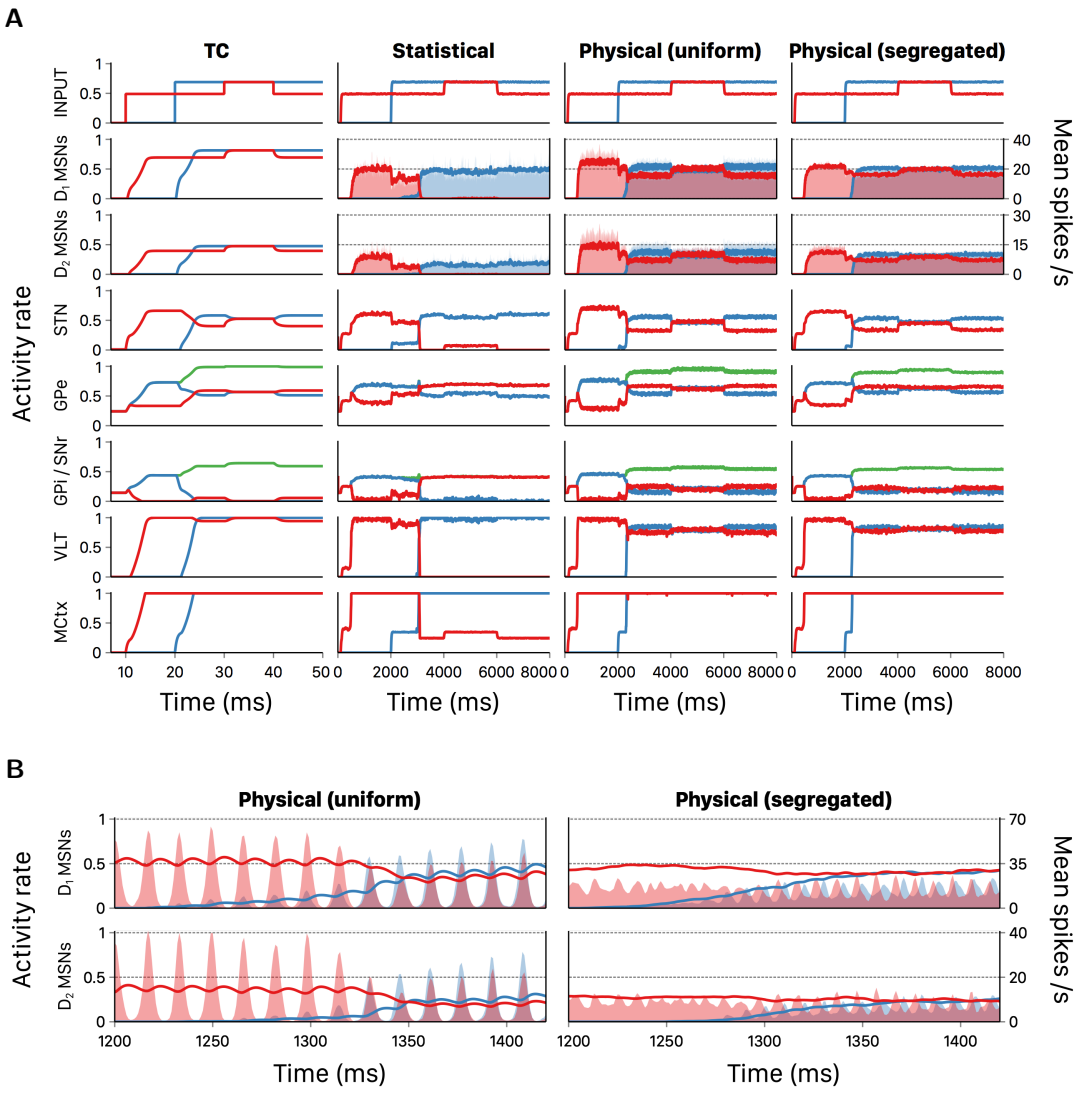

Figure 4.5: Comparison of the TC (Humphries and Gurney, 2002), statistical, and uniform and segregated physical models in response to the transient input experiment. (A) Comparison of activity rates in channels c_1 (red), c_2 (blue), and c_3 (green) in all versions of the model show similar behaviour, though the physical models allow simultaneous activation of multiple striatal action channels. Lines indicate population rate output and shaded regions indicate striatal mean spikes/s averaged over a 15 ms rolling window. **(B)** Expanded view of striatal populations in the physical model showing that input channel segregation reduces oscillatory spiking. Lines indicate population rate output and shaded regions indicate mean spikes/s averaged over a 2 ms rolling window.

Figure 4.5A shows a comparison of the transient selection experiment performed on the Humphries and Gurney (2002) TC model, statistical model, and both physical models. As already observed, inputs to the physical model elicit a slightly stronger striatal response. The second principal difference between statistical and physical models is that the statistical model only permits activation of a single action channel in striatum at any time while

both versions of the physical model show output more similar to the TC model in which both action channels are active in striatum simultaneously. Further testing (not shown) suggests that this is not merely due to increased effective input salience in the physical model; increasing input salience in the statistical model does not enable simultaneous selection of both channels, and decreasing input salience in the physical model still results in simultaneous selection. We may speculate that this is due to increased efficacy of inhibitory FSI–MSN and MSN–MSN connections in the statistical model (see Section 4.4.1), but a definitive answer requires further modelling work to examine this phenomenon in greater detail.

However, there are no unexpected differences in the rate outputs of any population in the segregated or uniform models; the primary difference between the two physical models is in the spiking activity of the striatal MSNs. **Figure 4.5B** shows an enhanced view of the differences in spiking behaviour between the uniform and segregated models using a 2 ms rolling mean; MSN spiking in the uniform model exhibits strong oscillations that are substantially reduced in the segregated model.

FSIs exert a significant inhibitory influence over MSNs (Koós and Tepper, 1999) and may mediate oscillatory activity (Sharott et al., 2009), so we endeavoured to explore this interneuron's role in the emergence of this apparently spontaneous oscillatory behaviour. To that end, we created two additional variants of the physical models; in the *no–FSI* model all FSI → MSN connections were lesioned to eliminate their influence in the model, and in the *2% FSI* model FSIs were generated alongside MSNs with a 2% probability, resulting in approximately double the standard FSI population to increase their influence.

Figure 4.6 shows oscillation power spectra in response to stimulation of a single input channel in uniform and segregated versions of models with all three FSI densities; power spectra were analysed from striatal spike trains in the Chronux MatLab package using multi-taper analysis (Bokil et al., 2010). **Figure 4.6A** shows power spectra for D_1 MSN activity in the 1% FSI uniform model in response to activation of a single action channel at varying saliences. All tested input levels generate strong γ-band oscillations at ∼60 Hz, confirming that activation of a single action channel is sufficient to generate oscillations. With no neuropeptides presesnt (filled curves) oscillation amplitude increased with input salience up to approximately 1,700 spikes/s while oscillation frequency remained fixed at ∼60 Hz, indicating that oscillations were not the result of entrainment to an external source. Including diffuse SP and enkephalin (hollow curves) increased the oscillation frequency somewhat to ∼68 Hz as well as greatly increasing their power.

Figure 4.6B *(upper panel)* shows power spectra for D_1 MSN activity in the uniform model with zero, 1%, or 2% FSIs in response to input of 1,500 spikes/s to a single action channel. Total lesioning of FSIs resulted in extremely strong ∼60 Hz oscillations, while the 1% FSI model shows strong but reduced oscillations at the same frequency. Oscillatory firing appears entirely eliminated in the 2% FSI model, suggesting that increased FSI density

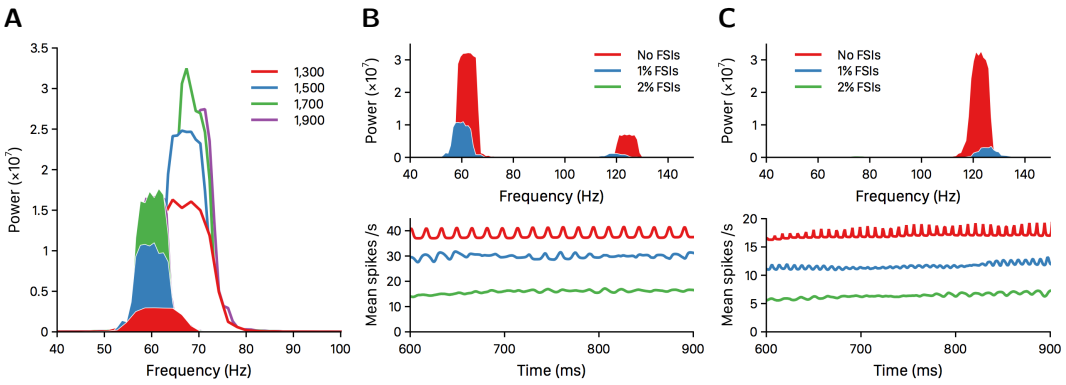

Figure 4.6: Oscillation power spectra and mean firing rate of D_1 MSNs in the uniform and segregated physical models. (A) Power spectra in the 1% FSI uniform model in response to 1 s of input to a single channel at 1,300–1,900 spikes/s with (hollow) and without (filled) diffuse neuropeptides show that input salience affects the amplitude but not the frequency of oscillations. Diffuse neuropeptides increase the frequency and amplitude of oscillations. Power spectra and mean spikes/s for **(B)** uniform and **(C)** segregated physical models without neuropeptides and with no FSIs (red), 1% FSIs (blue), and 2% FSIs (green) in response to 1 s of input to a single channel at 1,500 spikes/s shows active MSN density correlates with oscillation frequency and FSIs inhibit oscillations. A standard 1% FSI population promotes γ-band oscillations with relatively little effect on MSN firing rate. Mean spikes/s averaged over a 50 ms rolling window.

causes a reduction in oscillatory behaviour. However, compared to the no–FSI model the mean firing rate of MSNs is only marginally reduced with the inclusion of 1% FSIs *(lower panel)*, while an increase from 1% to 2% FSIs approximately halves the mean firing rate.

Figure 4.6C *(upper panel)* shows power spectra for D_1 MSN activity in the segregated model with zero, 1%, or 2% FSIs in response to input of 1,500 spikes/s to a single action channel. Total lesioning of FSIs produces very strong high-frequency oscillations of ∼120 Hz, which are dramatically reduced in the 1% FSI model and entirely absent in the 2% FSI model. Mean firing rate in all FSI conditions *(lower panel)* is lower than in the uniform model and is reduced by approximately equivalent amounts with the inclusion of 1% or 2% FSIs.

4.3.3 Input Segregation Does Not Affect Selection Transitions

To explore the impact of action channel segregation on selection, we simulated a selection transition between two action requests of equal salience under control, diffuse, and unidirectional neuropeptide configurations. **Figure 4.7** shows the results of this two-channel transition for the uniform and segregated physical models.

Both versions of the physical model again produce almost identical behavioural output from the motor cortex and highly similar rate outputs from individual neuron populations. The diffuse neuropeptide configuration increases the striatal response to both action requests,

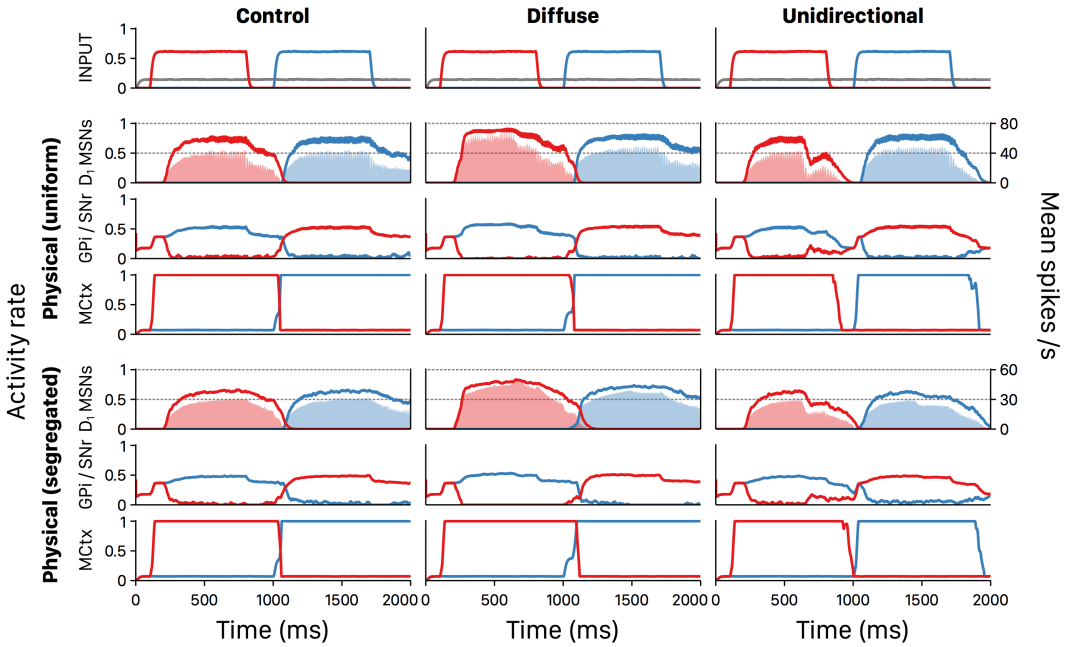

Figure 4.7: Input channel segregation does not affect switching in the physical model. Activity rates of key populations in channels c_1 (red) and c_2 (blue) following presentation of a two-channel sequence with input duration 700 ms and input salience 1,500 spikes/s. Striatal mean spikes/s averaged over a 15 ms rolling window.

while in the unidirectional configuration the onset of enkephalin's action causes a sudden drop in c_1 activity followed by a boost in c_2 activation as the result of SP influx. However, the physical segregation of action channels has a very limited effect on the converted striatal output rate and no discernible impact on selection in MCtx.

4.4 Discussion

We have expanded our striatal microcircuit model to accurately capture the connectivity and physical topography of biological tissue based on simulations of axonal and dendritic arborisation by Humphries et al. (2010). We have used this expanded microcircuit model in conjunction with the BG loop model from Section 3.3 to explore the impact of striatal topography on our previous neuropeptide modulation results, and of action channel segregation on striatal activity. By repeating key experiments from Section 3.6 we have shown that the inclusion of physical topography does not alter our previous findings, thus validating the statistical model as a tool for exploring selection and neuropeptide modulation. By manipulating the physical distribution of MSN populations representing two distinct action channels we have shown that the physical topography of the striatal microcircuit enables oscillatory spiking, the frequency and amplitude of which is modulated by the density of active MSNs and FSIs respectively.

4.4.1 Validation of Statistical Model

Repeating key experiments from Section 3.6 with the physical model allowed us to validate the results attained with the statistical model and explore any differences introduced. In both statistical and physical models the inclusion of diffuse neuropeptides facilitated the selection of an unordered action series (**Figure 4.2**), while directed SP connections facilitated the selection of an ordered action sequence and inhibited selection of a disordered sequence (**Figure 4.3**). Pruning SP connections enabled the selective inhibition of a specific distractor action in both statistical and physical models (**Figure 4.4**). The selection performance in all neuropeptide experiments was nearly identical using either model, providing a strong validation of the statistical model as a tool for probing the response of the striatal microcircuit to patterened inputs.

The validation of the statistical model permits some secondary observations. Principally, the similarity in output between statistical and physical models suggests that their structural differences do not significantly impact selection performance. Specifically, the physical organisation of MSN populations may not substantially affect their ability to respond effectively to cortical inputs. In the six-channel physical model used in Section 4.3.1, stimulation of MSN populations as sparse as 14,150 per mm^3 (16.$\dot{6}$% of all MSNs) provided a striatal response that consistently exceeded that of the statistical model. Thus, even sparsely distributed MSN populations can provide a robust collective response to a cortical action request of sufficient strength to strongly select that action. However, our model only activated a single striatal action channel at a time and so did not explore if competing inputs impact the responses of sparse populations.

Similarly, the nearly indistinguishable selection performance of the segregated and uniform physical models suggests that improved selection performance as the result of enhanced striatal response may not be a direct benefit of topographic accuracy. However, strength of striatal response is a single aspect of selection performance and physical topography may provide other direct or indirect selection benefits that the present model is not able to capture.

Some differences between the statistical and physical models were observed. In all experiments, the physical model's response to inputs was slightly enhanced compared to the statistical model. A similar result was observed by Tomkins et al. (2014), and we may postulate that this is due to a more accurate distribution of FSI $\rightarrow$ MSN connections in the physical model; while both models connect approximately 30 FSIs to each MSN (Table 3.1), the statistical model contains only 60 FSIs in total. Thus, in the statistical model it is more likely that a given FSI will make multiple connections to a single MSN, resulting in greater synchronicity and effectiveness of inhibitory inputs. Similarly, with a reduced MSN population the number of reciprocally–connected MSN–MSN pairs is likely to be artificially inflated, resulting in a greater amount of intra-population inhibition for active MSN groups.

Activation of the physical model was also found to cause significant oscillations in striatal activation, warranting further discussion.

4.4.2 Understanding Oscillations in the Microcircuit Model

Reducing the number of action channels in the physical model to two revealed unexpected oscillations in MSN spiking activity that were disrupted when MSN populations receiving inputs from each channel were physically segregated (**Figure 4.5B**). By presenting the model with inputs from a single action channel we showed that these oscillations did not arise from interactions between competing MSN populations and that the frequency of the oscillations was not related to the salience of the input (**Figure 4.6A**). Thus, stimulation of the striatal microcircuit model generates oscillatory spiking that is not contingent on tuned inputs or interactions between competing populations.

We further explored the causes for oscillatory spiking by modifying the number of FSIs in both the uniform and segregated models and providing them with input from a single action request. Lesioning FSIs resulted in exceptionally strong oscillations in both model subtypes, which were greatly diminished with a standard 1% FSI population and altogether eliminated with a 2% FSI model. However, the uniform model exhibited γ-band oscillations at $\sim$60 Hz while the segregated model exhibited high-frequency oscillations of $\sim$120 Hz (**Figures 4.6B,C**).

As segregating the model into two physically distinct regions disrupted the oscillations even when only a single channel was active, the salient feature distinguishing these two versions of the model was likely not the topographic organisation of MSN populations but the proportion of active MSNs in a given area. Activation of a single action channel targeted approximately the same number of neurons in both uniform and segregated models, but in the uniform model these were evenly distributed across the entire striatal volume while in the segregated model they were concentrated in one half. Thus, an action request in the uniform or segregated models activated 45% or 90% of the MSNs in a region respectively (allowing for the 10% of MSNs left as background in each). All FSIs across the striatal volume were targeted in all cases and so the relative density of active FSIs remained the same in both models.

Emergent oscillatory activity is therefore modulated by the density of active FSIs and MSNs. Specifically, oscillation frequency showed a positive correlation with the density of active MSNs, and oscillation amplitude showed a negative correlation with the proportion of striatal FSIs. The lack of observed oscillations in the recreation of results from Section 3.6 suggests that activation of only $\sim$17% of regional MSNs is insufficient to generate oscillations, placing a lower bound on the MSN density capable of supporting oscillatory spiking. Activation of 45% and 90% of MSNs generated oscillations of $\sim$60 Hz and $\sim$120 Hz respectively, and though this is suggestive of a simple linear relationship between active MSN density and oscillation frequency, two data points are insufficient to make such a claim.

We showed that an increase in the proportion of striatal FSIs corresponds to a reduction in the amplitude of oscillatory spiking (**Figures 4.6B,C**), but it is not immediately clear why this effect is more pronounced with a higher active MSN density. We may speculate that a higher density of active MSNs allows each active FSI in the region to target more MSNs, resulting in greater inhibitory synchronicity; this is similar to our proposed mechanism underlying the enhanced input response of the physical model over the statistical model.

We also showed that the presence of diffuse neuropeptides increased both the amplitude and frequency of striatal oscillations (**Figure 4.6A**). We have previously shown that diffuse striatal neuropeptides increase the effective salience of inputs (**Figure 3.6**) which likely accounts for the increased amplitude, but the reason for the increased frequency is not immediately apparent.

Alternative Explanations for Model Oscillations

We must also consider the possibility that the observed oscillations are a modelling artefact that do not represent a genuine phenomenon. For example, model inputs are simulated as a Poisson spike distribution sampled every 0.1 ms, but the input rate used in experiments exceeds 1,000 spikes/s and at times reaches 2,000 spikes/s. With a maximum of 1 spike event per timestep 10–20% of all timesteps will contain a spike, and this relatively high event frequency risks introducing unwanted regularity into the inputs that could prevent the formation of a true Poisson distribution. It is possible that this over-regularised input entrains MSN activity and thus generates oscillatory spikes.

We consider this to be an unlikely cause of oscillatory MSN activity as observed oscillations vary with active neural density rather than input frequency, but future work exploring this behaviour should still ensure the possibility is ruled out. This could be done in several ways. A short duration simulation with a 0.01 ms timestep to ensure a true Poisson input distribution could verify the continued presence of oscillations; alternatively, the model's input scheme could be modified to more closely resemble Humphries et al. (2009b) and thereby avoid the need for such high input rates. Finally, the model could also be supplied with perfectly regular inputs to observe if MSN spiking becomes entrained to an external rhythm.

Similarly, the model also currently includes a strong degree of synchronicity between inputs to D_1 and D_2 MSNs that could potentially lead to unwanted regularity and input entrainment. This should be corrected for future work on striatal oscillations.

4.4.3 Significance of Striatal Oscillations

Striatum Generates Oscillations

Striatal local field potential (LFP) oscillations have been frequently reported *in vivo*, but it can be difficult to show a correspondence between LFP measurements and local spiking activity as contributions to potentials may include afferent synaptic inputs, intrinsic membrane potentials, or spiking activity from a neighbouring region. Indeed, striatal oscillations are often coherent with oscillations in external populations such as amygdala (Popescu et al., 2009), cortex (Berke, 2009), or hippocampus (Tort et al., 2008). However, research suggests that the striatum is able to produce strong LFPs as the result of population spiking (Misgeld et al., 1979; Pennartz et al., 1990). Our results therefore reinforce suggestions that the striatum is able to generate oscillations locally (van der Meer et al., 2010)

It has been suggested that striatal oscillations are generated by FSI populations entrained to inputs (Sharott et al., 2009), but our results suggest that neither input entrainment (**Figure 4.6A**) nor FSIs (**Figures 4.6B,C**) are necessary to generate oscillatory spiking. However, there is strong evidence to suggest that entrainment does occur (Carmichael et al., 2017), though we did not explore the circumstances that may cause this.

Importance of γ-band Oscillations

γ-band LFP oscillations have been frequently recorded in rat ventral striatum (Kalenscher et al., 2010; van der Meer and Redish, 2009), and several authors have reported a link between γ-band striatal oscillations and key aspects of behaviour such as movement initiation (Masimore et al., 2005) or reward (Cohen et al., 2008; Tort et al., 2008). The observation of unexpected γ oscillations in our striatal microcircuit model is therefore especially interesting as it suggests that oscillation generation may be an important aspect of striatal microarchitecture functionality and that the timecourse of GABA may predispose such oscillations to the γ band (Buzsáki and Wang, 2012).

Our results also agree with research showing that FSIs selectively promote the activity of MSNs displaying γ-band oscillations while inhibiting other MSN activity (O'Hare et al., 2017). Though our limited dataset precludes a detailed analysis, **Figures 4.6B,C** show that a striatal population of 1% FSIs inhibits high-frequency oscillations substantially more than γ-band oscillations. Furthermore, Gittis et al. (2011) report that FSI hypofunction results in behavioural dysfunction without changing the average firing rate of MSNs, which could correspond with the abnormally strong oscillations and only marginally increased firing rate observed with lesioned FSIs in **Figure 4.6B**. FSI networks may be particularly well-suited to promote γ oscillations due to the dynamics of FSI–FSI gap junctions (Bartos et al., 2007; Traub et al., 2001), so FSIs may perform dual functions of moderating striatal

oscillations in general while selectively promoting those in the γ band. Further modelling work may help clarify the role of FSI gap junctions in promoting γ-band activity.

It is interesting to note that a higher active MSN density exhibited higher frequency oscillations but with a lower mean firing rate (**Figure 4.6C**); the striatal microcircuit model therefore appears to respond to inputs optimally when the active MSN density is relatively low. Thus, if future modelling work reveals that oscillation frequency correlates continuously with active MSN density then this optimal response density could represent a driving force underlying the topographic organisation of MSN populations.

Relationship Between Oscillation and Selection

The inputs to our model are biologically abstracted and we are therefore unable to draw conclusions about the relationship between oscillatory activity and inputs from specific neural structures. Furthermore, the spike–to–rate converted input for the BG loop model was designed to smooth out short-term variances in spiking activity, making our hybrid model ill-equipped to probe the effects of striatal oscillations on selection (**Figure 4.7**). Nonetheless, we may cautiously speculate as to the function of striatal oscillations.

It has been suggested that cortical γ-band oscillations may entrain populations encoding higher-order stimulus properties (Gray and Singer, 1989) and may function as a reference for temporal coding (Buzsáki and Chrobak, 1995). Striatal γ oscillations are "heterogeneous and regionally differentiated" (Kalenscher et al., 2010), and so we may therefore speculate that striatal γ oscillations generated by MSN populations bind together groups responding to action requests.

Such groups need not correspond directly to a single action request. Striatal γ oscillations have been found to relate to dissociable aspects of actions such as movement (Masimore et al., 2005) or reward (van der Meer and Redish, 2009), and models have previously shown that the BG loop is able to store ordinal sequence information in nested γ oscillations (Fukai, 1999), so coherently oscillating MSN groups may be able to represent anything from distinct features of single actions to entire action sequences. This is supported by recent work showing that the striatum encodes actions in localised, overlapping MSN populations (Klaus et al., 2017) and that entire action sequences can be encoded as a single unit (Martiros et al., 2018).

Chapter 5

General Discussion

In this thesis we have explored how striatal neuropeptides and striatal topography influence the selection of action sequences. We have utilised pre-existing models of a GABAergic striatal microcircuit and the basal ganglia–thalamocortical loop, and developed a novel model of the phenomenological effects of two neuropeptides to create a hybrid basal ganglia model to explore their effects.

Results specific to each experiment have been discussed in their respective chapters, so this chapter will discuss the strengths and weaknesses of the current approach, summarise the research findings, and make suggestions for future work.

5.1 Validity and Limitations of Modelling Approach

The question of validity is central to any modelling exercise; a model may be elegant and fascinating but it must still represent a real-world phenomenon to a useful degree, otherwise it is simply a toy.

However, assessing a model's validity is not an exact procedure. Validity is a measure of how accurately model inputs, processes, and outputs match their real-world counterparts, but as all models necessarily omit some features of the systems on which they are based the model's design goal must be considered when assessing its performance in each of these areas. Such omissions should be deliberate and justified, and the limitations these choices impose on the model should be considered. Given such omissions, is the model able to replicate the system's behaviour, explain observations, and make predictions in line with the design goal? If so, then we may reasonably declare the model to be valid.

The validity of neural modelling as a general practice is no longer seriously in question; a wide number of neural models at different scales have been created to a high degree of accuracy, and their explanations and predictions have often been indispensable. We must still ensure, however, that the specific models and practices employed here are valid for our

purposes and that our interactions with them do not impair their validity. We must also address the limits of our model's validity and highlight areas where it is lacking.

5.1.1 Neuropeptide Model

Our novel model of SP and enkephalin's glutamatergic EPSP modulation relies on the accuracy of the underlying neuron model of Humphries et al. (2009a) and available experimental data (Blomeley and Bracci, 2008, 2011; Blomeley et al., 2009). The Izhikevich point neurons used to construct the model MSNs and the striatal microcircuit are well-established and widely employed so the validity of the basic technique seems assured.

However, while we showed in Chapter 2 that our model is able to capture these neuropeptides' known effects on glutamatergic EPSPs, its validity is limited by two main factors: the paucity of experimental data currently available, and the decision to implement a purely phenomenological model. The Moyer et al. (2007) MSN model on which the reduced Humphries et al. model is based represents a synthesis of measurements from dozens of papers, while our neuropeptide model is based on data from only three. It is inevitable that further research will contribute additions and refinements to what is known about their effects that will make our model obsolete.

The phenomenological nature of the model further limits its validity to use with a specific MSN model and exploration of glutamate modulation. We cannot explain phenomena such as the observed delays between neuropeptide release and modulatory effect, or SP's direct depolarisation of MSNs (Blomeley and Bracci, 2008), nor can we reuse our neuropeptide models in other populations. A more detailed model simulating neuropeptide diffusion, their interactions with NK_1 and μ-opioid receptors, and the intracellular actions of these receptors may explain the processes underlying some of these phenomena, but the aforementioned lack of experimental data would make construction of such models itself a substantially predictive work.

Allowing for the limitations set by computational capacity and currently known data, we therefore consider our neuropeptide model to be of good validity for the task in which we have employed it.

5.1.2 Microcircuit Model

The biological fundamentals and connectivity underlying the microcircuit model are well-established, and it has been previously employed to explore the formation of cell assemblies (Humphries et al., 2009b) and selection in Huntington's disease (Tomkins et al., 2014). There are two principal concerns as to the model's validity; the use of the statistical model to simulate action selection and the omission of several interneuron classes.

Utilising simulations of neuronal connectivity (Humphries et al., 2010) allowed us to create a

relatively small-scale microcircuit model with only a few thousand neurons that maintained an appropriate number of each connection type. Although the number of connections was statistically accurate the reduced number of neurons meant their distribution could not be, and it was not clear if a model with no topography and an usual connectivity pattern would generate unexpected or unrealistic behaviour. The lack of topography also made it impossible to organise striatal inputs, with unknown consequences for selection.

By directly comparing the performance of a small-scale model with statistically accurate connectivity against a large-scale model with physically accurate connectivity we were able to probe the limits of the smaller model's validity and verify its performance in selection tasks. Where the mean spiking rate of MSNs and their response to glutamatergic inputs is the key consideration, the physical and statistical models consistently showed nearly identical responses in all situations thus validating the statistical approach. However, one interesting difference also emerged: population–level MSN activity in the physical model exhibited strong oscillatory behaviour. The generation of oscillations within the striatal microcircuit was a wholly unexpected development that highlights the limits of the statistical approach, and further work exploring the causes and effects of these oscillations would not be able to use the statistical model.

Several other simplifications or inaccuracies included in the microcircuit model warrant mentioning. As described here and in many other models, the only difference between D_1 and D_2 MSNs is their response to dopamine modulation and their output targets. However, D_1 and D_2 MSNs have distinct anatomical differences (Gertler et al., 2008) and receive cortical inputs differentiated by region (Wall et al., 2013) and neuron type (Lei et al., 2004). Connections between D_1 and D_2 MSNs are also differentiated by quantity and strength (Planert et al., 2010; Taverna et al., 2008), which may have important implications for the striatum's operation as a threshold detection device (Bahuguna et al., 2015). It is likely that additional aspects of striatal function are missed by models that fail to sufficiently distinguish these two distinct neural populations.

Similarly, the wholesale omission of interneuron classes is likely to have limited the validity of the current model in unknown ways. The impact of cholinergic interneurons is of particular interest; SP has been found to modulate the activity of cholinergic interneurons (Govindaiah et al., 2010), and TAN outputs are able to inhibit cortical input to MSNs (Ding et al., 2010) and are strongly correlated with the likelihood of a behavioural response to a stimulus (Blazquez et al., 2002). The significance of these features for sequence selection is unknown, and incorporation of this interneuron class should therefore be a priority for future iterations of the striatal microcircuit model.

The inclusion of biological features should not be pursued for its own sake, however. Our model utilises an NMDA saturation function developed to probe changes in Huntington's disease (Tomkins et al., 2012) although NMDA saturation is not a salient feature of sequence selection. While the inclusion of this feature is unlikely to have impaired the validity of our

model we consider it to have been an unnecessary addition.

5.1.3 Basal Ganglia–Thalamocortical Loop Model

The BG loop model is abstracted to a higher degree than the microcircuit model, and though many more detailed models of BG function and interactions exist (Helie et al., 2013; Schroll and Hamker, 2013) our model is primarily employed to interpret striatal output and enhance striatal validity. As a highly abstracted model it necessarily captures fewer biological features and could be considered intrinsically less valid for that reason; however, as a tool for processing the output of striatal computation it is the validity of the model's high-level structure and output that is important rather than the detailed simulation of individual populations or processes. The underlying connectivity, behaviour, and interactions of each population is in line with accepted research and so we consider the abstracted approach to be sound.

Two features are of particular importance: the analysis of motor cortex activity as a quantitative measure of selection, and the use of MCtx–VLT feedback as active memory to sustain selection. In principle, both of these features could be replicated in other ways — a quantitative analysis of MSN activity could determine activity sufficient for selection, and a modified input pattern could represent sustained selection — but these measures would each incrementally increase the strain on the validity of the microcircuit model. By incorporating them into the BG loop model, the proper functioning of each model reinforces the validity of the other; the microcircuit's model ability to replace rate-coded BG loop populations suggests it is correctly constructed an integrated, and the BG loop model's ability to select and sustain actions suggest its selection architecture is well implemented and responding appropriately.

Once again, the oscillatory behaviour of the physical model shows the limits of this validity. Oscillatory behaviour in striatum is common but has been linked to changes in behaviour that our hybrid model did not replicate. It is unclear if this is a limitation of the BG loop model or specifically of the spike–to–rate conversion created to integrate the two models, but in either case it illustrates that the validity of a given model depends on appropriately matching its complexity to its intended use. Striatal oscillations were an unexpected occurrence that the hybrid model was not designed to explore, and the BG loop architecture would need to be modified to properly study their implications.

5.1.4 Model Inputs

Although not properly a part of either model, the input to the entire hybrid model is fundamental to its operation and deserves similar scrutiny. Several assumptions have been made in supplying input to the model, including the use of a Poisson spike distribution, the rate, shape, and duration of inputs, and their distribution to MSN populations.

Several of these assumptions are not intended to accurately mimic biological systems and fall into the class of simplifications all models bear to make them tractable and useful. These include the use of Poisson spike trains, non-gradated inputs projected equally to both MSN subtypes, and each MSN's representation of a single, distinct action request. These simplifications are a matter of judgement, and as our research was intended to show the existence of a new phenomenon rather than probing the fine details of a known phenomenon, using idealised representations of features that are either irrelevant, of little interest, or for which insufficient data exists reduces the number of independent variables and the chances of introducing error or noise into the model.

The input connectivity to our model differs from Humphries et al. (2009b). We provide individual Poisson spiking inputs to each MSN rather than inputs representative of multiple afferent spike trains, and although our method results in an equivalent glutamatergic influx we forego the possibility of MSNs receiving multiple spikes in any timestep and thus somewhat alter the overall pattern of innervation. Our approach also introduces an undesirable degree of input synchronicity between striatal populations, which should be addressed before exploring oscillatory behaviour further. While low-level input patterning was not a primary focus of our model, for more biologically realistic input saliences, patterning, and for greater compatibility with similar models we would recommend updates to our input scheme in future iterations of the model.

The higher-level pattern of phasic inputs sustained by BG feedback is biologically inspired, however. Evidence suggests that transitions during sequence execution are marked by transient bursts of cortical activity (Averbeck et al., 2006; Dehaene et al., 2015) and it is plausible that positive MCtx–VLT feedback may support sustained selection (Chambers et al., 2005; Haber and McFarland, 2001). Our results reinforce the validity of this approach.

5.2 Research Contributions and Predictions

5.2.1 Phenomenological Neuropeptide Models

We have created a novel phenomenological model of the glutamate EPSP modulatory effects of substance P and enkephalin and validated it against available neurophysiological data. Our model captures the neuropeptides' facilitatory or inhibitory postsynaptic effects at short timescales of under 100–200 ms and at longer timescales of up to 2 s through changes to variables in a simple function that make it suitable for inclusion in large-scale striatal microcircuit models with low computational overhead.

5.2.2 Hybridisation of Basal Ganglia Models

We have integrated two computational models of differing types into a single hybrid model that allows us to simulate the membrane dynamics of individual striatal neurons and analyse this model's output using a simpler rate-coded model of BG activity. This hybridisation required the creation of interconnects to convert between rate-coded and spiking activity and allowed us to avoid spending computational resources on detailed spiking BG populations.

5.2.3 Striatal Neuropeptides Enhance Sequence Selection

By simulating sequential cortical action requests innervating striatal populations and the resultant neuropeptide release, we showed that the interacting effects of SP and enkephalin enhance the selection of actions forming part of a learned sequence. Specifically, we showed that absent any connectivity patterning, the presence of both neuropeptides facilitates the selection of subsequent actions, and that patterning following the inputs' semantic order allows the sequence ordering to be protected by inhibiting the selection of unordered action requests. We also showed that patterning SP projections allows for targeted inhibition of distracting actions, suggesting that complex SP patterning may permit or prevent the selection of specific action requests in different contexts. Our results thus predict that SP connections in the striatum are plastic and are modified as part of sequence learning or habituation. Specifically, we predict that SP connections between MSN populations representing actions in a sequence may be weakened or pruned if the successive selection of both actions would impair sequence execution.

Our results further predict that SP facilitates selection primarily by reinforcing the sustained selection of action requests, and secondarily by enhancing the salience of subsequent requests. Finally, our results predict that SP is required to overcome enkephalinergic inhibition, and that if our assumption of diffuse enkephalin projections throughout striatum is correct then selective disinhibition is the common action mode of the striatum as well as the basal ganglia in general.

Several potential experiments could test these hypotheses. Most directly, it would be enlightening to observe the impact of SP and/or enkephalin antagonists in the striatum, specifically in the contexts of action sequence learning and execution. Our results predict that blocking enkephalin should result in an increased incidence of disordered or incomplete sequence execution, and that blocking SP should cause an inability to fully execute individual actions. Additionally, a comparison of SP and enkephalin connectivity in different striatal regions and at different stages of sequence learning could reveal if SP network pruning is occurring during the process, and if enkephalin connectivity is static.

Further modelling work could also help refine these predictions; for example, an exploration of plausible learning rules guiding SP plasticity could generate predictions regarding input patterns that trigger connectivity modifications and how they occur.

5.2.4 Validation of Statistical Connectivity

We directly compared our small-scale model with no topography and statistically–derived connectivity against a large-scale striatal microcircuit with physically accurate topography and connectivity based on work by Humphries et al. (2010). Both models showed nearly identical time-averaged spiking behaviour in response to all selection tasks, thus validating the statistical approach for appropriate modelling tasks. Our results thus predict that the topographic organisation of MSN populations is not directly related to the salience of their collective response to inputs.

Although it may be possible to test this directly via targeted stimulation of MSN populations, it might be more fruitful to conduct further detailed modelling work to uncover specific benefits of topographic MSN organisation and selection functions that are impaired in its absence.

5.2.5 Striatal Microcircuit Activation Generates Oscillations

By stimulating the physical striatal microcircuit we unexpectedly showed that the microcircuit topography robustly generates oscillatory spiking. We showed that oscillatory activity does not arise as the result of competition between active populations and that its frequency is not related to the strength of spiking inputs. We further showed that the density of active MSNs and FSIs impacts the frequency and amplitude of evoked oscillations respectively, and that FSIs appear to selectively promote activity in MSN populations exhibiting γ-band oscillations. We speculate that γ-band oscillations may bind together MSN populations representing features of actions or action sequences and that the selective promotion of γ oscillations may constitute a driving force underlying the topographic organisation of striatal MSNs.

Further work on striatal γ oscillations should first explain their origins and explore these observations in detail, for which we propose a follow-up modelling study. Several pertinent issues should be explored, including the input, target, and topography parameters that support γ oscillation generation, and other striatal features that facilitate or inhibit it. Once the causes of striatal oscillations are better understood, experimental interventions can be developed to explore oscillatory behaviour *in vivo*.

5.3 Future Work

This is the first research of which we are aware to model the striatal effects of SP and enkephalin and to explore a co-operative role for the MSN network. There are therefore many opportunities for refinement and expansion, some of which we will discuss here.

5.3.1 Enhanced Microcircuit Validity

As discussed, validity is key to achieving accurate and useful results with any model. As such, there are several important details omitted from or simplified in the current model that could increase its validity and potentially reveal new functionality.

The inclusion of the cholinergic TAN should be a particular priority. TANs have several unique properties that make them of particular interest but also present a modelling challenge; their activity modulates MSN spiking (Franklin and Frank, 2015; Crittenden et al., 2017), is itself modulated by SP (Govindaiah et al., 2010), and has been linked to multiple aspects of learning and motivated behaviour (Aosaki et al., 1994; Apicella, 2007). However, their activity and influence differs between striatal regions (Yamada et al., 2004; Threlfell and Cragg, 2011), and few striatal models explicitly depict a specific sub-region. Nevertheless, useful models of TANs and their striatal function have previously been created (Ashby and Crossley, 2011; Stocco, 2012; Tan and Bullock, 2008) so their inclusion in future iterations of the model should be prioritised.

Greater differentiation of D_1 and D_2 MSN populations may also reveal important functionality and should be achievable with relatively little effort. In particular, the relative paucity of $D_1 \rightarrow D_2$ connections and weakness of D_1–originating connections (Taverna et al., 2008), and the FSI preference for connecting to D_1 MSNs (Gittis et al., 2010) imply significant structural asymmetries not reflected in most models and which may help the striatum to function as a threshold device (Bahuguna et al., 2015).

5.3.2 Exploration of Organisational Features

Given the topographic organisation of striatal tissue, the effects of striatal topography and input patterning are probably the most salient organisational features to explore further. We have briefly explored the impact of dividing the striatum into two neighbouring regions representing distinct action request targets, although as discussed the current model was ill-equipped to draw conclusions about the effects on selection. In addition to remedying this, it would be enlightening to observe the oscillatory behaviour of active striatal populations separated by a neutral region, with 'fuzzy' or overlapping boundaries, or with more complex topographic organisation.

Although hinted at by the emergent striatal oscillations in the physical model, we did not look into the propensity of the striatal microcircuit to dynamically form cell assemblies (Humphries et al., 2009b) that may encode network states (Carrillo-Reid et al., 2008) and facilitate switching (Ponzi and Wickens, 2010). The function of these cell assemblies, their relationship with oscillations and striatal neuropeptides, and the conditions that influence their formation are topics likely to be of significant interest.

Modifying the pattern of inputs to the model may also prove informative. The current

model consistently used idealised but unrealistic non-gradated inputs that projected to both D_1 and D_2 MSN populations equally, while neurophysiological recordings suggest that inputs are differentiated between MSN subpopulations (Wall et al., 2013). Additionally, transient cortical inputs are gradated (Averbeck et al., 2003) and the ramping signal likely relates to the selection of future actions (London et al., 2018). Maintaining clean selection in response to overlapping, gradated inputs is a key requirement of selection defined in Section 1.1.1 that was not tested in the current model and seems likely to be biologically significant, so should be explored in future research.

5.3.3 Plasticity

Modelling striatal plasticity may be the most challenging expansion of the current research, but also the most potentially rewarding. However, 'plasticity' is a broad category with multiple physiological manifestations, so we restrict our suggestions to those highlighted by the current research.

Our research predicts plasticity within the striatal SP network, and further experimental research is recommended to resolve this issue. If SP connections are plastic, modelling plausible learning rules could help refine the expected parameters of this plasticity. For example, what level of modification to the SP network is necessary to achieve selection benefits? If they are not, then future iterations of our model may explore the implications for sequence learning and execution. Additionally, the cholinergic TAN omitted from our model is thought to play a role in striatal plasticity (Calabresi et al., 1999) and is modulated by SP (Govindaiah et al., 2010), making this interneuron of great interest for such research in either case.

We have also proposed that striatal γ oscillations may bind together MSN populations representing related input features. The administration of dopaminergic agents influences switching between γ-50 and γ-80 oscillations (Berke, 2009), but the relative functions of these γ powers is complex and unclear (van der Meer et al., 2010). Our model used a static term to represent tonic dopamine and we did not explore the impact of modifying this or introducing phasic dopamine; it seems likely that an investigation of striatal plasticity causing (or caused by) γ oscillations will need to include a more detailed dopamine model to obtain satisfactory results.

5.4 Conclusion

Our results strongly imply an under-appreciated role for striatal neuropeptides as messengers within the striatal complex. Rather than playing a limited modulatory role, they appear to have a significant influence on the selection of striatal inputs; in particular, SP network pruning can protect action sequences from distractions and disordered execution. As habits

consist of ordinal action sequences requiring minimal concentration we predict that SP plasticity is essential for habit learning, providing a specific hypothesis for future research.

Conversely, our findings suggest that the topographic organisation of MSN populations has relatively little impact on the strength of their response to cortical inputs. Instead of relying on entrainment to external populations, striatal oscillations emerge from innervation of the topographic microcircuit and are tuned by the density of active neurons. These oscillations appear to not impact selection directly but may instead bind together MSN groups representing congruent features of actions or sequences, potentially suggesting another means of encoding information within the striatum.

We have not shown a direct link between neuropeptide release and oscillation generation, other than a slight modulation of oscillation frequency and amplitude. This may be due to omissions in our model, or it may be that no further link exists; if these two features relate to encoding separate aspects of action requests then their interaction may be informationally inefficient and undesired. The observed modulation of oscillations may also prove to be significant for reasons not yet known, and so our assessment of the relationship between these two findings for now remains a scientifically satisfying "we don't know".

Appendix A

Model Descriptions

Following are full descriptions of the striatal microcircuit, neuropeptide, and basal ganglia–thalamocortical loop models, including notations for different versions and all variables used, following the format proposed by Nordlie et al. (2009). Items applicable to specific model subtypes will be indicated as relevant to the *statistical* model (SM) or the *physical* model (PM).

A.1 Striatal Microcircuit Model

Table A.1: Striatal microcircuit: Summary

	***Statistical* model (SM)**	***Physical* model (PM)**
Populations	Three: D_1 MSNs, D_2 MSNs, and FSIs	
Topology	None	Neurons given random co-ordinates with enforced minimum distance
Connectivity		
↳ *GABA*	All–to–all, pruned according to statistical probability	Probabilistic, according to distance falloff and connection type
↳ *Neuropeptide*	GABA connections pruned according to probability profile	
↳ *To BG loop*	Fixed channel–convergent	
Neuron model	Modified Izhikevich point neurons with reset condition and spiking threshold	
Channel model	—	—
Synapse model	Conductance-based single-exponential with dopamine and neuropeptide modifiers	
Plasticity	—	—
Input	Independent Poisson spike sources project to each neuron MCtx output from BG–thalamocortical loop	
Measurements	Raster log of spikes Mean spike frequency converted to rate output	

Table A.2: Striatal microcircuit: Populations

Name	Type	Size (SM)	Size (PM)	Organisation
D_1 MSNs (d1)	Modified Izhikevich	3,000	42,450	Six groups of 500 (SM)
D_2 MSNs (d2)		3,000	42,450	or 7,075 (PM) neurons
FSIs (fs)	Modified nonlinear Izhikevich	60	∼849	—

Table A.3: Striatal microcircuit: Connectivity

Connection	Algorithm
MSN → MSN ↳ GABA	For d1 → d1, d1 → d2, d2 → d1, d2 → d2: 1. Create all–to–all connections 2. Assign each connection value $R \sim U([0,1])$ 3. Connection probability $P = \frac{\text{Number of afferent connections}^*}{\text{Number of target neurons}}$ 4. Remove connections where $R \leq P$
↳ Substance P ↳ *Diffuse* ↳ *Pruned* ↳ *Unidirectional*	For d1 → d1, d1 → d2: All GABA connections co-release SP All GABA connections except $d1_1 \rightarrow d1_6$, $d1_1 \rightarrow d2_6$ co-release SP Only GABA connections $d1_c \rightarrow d1_{c+1}$, $d1_c \rightarrow d2_{c+1}$ where $c < 4$ co-release SP
↳ Enkephalin	For d2 → d1, d2 → d2: All GABA connections co-release enkephalin
FSI → MSN	For fs → d1, fs → d2: As MSN → MSN (GABA)
FSI → FSI ↳ GABA	For fs → fs: As MSN → MSN (GABA)
↳ Gap junction	As MSN → MSN (GABA) except: 1. Begin with all fs → fs GABA connections

*Values for expected number of afferent connections taken from Table 5, Humphries et al. (2010)

Table A.4: Striatal microcircuit: Neuron models

Name	Dynamics	Modifications
D_1 MSN	$C\dot{v} = k(v - v_r)(v - v_t) - u + I$ $\dot{u} = a\left[b(v - v_r) - u\right]$	$v_r \leftarrow v_r(1 + K\phi_1)$ $d \leftarrow d(1 - L\phi_1)$
D_2 MSN	$I = I_{\text{ampa}} + I_{\text{gaba}} + B(v)I_{\text{nmda}}$ Reset: if $v > v_{\text{peak}}$ then $v \leftarrow c, u \leftarrow u + d$	$k \leftarrow k(1 - \alpha\phi_2)$
FSI	$C\dot{v} = k(v - v_r)(v - v_t) - u + I$ $\dot{u} = \begin{cases} -au & \text{if } v < v_b \\ -a\left[b(v - v_b)^3 - u\right] & \text{if } v \geq v_b \end{cases}$ $I = I_{\text{ampa}} + I_{\text{gaba}}$ Reset: if $v > v_{\text{peak}}$ then $v \leftarrow c, u \leftarrow u + d$	$v_r \leftarrow v_r(1 - \eta\phi_1)$
Rate–to–spike	Each timestep, emit spike if $y_i^{\text{mc}} r_{\max} \tau_{\text{bg}} > P$	$P \sim U([0,1])$

Table A.5: Striatal microcircuit: Synaptic models

Type	**Dynamics**
Synapse	$I_z = \bar{g}_z h_z (E - v)$ $\dot{h}_z = \frac{-h_z}{\tau_z}$, and $h_z(t) \leftarrow h_z(t) + \left[1 - \frac{h_z(t)}{\omega_z}\right] S_z(t)$ Where z is AMPA, GABA or NMDA and $S_z(t)$ is the number of spikes arriving at z receptors at time t $B(v) = \frac{1}{1 + \frac{[\mathrm{Mg}^{2+}]_0}{3.57} \exp(-0.062v)}$
Neuropeptide	$A_p(t) = \sum_i S_p \left[\exp\left(\frac{-(t-t_i)}{\tau_p^f}\right) - \exp\left(\frac{-(t-t_i)}{\tau_p^r}\right)\right]$ $N_p(t) = \beta_p \left[1 - \exp\left(-\frac{A_p(t)}{\lambda_p}\right)^{\kappa_p}\right]$ Where p is SP or enkephalin and S_p is the number of spikes causing a release of neuropeptide p
Gap junction	$\tau \dot{v}_{ij}^* = (v_i - v_{ij}^*) + (v_j - v_{ij}^*)$
Name	**Modifications**
MSN	$I_z = I_z \left[1 + N_{\mathrm{sp}}(t - \tau_{\mathrm{sp}}^d)\right] \left[1 - N_{\mathrm{enk}}(t - \tau_{\mathrm{enk}}^d)\right]$ Where z is AMPA or NMDA
↳ D_1	$I_{\mathrm{nmda}} = I_{\mathrm{nmda}}(1 + \beta_1 \phi_1)$
↳ D_2	$I_{\mathrm{ampa}} = I_{\mathrm{ampa}}(1 - \beta_2 \phi_2)$
FSI	$I_{\mathrm{gaba}} = I_{\mathrm{gaba}}(1 - \epsilon_2 \phi_2)$

Table A.6: Striatal microcircuit: Inputs

Input	**Targets**		
	D_1 MSN (d1)	D_2 MSN (d2)	FSI (fs)
Sensory cortex (sc)	$\mathrm{sc}_c^i \rightarrow \mathrm{d1}_c^i$, $1 \leq i \leq n$	$\mathrm{sc}_c^i \rightarrow \mathrm{d2}_c^i$, $1 \leq i \leq n$	$\mathrm{sc}_c^i \rightarrow \mathrm{fs}^i$, $1 \leq i \leq n$
Motor cortex (mc)	$\mathrm{mc}_c \rightarrow \mathrm{d1}_c^i$, $1 \leq i \leq n$	$\mathrm{mc}_c \rightarrow \mathrm{d2}_c^i$, $1 \leq i \leq n$	$\mathrm{mc}_c \rightarrow \mathrm{fs}^i$, $1 \leq i \leq n$
	Where n is the number of neurons in an action channel		Where n is the number of FSIs

Table A.7: Str. microcircuit: MSN properties

Par.	**Value**	**Source**
a	0.01	Mahon et al. (2000)
b	−20	Izhikevich (2007)
c	−55 mV	
k	1	
v_r	−80 mV	
v_{peak}	40 mV	
C	15.2 pF	Humphries et al. (2009a)
d	91	
K	0.0289	
L	0.331	
v_t	−29.7 mV	
α	0.032	

Table A.8: Str. microcircuit: FSI properties

Par.	**Value**	**Source**
a	0.2	Izhikevich (2007)
b	0.025	
d	0	
k	1	
v_{peak}	25 mV	
v_b	−55 mV	
C	80 pF	Tateno et al. (2004)
c	−60 mV	
v_r	−70 mV	
v_t	−50 mV	
ϵ	0.625	Fits Gorelova et al. (2002)
η	0.1	Fits Bracci et al. (2002)

Table A.9: Striatal microcircuit: Synapse properties

Parameter	Value		Source
$E_{\text{ampa}}, E_{\text{nmda}}$	0	mV	
E_{gaba}	−60	mV	
τ_{ampa}	6	ms	Moyer et al. (2007)
τ_{nmda}	160	ms	
τ_{gaba}	4	ms	
τ_{bg}	0.1	ms	Humphries and Gurney (2007)
$\tau_{\text{fs–gap}}$	5	ms	Fitted to Galarreta and Hestrin (1999)
$[\text{MG}^{2+}]_0$	1	mM	Jahr and Stevens (1990)
g_{ampa} CTX–MSN	0.4	nS	Tomkins et al. (2014)
g_{ampa} CTX–FSI	1	nS	Fits linear rise in EPSC data from Gittis et al. (2010)
g_{nmda}	0.2	nS	Fixed by maintaining the Moyer et al. (2007) 2:1 AMPA:NMDA ratio
g_{gaba} MSN–MSN	0.75	nS	Koós et al. (2004)
g_{gaba} FSI–MSN	3.75	nS	Tomkins et al. (2014)
g_{gaba} FSI–FSI	1.1	nS	Gittis et al. (2010)
$g_{\text{fs–gap}}$	5	nS	Fitted to Galarreta and Hestrin (1999)
β_1	0.5		
β_2	0.3		Tomkins et al. (2014)
ϕ_1, ϕ_2	0.3		
ω_{ampa}, ω_{gaba}	2000		
ω_{nmda}	600		Tuning (see text)
r_{max}	2000		

Table A.10: Striatal microcircuit: Neuropeptide properties

Par.	Value	Source
β_{sp}	0.47	Blomeley and Bracci (2008)
τ^r_{sp}	10 ms	
τ^f_{sp}	200 ms	
τ^d_{sp}	40 ms	Tuning (see text)
λ_{sp}	5.5	
κ_{sp}	2.5	
β_{enk}	0.3	Blomeley and Bracci (2011)
τ^r_{enk}	15 ms	
τ^f_{enk}	300 ms	
τ^d_{enk}	400 ms	Tuning (see text)
λ_{enk}	4.5	
κ_{enk}	1	

A.2 Basal Ganglia–Thalamocortical Loop Model

Table A.11: Basal ganglia–thalamocortical loop: Summary

Populations	Five: STN, GPe, GPi/SNr, VLT, MCtx
Topology	—
Connectivity ↳ *To striatum*	One–to–one and all–to–all (STN → GPe, STN → GPi/SNr) Fixed channel–divergent
Neuron model	Leaky integrator
Channel models **Synapse model** **Plasticity**	—
Input	Rate–converted Poisson spike sources project to MCtx and STN Rate–converted spiking MSN output from striatum projects to GPe and GPi/SNr
Measurements	Activity rate output

Table A.12: Basal ganglia–thalamocortical loop: Populations

Name	Type	Size	Organisation
STN (stn) GPe (gp) GPi/SNr (snr) VLT (vlt) MCtx (mc)	Leaky integrator	6	Six action channels $c_1 \dots c_6$ of 1 neuron each

Table A.13: Basal ganglia–thalamocortical loop: Connectivity

Connection	Type
STN → GPe STN → GPi/SNr	All–to–all
GPe → STN GPe → GPi/SNr GPi/SNr → VLT VLT → MCtx MCtx → VLT MCtx → STN	One–to–one

Table A.14: Basal ganglia–thalamocortical loop: Neuron models

Type	**Dynamics**
Activation	$\dot{a} = k(a - u) + u$
Output	$y(t) = F(a(t), \theta) = \begin{cases} 0 & \text{if } a(t) \leq \theta \\ a(t) - \theta & \text{if } \theta < a(t) < 1 - \theta \\ 1 & \text{if } a(t) \geq 1 - \theta \end{cases}$
Spike–to–rate	$r_i^s(t) = \sum_i S_s \left[\exp\left(\frac{-(t-t_i)}{\tau_s^f}\right) - \exp\left(\frac{-(t-t_i)}{\tau_s^r}\right)\right]$ $y_i^s(t) = 1 - \exp\left(-\frac{r_s(t)}{\lambda_s}\right)^{\kappa_s}$ Where s is d1, d2, or sc and S_s is the number of spikes arriving from population s
Population	**Dynamics**
D1 (d1′)*	$u_i^{\mathrm{d1'}} = (w_{\mathrm{d1}} y_i^{\mathrm{sc}} + w_{\mathrm{d1}} y_i^{\mathrm{mc}})(1 + \chi_{\mathrm{d1'}})$ $y_i^{\mathrm{d1'}} = F\left(a_i^{\mathrm{d1'}}, 0.2\right)$
D2 (d2′)*	$u_i^{\mathrm{d2'}} = (w_{\mathrm{d2}} y_i^{\mathrm{sc}} + w_{\mathrm{d2}} y_i^{\mathrm{mc}})(1 - \chi_{\mathrm{d2'}})$ $y_i^{\mathrm{d2'}} = F\left(a_i^{\mathrm{d2'}}, 0.2\right)$
STN (stn)	$u_i^{\mathrm{stn}} = w_{\mathrm{stn}} y_i^{\mathrm{sc}} + w_{\mathrm{stn}} y_i^{\mathrm{mc}} + w_{\mathrm{gp}} y_i^{\mathrm{gp}}$ $y_i^{\mathrm{stn}} = F\left(a_i^{\mathrm{stn}}, -0.25\right)$
GPe (gp)	$u_i^{\mathrm{gp}} = w_{\mathrm{gp}} \sum_j^n y_j^{\mathrm{stn}} - y_i^{\mathrm{d2}}$ $y_i^{\mathrm{gp}} = F\left(a_i^{\mathrm{gp}}, -0.2\right)$
GPi/SNr (snr)	$u_i^{\mathrm{snr}} = w_{\mathrm{snr}} \sum_j^n y_j^{\mathrm{stn}} - y_i^{\mathrm{d1}} - w_{\mathrm{gp}} y_i^{\mathrm{gp}}$ $y_i^{\mathrm{snr}} = F\left(a_i^{\mathrm{snr}}, -0.2\right)$
VLT (vlt)	$u_i^{\mathrm{vlt}} = w_{\mathrm{mc}} y_i^{\mathrm{mc}} + w_{\mathrm{snr}} y_i^{\mathrm{snr}}$ $y_i^{\mathrm{vlt}} = F\left(a_i^{\mathrm{vlt}}, 0\right)$
MCtx (mc)	$u_i^{\mathrm{mc}} = w_{\mathrm{mc}} y_i^{\mathrm{sc}} + w_{\mathrm{vlt}} y_i^{\mathrm{vlt}}$ $y_i^{\mathrm{mc}} = F\left(a_i^{\mathrm{mc}}, 0\right)$

*Leaky integrator populations D1 and D2 are not used beyond model calibration

Table A.15: Basal ganglia–thalamocortical loop: Inputs

Input	**Targets**	
	STN (stn)	MCtx (mc)
Sensory cortex (sc)	$\mathrm{sc}_c^i \rightarrow \mathrm{stn}_c$, $1 \leq i \leq 500$	$\mathrm{sc}_c^i \rightarrow \mathrm{mc}_c$, $1 \leq i \leq 500$
	GPi/SNr (snr)	GPe (gp)
D_1 / D_2 MSNs (d1 / d2)	$\mathrm{d1}_c^i \rightarrow \mathrm{snr}_c$, $1 \leq i \leq 500$	$\mathrm{d2}_c^i \rightarrow \mathrm{gp}_c$, $1 \leq i \leq 500$

Table A.16: Basal ganglia–thalamocortical loop: Weights and properties

Param.	**Value**	**Source**
$w_{\text{sc–mc}}$	0.5	Tuning (see text)
$w_{\text{vlt–mc}}$	1.05	
$w_{\text{sc–d1}}$	0.5	Humphries and Gurney (2002)
$w_{\text{sc–d2}}$	0.5	
$w_{\text{sc–stn}}$	0.5	
$w_{\text{mc–d1}}$	0.5	
$w_{\text{mc–d2}}$	0.5	
$w_{\text{mc–stn}}$	0.5	
$w_{\text{mc–vlt}}$	1	
$w_{\text{d1–snr}}$	-1	
$w_{\text{d2–gp}}$	-1	
$w_{\text{stn–snr}}$	0.8	
$w_{\text{stn–gp}}$	0.8	
$w_{\text{gp–stn}}$	-1	
$w_{\text{gp–snr}}$	-0.4	
$w_{\text{snr–vlt}}$	-1	
$\chi_{\text{d1}'}$, $\chi_{\text{d2}'}$	2	
k	0.9608	
τ^r_{d1}, τ^r_{d2}, τ^r_{sc}	9 ms	Tuning (See text)
τ^f_{d1}, τ^f_{d2}, τ^f_{sc}	10 ms	
λ_{d1}, λ_{d2}	15	
λ_{sc}	850	
κ_{d1}, κ_{d2}	1	
κ_{sc}	1.5	

Bibliography

Agnoli, L., Mainolfi, P., Invernizzi, R. W., and Carli, M. (2013). Dopamine D1-Like and D2-Like Receptors in the Dorsal Striatum Control Different Aspects of Attentional Performance in the Five-Choice Serial Reaction Time Task Under a Condition of Increased Activity of Corticostriatal Inputs. *Neuropsychopharmacology*, 38(5):701–714.

Agostino, R., Berardelli, A., Formica, A., Accornero, N., and Manfredi, M. (1992). Sequential Arm Movements in Patients with Parkinson's Disease, Huntington's Disease and Dystonia. *Brain*, 115(5):1481–1495.

Aizman, O., Brismar, H., Uhlén, P., Zettergren, E., Levey, A. I., Forssberg, H., Greengard, P., and Aperia, A. (2000). Anatomical and physiological evidence for D1 and D2 dopamine receptor colocalization in neostriatal neurons. *Nature neuroscience*, 3(3).

Albin, R. L., Young, A. B., and Penney, J. B. (1989). The functional anatomy of basal ganglia disorders. *Trends in Neurosciences*, 12(10):366–375.

Alexander, G. E. and Crutcher, M. D. (1990). Functional architecture of basal ganglia circuits: neural substrates of parallel processing. *Trends in Neurosciences*, 13(7):266–271.

Alloway, K. D., Smith, J. B., Mowery, T. M., and Watson, G. D. R. (2017). Sensory Processing in the Dorsolateral Striatum: The Contribution of Thalamostriatal Pathways. *Frontiers in Systems Neuroscience*, 11.

Aosaki, T., Tsubokawa, H., Ishida, A., Watanabe, K., Graybiel, A. M., and Kimura, M. (1994). Responses of tonically active neurons in the primate's striatum undergo systematic changes during behavioral sensorimotor conditioning. *Journal of Neuroscience*, 14(6):3969–3984.

Apicella, P. (2007). Leading tonically active neurons of the striatum from reward detection to context recognition. *Trends in Neurosciences*, 30(6):299–306.

Aron, A. R. and Poldrack, R. A. (2006). Cortical and Subcortical Contributions to Stop Signal Response Inhibition: Role of the Subthalamic Nucleus. *Journal of Neuroscience*, 26(9):2424–2433.

Ashby, F. G. and Crossley, M. J. (2011). A Computational Model of How Cholinergic Interneurons Protect Striatal-dependent Learning. *Journal of Cognitive Neuroscience*, 23(6):1549–1566.

Averbeck, B. B., Chafee, M. V., Crowe, D. A., and Georgopoulos, A. P. (2003). Neural activity in prefrontal cortex during copying geometrical shapes. I. Single cells encode shape, sequence, and metric parameters. *Experimental Brain Research*, 150(2):127–141.

Averbeck, B. B., Sohn, J.-W., and Lee, D. (2006). Activity in prefrontal cortex during dynamic selection of action sequences. *Nature Neuroscience*, 9(2):276–282.

Bahuguna, J., Aertsen, A., and Kumar, A. (2015). Existence and Control of Go/No-Go Decision Transition Threshold in the Striatum. *PLoS Comput Biol*, 11(4):e1004233.

Balleine, B. W., Dezfouli, A., Ito, M., and Doya, K. (2015). Hierarchical control of goal-directed action in the cortical–basal ganglia network. *Current Opinion in Behavioral Sciences*, 5:1–7.

Balleine, B. W. and O'Doherty, J. P. (2010). Human and Rodent Homologies in Action Control: Corticostriatal Determinants of Goal-Directed and Habitual Action. *Neuropsychopharmacology*, 35(1):48–69.

Banghart, M. R., Neufeld, S. Q., Wong, N. C., and Sabatini, B. L. (2015). Enkephalin Disinhibits Mu Opioid Receptor-Rich Striatal Patches via Delta Opioid Receptors. *Neuron*.

Bar-Gad, I., Morris, G., and Bergman, H. (2003). Information processing, dimensionality reduction and reinforcement learning in the basal ganglia. *Progress in Neurobiology*, 71(6):439–473.

Bartos, M., Vida, I., and Jonas, P. (2007). Synaptic mechanisms of synchronized gamma oscillations in inhibitory interneuron networks. *Nature Reviews Neuroscience*, 8(1):45–56.

Baudic, S., Maison, P., Dolbeau, G., Boissé, M.-F., Bartolomeo, P., Dalla Barba, G., Traykov, L., and Bachoud-Lévi, A.-C. (2006). Cognitive Impairment Related to Apathy in Early Huntington's Disease. *Dementia and Geriatric Cognitive Disorders*, 21(5-6):316–321.

Bednark, J. G., Campbell, M. E. J., and Cunnington, R. (2015). Basal ganglia and cortical networks for sequential ordering and rhythm of complex movements. *Frontiers in Human Neuroscience*, page 421.

Beiser, D. G. and Houk, J. C. (1998). Model of Cortical-Basal Ganglionic Processing: Encoding the Serial Order of Sensory Events. *Journal of Neurophysiology*, 79(6):3168–3188.

Beiser, D. G., Hua, S. E., and Houk, J. C. (1997). Network models of the basal ganglia. *Current Opinion in Neurobiology*, 7(2):185–190.

Benecke, R., Rothwell, J. C., Dick, J. P. R., Day, B. L., and Marsden, C. D. (1987). Disturbance of Sequential Movements in Patients with Parkinson's Disease. *Brain*, 110(2):361–379.

Bennett, B. D. and Bolam, J. P. (1994). Synaptic input and output of parvalbumin-immunoreactive neurons in the neostriatum of the rat. *Neuroscience*, 62(3):707–719.

Berke, J. D. (2008). Uncoordinated Firing Rate Changes of Striatal Fast-Spiking Interneurons during Behavioral Task Performance. *Journal of Neuroscience*, 28(40):10075–10080.

Berke, J. D. (2009). Fast oscillations in cortical-striatal networks switch frequency following rewarding events and stimulant drugs. *European Journal of Neuroscience*, 30(5):848–859.

Berns, G. S. and Sejnowski, T. J. (1996). How the Basal Ganglia Make Decisions. In Damasio, A., Damasio, H., and Christen, Y., editors, *Neurobiology of Decision-Making*, Research and Perspectives in Neurosciences, pages 101–113. Springer, Berlin, Heidelberg.

Berns, G. S. and Sejnowski, T. J. (1998). A Computational Model of How the Basal Ganglia Produce Sequences. *Journal of Cognitive Neuroscience*, 10(1):108–121.

Biezonski, D. K., Trifilieff, P., Meszaros, J., Javitch, J. A., and Kellendonk, C. (2015). Evidence for limited D1 and D2 receptor coexpression and colocalization within the dorsal striatum of the neonatal mouse. *Journal of Comparative Neurology*, 523(8):1175–1189.

Björklund, A. and Dunnett, S. B. (2007). Dopamine neuron systems in the brain: an update. *Trends in Neurosciences*, 30(5):194–202.

Blackwell, K. T., Czubayko, U., and Plenz, D. (2003). Quantitative Estimate of Synaptic Inputs to Striatal Neurons during Up and Down States In Vitro. *Journal of Neuroscience*, 23(27):9123–9132.

Blazquez, P. M., Fujii, N., Kojima, J., and Graybiel, A. M. (2002). A Network Representation of Response Probability in the Striatum. *Neuron*, 33(6):973–982.

Blomeley, C. and Bracci, E. (2008). Substance P depolarizes striatal projection neurons and facilitates their glutamatergic inputs. *The Journal of Physiology*, 586(8):2143–2155.

Blomeley, C. P. and Bracci, E. (2011). Opioidergic Interactions between Striatal Projection Neurons. *The Journal of Neuroscience*, 31(38):13346–13356.

Blomeley, C. P., Kehoe, L. A., and Bracci, E. (2009). Substance P Mediates Excitatory Interactions between Striatal Projection Neurons. *The Journal of Neuroscience*, 29(15):4953–4963.

Bogacz, R., Wagenmakers, E.-J., Forstmann, B. U., and Nieuwenhuis, S. (2010). The neural basis of the speed–accuracy tradeoff. *Trends in Neurosciences*, 33(1):10–16.

Bokil, H., Andrews, P., Kulkarni, J. E., Mehta, S., and Mitra, P. (2010). Chronux: A Platform for Analyzing Neural Signals. *Journal of neuroscience methods*, 192(1):146–151.

Bolam, J. P., Bergman, H., Graybiel, A. M., Kimura, M., Plenz, D., Seung, H. S., Surmeier, D. J., and Wickens, J. R. (2006). Microcircuits in the Striatum. In Grillner, S. and Graybiel, A. M., editors, *Microcircuits: The Interface Between Neurons and Global Brain Function*, pages 165–190. MIT Press, Cambridge, MA.

Bolam, J. P., Hanley, J. J., Booth, P. a. C., and Bevan, M. D. (2000). Synaptic organisation of the basal ganglia. *The Journal of Anatomy*, 196(4):527–542.

Bolam, J. P., Izzo, P. N., and Graybiel, A. M. (1988). Cellular substrate of the histochemically defined striosome/matrix system of the caudate nucleus: A combined golgi and immunocytochemical study in cat and ferret. *Neuroscience*, 24(3):853–875.

Botvinick, M. M., Niv, Y., and Barto, A. C. (2009). Hierarchically organized behavior and its neural foundations: A reinforcement learning perspective. *Cognition*, 113(3):262–280.

Bracci, E., Centonze, D., Bernardi, G., and Calabresi, P. (2002). Dopamine Excites Fast-Spiking Interneurons in the Striatum. *Journal of Neurophysiology*, 87(4):2190–2194.

Brimblecombe, K. R. and Cragg, S. J. (2015). Substance P Weights Striatal Dopamine Transmission Differently within the Striosome-Matrix Axis. *The Journal of Neuroscience*, 35(24):9017–9023.

Brog, J. S., Salyapongse, A., Deutch, A. Y., and Zahm, D. S. (1993). The patterns of afferent innervation of the core and shell in the "Accumbens" part of the rat ventral striatum: Immunohistochemical detection of retrogradely transported fluoro-gold. *Journal of Comparative Neurology*, 338(2):255–278.

Bromberg-Martin, E. S., Matsumoto, M., and Hikosaka, O. (2010). Dopamine in Motivational Control: Rewarding, Aversive, and Alerting. *Neuron*, 68(5):815–834.

Brown, L. L., Smith, D. M., and Goldbloom, L. M. (1998). Organizing principles of cortical integration in the rat neostriatum: Corticostriate map of the body surface is an ordered lattice of curved laminae and radial points. *The Journal of Comparative Neurology*, 392(4):468–488.

Buxton, D., Bracci, E., Overton, P. G., and Gurney, K. (2017). Striatal Neuropeptides Enhance Selection and Rejection of Sequential Actions. *Frontiers in Computational Neuroscience*, 11.

Buzsáki, G. and Chrobak, J. J. (1995). Temporal structure in spatially organized neuronal ensembles: a role for interneuronal networks. *Current Opinion in Neurobiology*, 5(4):504–510.

Buzsáki, G. and Wang, X.-J. (2012). Mechanisms of Gamma Oscillations. *Annual Review of Neuroscience*, 35(1):203–225.

Cador, M., Kelley, A. E., Le Moal, M., and Stinus, L. (1986). Ventral tegmental area infusion of substance P, neurotensin and enkephalin: Differential effects on feeding behavior. *Neuroscience*, 18(3):659–669.

Calabresi, P., Gubellini, P., Centonze, D., Sancesario, G., Morello, M., Giorgi, M., Pisani, A., and Bernardi, G. (1999). A Critical Role of the Nitric Oxide/cGMP Pathway in Corticostriatal Long-Term Depression. *Journal of Neuroscience*, 19(7):2489–2499.

Calaprice, A. (2010). *The Ultimate Quotable Einstein*. Princeton University Press, New Jersey.

Cardinal, R. N., Parkinson, J. A., Hall, J., and Everitt, B. J. (2002). Emotion and motivation: the role of the amygdala, ventral striatum, and prefrontal cortex. *Neuroscience & Biobehavioral Reviews*, 26(3):321–352.

Carmichael, J. E., Gmaz, J. M., and Meer, M. A. A. v. d. (2017). Gamma Oscillations in the Rat Ventral Striatum Originate in the Piriform Cortex. *Journal of Neuroscience*, 37(33):7962–7974.

Carrillo-Reid, L., Hernández-López, S., Tapia, D., Galarraga, E., and Bargas, J. (2011). Dopaminergic Modulation of the Striatal Microcircuit: Receptor-Specific Configuration of Cell Assemblies. *The Journal of Neuroscience*, 31(42):14972–14983.

Carrillo-Reid, L., Tecuapetla, F., Tapia, D., Hernández-Cruz, A., Galarraga, E., Drucker-Colin, R., and Bargas, J. (2008). Encoding Network States by Striatal Cell Assemblies. *Journal of Neurophysiology*, 99(3):1435–1450.

Centonze, D., Grande, C., Usiello, A., Gubellini, P., Erbs, E., Martín, A. B., Pisani, A., Tognazzi, N., Bernardi, G., Moratalla, R., Borrelli, E., and Calabresi, P. (2003). Receptor Subtypes Involved in the Presynaptic and Postsynaptic Actions of Dopamine on Striatal Interneurons. *The Journal of Neuroscience*, 23(15):6245–6254.

Chambers, J. M., Gurney, K., Humphries, M., and Prescott, A. (2005). Mechanisms of choice in the primate brain: a quick look at positive feedback. In *Modelling Natural Action Selection: Proceedings of an International Workshop*, pages 45–52, Edinburgh, UK. Citeseer.

Chersi, F., Mirolli, M., Pezzulo, G., and Baldassarre, G. (2013). A spiking neuron model of the cortico-basal ganglia circuits for goal-directed and habitual action learning. *Neural Networks*, 41:212–224.

Clements, J. D., Lester, R. A., Tong, G., Jahr, C. E., and Westbrook, G. L. (1992). The time course of glutamate in the synaptic cleft. *Science*, 258(5087):1498–1501.

Coffey, K. R., Nader, M., and West, M. O. (2016). Single body parts are processed by individual neurons in the mouse dorsolateral striatum. *Brain Research*, 1636:200–207.

Cohen, M. X., Axmacher, N., Lenartz, D., Elger, C. E., Sturm, V., and Schlaepfer, T. E. (2008). Good Vibrations: Cross-frequency Coupling in the Human Nucleus Accumbens during Reward Processing. *Journal of Cognitive Neuroscience*, 21(5):875–889.

Cope, A. J., Richmond, P., James, S. S., Gurney, K., and Allerton, D. J. (2017). SpineCreator: a Graphical User Interface for the Creation of Layered Neural Models. *Neuroinformatics*, 15(1):25–40.

Crittenden, J. R., Lacey, C. J., Weng, F.-J., Garrison, C. E., Gibson, D. J., Lin, Y., and Graybiel, A. M. (2017). Striatal Cholinergic Interneurons Modulate Spike-Timing in Striosomes and Matrix by an Amphetamine-Sensitive Mechanism. *Frontiers in Neuroanatomy*, 11.

Cui, G., Jun, S. B., Jin, X., Pham, M. D., Vogel, S. S., Lovinger, D. M., and Costa, R. M. (2013). Concurrent activation of striatal direct and indirect pathways during action initiation. *Nature*, 494(7436):238–242.

de Tommaso, M., Specchio, N., Sciruicchio, V., Difruscolo, O., and Specchio, L. M. (2004). Effects of rivastigmine on motor and cognitive impairment in Huntington's disease. *Movement Disorders*, 19(12):1516–1518.

Dehaene, S., Meyniel, F., Wacongne, C., Wang, L., and Pallier, C. (2015). The Neural Representation of Sequences: From Transition Probabilities to Algebraic Patterns and Linguistic Trees. *Neuron*, 88(1):2–19.

DeLong, M. R. (1990). Primate models of movement disorders of basal ganglia origin. *Trends in Neurosciences*, 13(7):281–285.

Deniau, J. M., Mailly, P., Maurice, N., and Charpier, S. (2007). The pars reticulata of the substantia nigra: a window to basal ganglia output. *Progress in Brain Research*, 160:151–172.

Devan, B. D., Hong, N. S., and McDonald, R. J. (2011). Parallel associative processing in the dorsal striatum: Segregation of stimulus–response and cognitive control subregions. *Neurobiology of Learning and Memory*, 96(2):95–120.

Dezfouli, A. and 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.

Dezfouli, A., Lingawi, N. W., and Balleine, B. W. (2014). Habits as action sequences: hierarchical action control and changes in outcome value. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 369(1655):20130482.

Dickinson, A. (1985). Actions and habits: the development of behavioural autonomy. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 308(1135):67–78.

Ding, J. B., Guzman, J. N., Peterson, J. D., Goldberg, J. A., and Surmeier, D. J. (2010). Thalamic Gating of Corticostriatal Signaling by Cholinergic Interneurons. *Neuron*, 67(2):294–307.

Do, J., Kim, J.-I., Bakes, J., Lee, K., and Kaang, B.-K. (2013). Functional roles of neurotransmitters and neuromodulators in the dorsal striatum. *Learning & Memory*, 20(1):21–28.

Doig, N. M., Moss, J., and Bolam, J. P. (2010). Cortical and Thalamic Innervation of Direct and Indirect Pathway Medium-Sized Spiny Neurons in Mouse Striatum. *Journal of Neuroscience*, 30(44):14610–14618.

Domenger, D. and Schwarting, R. K. W. (2006). The serial reaction time task in the rat: Effects of D1 and D2 dopamine-receptor antagonists. *Behavioural Brain Research*, 175(2):212–222.

Dommett, E., Coizet, V., Blaha, C. D., Martindale, J., Lefebvre, V., Walton, N., Mayhew, J. E. W., Overton, P. G., and Redgrave, P. (2005). How Visual Stimuli Activate Dopaminergic Neurons at Short Latency. *Science*, 307(5714):1476–1479.

Donoghue, J. P. and Herkenham, M. (1986). Neostriatal projections from individual cortical fields conform to histochemically distinct striatal compartments in the rat. *Brain Research*, 365(2):397–403.

Doya, K. (1999). What are the computations of the cerebellum, the basal ganglia and the cerebral cortex? *Neural Networks*, 12(7–8):961–974.

Draganski, B., Kherif, F., Kloppel, S., Cook, P. A., Alexander, D. C., Parker, G. J. M., Deichmann, R., Ashburner, J., and Frackowiak, R. S. J. (2008). Evidence for Segregated and Integrative Connectivity Patterns in the Human Basal Ganglia. *Journal of Neuroscience*, 28(28):7143–7152.

Duff, K., Paulsen, J. S., Beglinger, L. J., Langbehn, D. R., Wang, C., Stout, J. C., Ross, C. A., Aylward, E., Carlozzi, N. E., Queller, S., and others (2010). "Frontal" behaviors before the diagnosis of Huntington's disease and their relationship to markers of disease progression: evidence of early lack of awareness. *The Journal of neuropsychiatry and clinical neurosciences*, 22(2):196–207.

Durieux, P. F., Bearzatto, B., Guiducci, S., Buch, T., Waisman, A., Zoli, M., Schiffmann, S. N., and de Kerchove d'Exaerde, A. (2009). D2r striatopallidal neurons inhibit both locomotor and drug reward processes. *Nature Neuroscience*, 12(4):393–395.

Durston, S., Tottenham, N. T., Thomas, K. M., Davidson, M. C., Eigsti, I.-M., Yang, Y., Ulug, A. M., and Casey, B. J. (2003). Differential patterns of striatal activation in young children with and without ADHD. *Biological Psychiatry*, 53(10):871–878.

El Massioui, N., Lamirault, C., Yagüe, S., Adjeroud, N., Garces, D., Maillard, A., Tallot, L., Yu-Taeger, L., Riess, O., Allain, P., Nguyen, H. P., von Hörsten, S., and Doyère, V. (2016).

Impaired Decision Making and Loss of Inhibitory-Control in a Rat Model of Huntington Disease. *Frontiers in Behavioral Neuroscience*, 10.

Epstein, J. M. (2008). Why Model? *Journal of Artificial Societies and Social Simulation*, 11(4):12.

Fields, H. L., Hjelmstad, G. O., Margolis, E. B., and Nicola, S. M. (2007). Ventral Tegmental Area Neurons in Learned Appetitive Behavior and Positive Reinforcement. *Annual Review of Neuroscience*, 30(1):289–316.

Flaherty, A. W. and Graybiel, A. M. (1993). Output architecture of the primate putamen. *The Journal of Neuroscience*, 13(8):3222–3237.

Forstmann, B. U., Dutilh, G., Brown, S., Neumann, J., Cramon, D. Y. v., Ridderinkhof, K. R., and Wagenmakers, E.-J. (2008). Striatum and pre-SMA facilitate decision-making under time pressure. *Proceedings of the National Academy of Sciences*, 105(45):17538–17542.

François, C., Percheron, G., Parent, A., Sadikot, A. F., Fenelon, G., and Yelnik, J. (1991). Topography of the projection from the central complex of the thalamus to the sensorimotor striatal territory in monkeys. *Journal of Comparative Neurology*, 305(1):17–34.

Frank, M. J. (2006). Hold your horses: A dynamic computational role for the subthalamic nucleus in decision making. *Neural Networks*, 19(8):1120–1136.

Franklin, N. T. and Frank, M. J. (2015). A cholinergic feedback circuit to regulate striatal population uncertainty and optimize reinforcement learning. *eLife*, 4.

Freeze, B. S., Kravitz, A. V., Hammack, N., Berke, J. D., and Kreitzer, A. C. (2013). Control of Basal Ganglia Output by Direct and Indirect Pathway Projection Neurons. *The Journal of Neuroscience*, 33(47):18531–18539.

Friedman, A., Homma, D., Gibb, L. G., Amemori, K.-i., Rubin, S. J., Hood, A. S., Riad, M. H., and Graybiel, A. M. (2015). A Corticostriatal Path Targeting Striosomes Controls Decision-Making under Conflict. *Cell*, 161(6):1320–1333.

Fujii, N. and Graybiel, A. M. (2003). Representation of Action Sequence Boundaries by Macaque Prefrontal Cortical Neurons. *Science*, 301(5637):1246–1249.

Fukai, T. (1999). Sequence generation in arbitrary temporal patterns from theta-nested gamma oscillations: a model of the basal ganglia–thalamo-cortical loops. *Neural Networks*, 12(7):975–987.

Galarreta, M. and Hestrin, S. (1999). A network of fast-spiking cells in the neocortex connected by electrical synapses. *Nature*, 402(6757):72–75.

Gerfen, C. R. (1984). The neostriatal mosaic: compartmentalization of corticostriatal input and striatonigral output systems. *Nature*, 311(5985):461–464.

Gerfen, C. R. and Surmeier, D. J. (2011). Modulation of striatal projection systems by dopamine. *Annual review of neuroscience*, 34:441–466.

Gerfen, C. R. and Wilson, C. J. (1996). Chapter II The basal ganglia. *Handbook of Chemical Neuroanatomy*, 12:371–468.

Gertler, T. S., Chan, C. S., and Surmeier, D. J. (2008). Dichotomous Anatomical Properties of Adult Striatal Medium Spiny Neurons. *Journal of Neuroscience*, 28(43):10814–10824.

Gillies, A. and Arbuthnott, G. (2000). Computational models of the basal ganglia. *Movement disorders : official journal of the Movement Disorder Society*, 15(5):762–770.

Gittis, A. H. and Kreitzer, A. C. (2012). Striatal microcircuitry and movement disorders. *Trends in Neurosciences*, 35(9):557–564.

Gittis, A. H., Leventhal, D. K., Fensterheim, B. A., Pettibone, J. R., Berke, J. D., and Kreitzer, A. C. (2011). Selective Inhibition of Striatal Fast-Spiking Interneurons Causes Dyskinesias. *Journal of Neuroscience*, 31(44):15727–15731.

Gittis, A. H., Nelson, A. B., Thwin, M. T., Palop, J. J., and Kreitzer, A. C. (2010). Distinct Roles of GABAergic Interneurons in the Regulation of Striatal Output Pathways. *The Journal of Neuroscience*, 30(6):2223–2234.

Gobel, E. W., Sanchez, D. J., and Reber, P. J. (2011). Integration of Temporal and Ordinal Information During Serial Interception Sequence Learning. *Journal of Experimental Psychology*, 37(4):994–1000.

Gonzalez-Nicolini, V. and McGinty, J. F. (2002). NK-1 receptor blockade decreases amphetamine-induced behavior and neuropeptide mRNA expression in the striatum. *Brain Research*, 931(1):41–49.

Gorelova, N., Seamans, J. K., and Yang, C. R. (2002). Mechanisms of Dopamine Activation of Fast-Spiking Interneurons That Exert Inhibition in Rat Prefrontal Cortex. *Journal of Neurophysiology*, 88(6):3150–3166.

Govindaiah, G., Wang, Y., and Cox, C. L. (2010). Substance P selectively modulates GABAA receptor-mediated synaptic transmission in striatal cholinergic interneurons. *Neuropharmacology*, 58(2):413–422.

Grace, A. A. (1991). The cortical regulation of dopamine system responsivity: A hypothesis regarding its role in the etiology of schizophrenia. *Schizophrenia Research*, 4(3):345.

Graveland, G. A. and Difiglia, M. (1985). The frequency and distribution of medium-sized neurons with indented nuclei in the primate and rodent neostriatum. *Brain Research*, 327(1):307–311.

Gray, C. M. and Singer, W. (1989). Stimulus-specific neuronal oscillations in orientation columns of cat visual cortex. *Proceedings of the National Academy of Sciences*, 86(5):1698–1702.

Graybiel, A. M. (1998). The Basal Ganglia and Chunking of Action Repertoires. *Neurobiology of Learning and Memory*, 70(1–2):119–136.

Graybiel, A. M. (2008). Habits, Rituals, and the Evaluative Brain. *Annual Review of Neuroscience*, 31(1):359–387.

Graybiel, A. M. and Grafton, S. T. (2015). The Striatum: Where Skills and Habits Meet. *Cold Spring Harbor Perspectives in Biology*, 7(8):a021691.

Graybiel, A. M. and Ragsdale, C. W. (1978). Histochemically distinct compartments in the striatum of human, monkeys, and cat demonstrated by acetylthiocholinesterase staining. *Proceedings of the National Academy of Sciences*, 75(11):5723–5726.

Grillner, S., Markram, H., De Schutter, E., Silberberg, G., and LeBeau, F. E. N. (2005). Microcircuits in action – from CPGs to neocortex. *Trends in Neurosciences*, 28(10):525–533.

Grillner, S. and Robertson, B. (2016). The Basal Ganglia Over 500 Million Years. *Current Biology*, 26(20):R1088–R1100.

Groves, P. M. (1983). A theory of the functional organization of the neostriatum and the neostriatal control of voluntary movement. *Brain Research Reviews*, 5(2):109–132.

Gruber, A. J., Solla, S. A., Surmeier, D. J., and Houk, J. C. (2003). Modulation of Striatal Single Units by Expected Reward: A Spiny Neuron Model Displaying Dopamine-Induced Bistability. *Journal of Neurophysiology*, 90(2):1095–1114.

Gurney, K., Prescott, T. J., and Redgrave, P. (2001a). A computational model of action selection in the basal ganglia. I. A new functional anatomy. *Biological cybernetics*, 84(6):401–410.

Gurney, K., Prescott, T. J., and Redgrave, P. (2001b). A computational model of action selection in the basal ganglia. II. Analysis and simulation of behaviour. *Biological cybernetics*, 84(6):411–423.

Gurney, K. N. and Overton, P. G. (2004). A model of short and long range selective processing in neostriatum. *Neurocomputing*, 58(Supplement C):555–562.

Haber, S. and McFarland, N. R. (2001). The Place of the Thalamus in Frontal Cortical-Basal Ganglia Circuits. *The Neuroscientist*, 7(4):315–324.

Haber, S. N. and Calzavara, R. (2009). The cortico-basal ganglia integrative network: The role of the thalamus. *Brain Research Bulletin*, 78(2):69–74.

Hamilton, J. M., Salmon, D. P., Corey-Bloom, J., Gamst, A., Paulsen, J. S., Jerkins, S., Jacobson, M. W., and Peavy, G. (2003). Behavioural abnormalities contribute to functional decline in Huntington's disease. *Journal of Neurology, Neurosurgery & Psychiatry*, 74(1):120–122.

Hasenöhrl, R. U., Souza-Silva, M. A. D., Nikolaus, S., Tomaz, C., Brandao, M. L., Schwarting, R. K. W., and Huston, J. P. (2000). Substance P and its role in neural mechanisms governing learning, anxiety and functional recovery. *Neuropeptides*, 34(5):272–280.

Helie, S., Chakravarthy, S., and Moustafa, A. A. (2013). Exploring the cognitive and motor functions of the basal ganglia: an integrative review of computational cognitive neuroscience models. *Frontiers in Computational Neuroscience*, 7.

Herkenham, M. and Pert, C. B. (1981). Mosaic distribution of opiate receptors, parafascicular projections and acetylcholinesterase in rat striatum. *Nature*, 291(5814):415–418.

Hernández-López, S., Bargas, J., Surmeier, D. J., Reyes, A., and Galarraga, E. (1997). D1 Receptor Activation Enhances Evoked Discharge in Neostriatal Medium Spiny Neurons by Modulating an L-Type Ca2+ Conductance. *Journal of Neuroscience*, 17(9):3334–3342.

Higley, M. J. and Sabatini, B. L. (2010). Competitive regulation of synaptic Ca2+ influx by D2 dopamine and A2a adenosine receptors. *Nature Neuroscience*, 13(8):958–966.

Hikida, T., Kimura, K., Wada, N., Funabiki, K., and Nakanishi, S. (2010). Distinct Roles of Synaptic Transmission in Direct and Indirect Striatal Pathways to Reward and Aversive Behavior. *Neuron*, 66(6):896–907.

Hikosaka, O., Ghazizadeh, A., Griggs, W., and Amita, H. (2017). Parallel basal ganglia circuits for decision making. *Journal of Neural Transmission*, pages 1–15.

Hjorth, J., Blackwell, K. T., and Kotaleski, J. H. (2009). Gap Junctions between Striatal Fast-Spiking Interneurons Regulate Spiking Activity and Synchronization as a Function of Cortical Activity. *The Journal of Neuroscience*, 29(16):5276–5286.

Hodgkin, A. L. and Huxley, A. F. (1952). A quantitative description of membrane current and its application to conduction and excitation in nerve. *The Journal of Physiology*, 117(4):500–544.

Hoover, J. E. and Strick, P. L. (1993). Multiple output channels in the basal ganglia. *Science*, 259(5096):819–821.

Humphries, M. D. and Gurney, K. (2007). Solution Methods for a New Class of Simple Model Neurons. *Neural Computation*, 19(12):3216–3225.

Humphries, M. D. and Gurney, K. N. (2002). The role of intra-thalamic and thalamocortical circuits in action selection. *Network: Computation in Neural Systems*, 13(1):131–156.

Humphries, M. D., Lepora, N., Wood, R., and Gurney, K. (2009a). Capturing dopaminergic modulation and bimodal membrane behaviour of striatal medium spiny neurons in accurate, reduced models. *Frontiers in Computational Neuroscience*, 3.

Humphries, M. D., Stewart, R. D., and Gurney, K. N. (2006). A Physiologically Plausible Model of Action Selection and Oscillatory Activity in the Basal Ganglia. *The Journal of Neuroscience*, 26(50):12921–12942.

Humphries, M. D., Wood, R., and Gurney, K. (2009b). Dopamine-modulated dynamic cell assemblies generated by the GABAergic striatal microcircuit. *Neural Networks*, 22(8):1174–1188.

Humphries, M. D., Wood, R., and Gurney, K. (2010). Reconstructing the Three-Dimensional GABAergic Microcircuit of the Striatum. *PLoS Computational Biology*, 6(11):e1001011.

Hunnicutt, B. J., Jongbloets, B. C., Birdsong, W. T., Gertz, K. J., Zhong, H., and Mao, T. (2016). A comprehensive excitatory input map of the striatum reveals novel functional organization. *eLife*, 5.

Huston, J. P. and Hasenöhrl, R. U. (1995). The role of neuropeptides in learning: focus on the neurokinin substance P. *Behavioural Brain Research*, 66(1–2):117–127.

Isomura, Y., Takekawa, T., Harukuni, R., Handa, T., Aizawa, H., Takada, M., and Fukai, T. (2013). Reward-Modulated Motor Information in Identified Striatum Neurons. *Journal of Neuroscience*, 33(25):10209–10220.

Izhikevich, E. M. (2003). Simple model of spiking neurons. *IEEE Transactions on neural networks*, 14(6):1569–1572.

Izhikevich, E. M. (2007). *Dynamical Systems in Neuroscience*. MIT Press, Cambridge, MA.

Jaeger, D., Kita, H., and Wilson, C. J. (1994). Surround inhibition among projection neurons is weak or nonexistent in the rat neostriatum. *Journal of Neurophysiology*, 72(5):2555–2558.

Jahr, C. E. and Stevens, C. F. (1990). A quantitative description of NMDA receptor-channel kinetic behavior. *The Journal of Neuroscience*, 10(6):1830–1837.

Jáidar, O., Carrillo-Reid, L., Hernández, A., Drucker-Colín, R., Bargas, J., and Hernández-Cruz, A. (2010). Dynamics of the Parkinsonian Striatal Microcircuit: Entrainment into a Dominant Network State. *Journal of Neuroscience*, 30(34):11326–11336.

James, T. A. and Starr, M. S. (1977). Behavioural and biochemical effects of substance P injected into the substantia nigra of the rat. *Journal of Pharmacy and Pharmacology*, 29(1):181–182.

Jin, X. and Costa, R. M. (2010). Start/stop signals emerge in nigrostriatal circuits during sequence learning. *Nature*, 466(7305):457–462.

Jin, X., Tecuapetla, F., and Costa, R. M. (2014). Basal ganglia subcircuits distinctively encode the parsing and concatenation of action sequences. *Nature Neuroscience*, 17(3):423–430.

Kalenscher, T., Lansink, C. S., Lankelma, J. V., and Pennartz, C. M. A. (2010). Reward-Associated Gamma Oscillations in Ventral Striatum Are Regionally Differentiated and Modulate Local Firing Activity. *Journal of Neurophysiology*, 103(3):1658–1672.

Kalivas, P. W., Widerlöv, E., Stanley, D., Breese, G., and Prange, A. J. (1983). Enkephalin action on the mesolimbic system: a dopamine-dependent and a dopamine-independent increase in locomotor activity. *Journal of Pharmacology and Experimental Therapeutics*, 227(1):229–237.

Kalkhoven, C., Sennef, C., Peeters, A., and van den Bos, R. (2014). Risk-taking and pathological gambling behavior in Huntington's disease. *Frontiers in Behavioral Neuroscience*, 8.

Kaneda, K., Nambu, A., Tokuno, H., and Takada, M. (2002). Differential Processing Patterns of Motor Information Via Striatopallidal and Striatonigral Projections. *Journal of Neurophysiology*, 88(3):1420–1432.

Kasanetz, F., Riquelme, L. A., O'Donnell, P., and Murer, M. G. (2006). Turning off cortical ensembles stops striatal Up states and elicits phase perturbations in cortical and striatal slow oscillations in rat in vivo. *The Journal of Physiology*, 577(1):97–113.

Kawaguchi, Y. (1993). Physiological, morphological, and histochemical characterization of three classes of interneurons in rat neostriatum. *Journal of Neuroscience*, 13(11):4908–4923.

Kelley, A. E. (2004). Ventral striatal control of appetitive motivation: role in ingestive behavior and reward-related learning. *Neuroscience & Biobehavioral Reviews*, 27(8):765–776.

Kelley, A. E., Cador, M., Stinus, L., and Le Moal, M. (1989). Neurotensin, substance P, neurokinin-α, and enkephalin: injection into ventral tegmental area in the rat produces differential effects on operant responding. *Psychopharmacology*, 97(2):243–252.

Kelley, A. E. and Iversen, S. D. (1979). Substance P infusion into substantia nigra of the rat: Behavioural analysis and involvement of striatal dopamine. *European Journal of Pharmacology*, 60(2–3):171–179.

Kemp, J. M. (1968). Observations on the caudate nucleus of the cat impregnated with the Golgi method. *Brain Research*, 11(2):467–470.

Keramati, M., Dezfouli, A., and Piray, P. (2011). Speed/Accuracy Trade-Off between the Habitual and the Goal-Directed Processes. *PLoS Computational Biology*, 7(5):e1002055.

Kim, H. F. and Hikosaka, O. (2015). Parallel basal ganglia circuits for voluntary and automatic behaviour to reach rewards. *Brain*, 138(7):1776–1800.

Kishore, A., Panikar, D., Balakrishnan, S., Joseph, S., and Sarma, S. (2000). Evidence of functional somatotopy in GPi from results of pallidotomy. *Brain*, 123(12):2491–2500.

Kita, H., Kita, T., and Kitai, S. T. (1985). Active membrane properties of rat neostriatal neurons in an in vitro slice preparation. *Experimental Brain Research*, 60(1):54–62.

Kita, H., Kosaka, T., and Heizmann, C. W. (1990). Parvalbumin-immunoreactive neurons in the rat neostriatum: a light and electron microscopic study. *Brain Research*, 536(1–2):1–15.

Kitai, S. T. and Surmeier, D. J. (1993). Cholinergic and dopaminergic modulation of potassium conductances in neostriatal neurons. *Advances in neurology*, 60:40–52.

Klaus, A., Martins, G. J., Paixao, V. B., Zhou, P., Paninski, L., and Costa, R. M. (2017). The Spatiotemporal Organization of the Striatum Encodes Action Space. *Neuron*, 95(5):1171–1180.e7.

Knutson, B., Delgado, M. R., and Phillips, P. E. M. (2009). Representation of subjective value in the striatum. In Glimcher, P. W., Camerer, C. F., Fehr, E., and Poldrack, R. A., editors, *Neuroeconomics: Decision Making and the Brain*, pages 389–406. Academic Press.

Koechlin, E. and Burnod, Y. (1996). Dual Population Coding in the Neocortex: A Model of Interaction between Representation and Attention in the Visual Cortex. *Journal of Cognitive Neuroscience*, 8(4):353–370.

Koós, T. and Tepper, J. M. (1999). Inhibitory control of neostriatal projection neurons by GABAergic interneurons. *Nature Neuroscience*, 2(5):467–472.

Koós, T., Tepper, J. M., and Wilson, C. J. (2004). Comparison of IPSCs Evoked by Spiny and Fast-Spiking Neurons in the Neostriatum. *The Journal of Neuroscience*, 24(36):7916–7922.

Kotaleski, J. H., Plenz, D., and Blackwell, K. T. (2006). Using potassium currents to solve signal-to-noise problems in inhibitory feedforward networks of the striatum. *Journal of Neurophysiology, Journal of neurophysiology*, 95, 95(1, 1):331, 331–341.

Kravitz, A. V., Freeze, B. S., Parker, P. R. L., Kay, K., Thwin, M. T., Deisseroth, K., and Kreitzer, A. C. (2010). Regulation of parkinsonian motor behaviours by optogenetic control of basal ganglia circuitry. *Nature*, 466(7306):622–626.

Kravitz, A. V., Tye, L. D., and Kreitzer, A. C. (2012). Distinct roles for direct and indirect pathway striatal neurons in reinforcement. *Nature neuroscience*, 15(6):816–818.

Kreitzer, A. C. (2009). Physiology and Pharmacology of Striatal Neurons. *Annual Review of Neuroscience*, 32(1):127–147.

Kreitzer, A. C. and Malenka, R. C. (2007). Endocannabinoid-mediated rescue of striatal LTD and motor deficits in Parkinson's disease models. *Nature*, 445(7128):643–647.

Laitinen, L. V., Bergenheim, A. T., and Hariz, M. I. (1992). Ventroposterolateral Pallidotomy Can Abolish All Parkinsonian Symptoms. *Stereotactic and Functional Neurosurgery*, 58(1-4):14–21.

Lanciego, J. L., Luquin, N., and Obeso, J. A. (2012). Functional Neuroanatomy of the Basal Ganglia. *Cold Spring Harbor Perspectives in Medicine*, 2(12):a009621.

Lapper, S. R. and Bolam, J. P. (1992). Input from the frontal cortex and the parafascicular nucleus to cholinergic interneurons in the dorsal striatum of the rat. *Neuroscience*, 51(3):533–545.

Lapper, S. R., Smith, Y., Sadikot, A. F., Parent, A., and Bolam, J. P. (1992). Cortical input to parvalbumin-immunoreactive neurones in the putamen of the squirrel monkey. *Brain Research*, 580(1):215–224.

Lehéricy, S., Bardinet, E., Tremblay, L., Van de Moortele, P.-F., Pochon, J.-B., Dormont, D., Kim, D.-S., Yelnik, J., and Ugurbil, K. (2006). Motor control in basal ganglia circuits using fMRI and brain atlas approaches. *Cerebral Cortex*, 16(2):149–161.

Lei, W., Jiao, Y., Mar, N. D., and Reiner, A. (2004). Evidence for Differential Cortical Input to Direct Pathway versus Indirect Pathway Striatal Projection Neurons in Rats. *The Journal of Neuroscience*, 24(38):8289–8299.

London, T. D., Licholai, J. A., Szczot, I., Ali, M. A., LeBlanc, K. H., Fobbs, W. C., and Kravitz, A. V. (2018). Coordinated Ramping of Dorsal Striatal Pathways preceding Food Approach and Consumption. *Journal of Neuroscience*, 38(14):3547–3558.

Luk, K. C. and Sadikot, A. F. (2001). GABA promotes survival but not proliferation of parvalbumin-immunoreactive interneurons in rodent neostriatum: an in vivo study with stereology. *Neuroscience*, 104(1):93–103.

Magill, P. J., Bolam, J. P., and Bevan, M. D. (2001). Dopamine regulates the impact of the cerebral cortex on the subthalamic nucleus–globus pallidus network. *Neuroscience*, 106(2):313–330.

Mahon, S., Deniau, J.-M., Charpier, S., and Delord, B. (2000). Role of a Striatal Slowly Inactivating Potassium Current in Short-Term Facilitation of Corticostriatal Inputs: A Computer Simulation Study. *Learning & Memory*, 7(5):357–362.

Mana, S. and Chevalier, G. (2001). The fine organization of nigro-collicular channels with additional observations of their relationships with acetylcholinesterase in the rat. *Neuroscience*, 106(2):357–374.

Mannella, F., Gurney, K., and Baldassarre, G. (2013). The nucleus accumbens as a nexus between values and goals in goal-directed behavior: a review and a new hypothesis. *Frontiers in Behavioral Neuroscience*, 7.

Mansfield, E. L., Karayanidis, F., Jamadar, S., Heathcote, A., and Forstmann, B. U. (2011). Adjustments of Response Threshold during Task Switching: A Model-Based Functional Magnetic Resonance Imaging Study. *Journal of Neuroscience*, 31(41):14688–14692.

Martiros, N., Burgess, A. A., and Graybiel, A. M. (2018). Inversely Active Striatal Projection Neurons and Interneurons Selectively Delimit Useful Behavioral Sequences. *Current Biology*, 28(4):560–573.e5.

Masimore, B., Schmitzer-Torbert, N. C., Kakalios, J., and David Redish, A. (2005). Transient striatal γ local field potentials signal movement initiation in rats. *NeuroReport*, 16(18):2021.

Matsumoto, N., Minamimoto, T., Graybiel, A. M., and Kimura, M. (2001). Neurons in the Thalamic CM-Pf Complex Supply Striatal Neurons With Information About Behaviorally Significant Sensory Events. *Journal of Neurophysiology*, 85(2):960–976.

McFarland, D. (1989). *Problems of Animal Behaviour*. Longman Scientific & Technical; Wiley, Harlow, Essex, England; New York. OCLC: 18106007.

McFarland, N. R. and Haber, S. N. (2000). Convergent inputs from thalamic motor nuclei and frontal cortical areas to the dorsal striatum in the primate. *Journal of Neuroscience*, 20(10):3798–3813.

McGeorge, A. J. and Faull, R. L. M. (1989). The organization of the projection from the cerebral cortex to the striatum in the rat. *Neuroscience*, 29(3):503–537.

McHaffie, J. G., Stanford, T. R., Stein, B. E., Coizet, V., and Redgrave, P. (2005). Subcortical loops through the basal ganglia. *Trends in Neurosciences*, 28(8):401–407.

Medina, L. and Reiner, A. (1995). Neurotransmitter Organization and Connectivity of the Basal Ganglia in Vertebrates: Implications for the Evolution of Basal Ganglia (Part 1 of 2). *Brain, Behavior and Evolution*, 46(4-5):235–246.

Mehler-Wex, C., Riederer, P., and Gerlach, M. (2006). Dopaminergic dysbalance in distinct basal ganglia neurocircuits: Implications for the pathophysiology of parkinson's disease, schizophrenia and attention deficit hyperactivity disorder. *Neurotoxicity Research*, 10(3-4):167–179.

Mengual, E., de las Heras, S., Erro, E., Lanciego, J. L., and Giménez-Amaya, J. M. (1999). Thalamic interaction between the input and the output systems of the basal ganglia. *Journal of Chemical Neuroanatomy*, 16(3):187–200.

Middleton, F. A. and Strick, P. L. (2000). Basal Ganglia Output and Cognition: Evidence from Anatomical, Behavioral, and Clinical Studies. *Brain and Cognition*, 42(2):183–200.

Middleton, F. A. and Strick, P. L. (2002). Basal-ganglia 'Projections' to the Prefrontal Cortex of the Primate. *Cerebral Cortex*, 12(9):926–935.

Minamimoto, T. and Kimura, M. (2002). Participation of the Thalamic CM-Pf Complex in Attentional Orienting. *Journal of Neurophysiology*, 87(6):3090–3101.

Mink, J. W. and Thach, W. T. (1993). Basal ganglia intrinsic circuits and their role in behavior. *Current Opinion in Neurobiology*, 3(6):950–957.

Misgeld, U., Okada, Y., and Hassler, R. (1979). Locally evoked potentials in slices of rat neostriatum: A tool for the investigation of intrinsic excitatory processes. *Experimental Brain Research*, 34(3):575–590.

Mitchell, I. J., Cooper, A. J., and Griffiths, M. R. (1999). The selective vulnerability of striatopallidal neurons. *Progress in Neurobiology*, 59(6):691–719.

Mitchinson, B., Chan, T.-S., Chambers, J., Pearson, M., Humphries, M., Fox, C., Gurney, K., and Prescott, T. J. (2010). BRAHMS: Novel middleware for integrated systems computation. *Advanced Engineering Informatics*, 24(1):49–61.

Morton, A. J., Hunt, M. J., Hodges, A. K., Lewis, P. D., Redfern, A. J., Dunnett, S. B., and Jones, L. (2005). A combination drug therapy improves cognition and reverses gene expression changes in a mouse model of Huntington's disease. *European Journal of Neuroscience*, 21(4):855–870.

Moyer, J. T., Wolf, J. A., and Finkel, L. H. (2007). Effects of Dopaminergic Modulation on the Integrative Properties of the Ventral Striatal Medium Spiny Neuron. *Journal of Neurophysiology*, 98(6):3731–3748.

Muñoz, M. and Coveñas, R. (2014). Involvement of substance P and the NK-1 receptor in human pathology. *Amino Acids*, 46(7):1727–1750.

Murtra, P., Sheasby, A. M., Hunt, S. P., and De Felipe, C. (2000). Rewarding effects of opiates are absent in mice lacking the receptor for substance P. *Nature*, 405(6783):180–183.

Naarding, P., Janzing, J. G., Eling, P., van der Werf, S., and Kremer, B. (2009). Apathy is not depression in Huntington's disease. *The Journal of neuropsychiatry and clinical neurosciences*, 21(3):266–270.

Nakahara, H., Itoh, H., Kawagoe, R., Takikawa, Y., and Hikosaka, O. (2004). Dopamine Neurons Can Represent Context-Dependent Prediction Error. *Neuron*, 41(2):269–280.

Nakano, K., Kayahara, T., Tsutsumi, T., and Ushiro, H. (2000). Neural circuits and functional organization of the striatum. *Journal of Neurology*, 247(5):V1–V15.

Nambu, A., Tokuno, H., and Takada, M. (2002). Functional significance of the cortico–subthalamo–pallidal 'hyperdirect' pathway. *Neuroscience Research*, 43(2):111–117.

Napier, T. C., Mitrovic, I., Churchill, L., Klitenick, M. A., Lu, X. Y., and Kalivas, P. W. (1995). Substance P in the ventral pallidum: Projection from the ventral striatum, and electrophysiological and behavioral cinsequences of pallidal substance P. *Neuroscience*, 69(1):59–70.

Nicola, S. M., Hopf, F. W., and Hjelmstad, G. O. (2004). Contrast enhancement: a physiological effect of striatal dopamine? *Cell and Tissue Research*, 318(1):93–106.

Nicola, S. M., Surmeier, D. J., and Malenka, R. C. (2000). Dopaminergic Modulation of Neuronal Excitability in the Striatum and Nucleus Accumbens. *Annual Review of Neuroscience*, 23(1):185–215.

Nisenbaum, E. S., Xu, Z. C., and Wilson, C. J. (1994). Contribution of a slowly inactivating potassium current to the transition to firing of neostriatal spiny projection neurons. *Journal of Neurophysiology*, 71(3):1174–1189.

Nordlie, E., Gewaltig, M.-O., and Plesser, H. E. (2009). Towards Reproducible Descriptions of Neuronal Network Models. *PLOS Comput Biol*, 5(8):e1000456.

O'Hare, J. K., Ade, K. K., Sukharnikova, T., Van Hooser, S. D., Palmeri, M. L., Yin, H. H., and Calakos, N. (2016). Pathway-Specific Striatal Substrates for Habitual Behavior. *Neuron*, 89(3):472–479.

O'Hare, J. K., Li, H., Kim, N., Gaidis, E., Ade, K., Beck, J., Yin, H., and Calakos, N. (2017). Striatal fast-spiking interneurons selectively modulate circuit output and are required for habitual behavior. *eLife*, 6.

Oorschot, D. E. (1996). Total number of neurons in the neostriatal, pallidal, subthalamic, and substantia nigral nuclei of the rat basal ganglia: A stereological study using the cavalieri and optical disector methods. *The Journal of Comparative Neurology*, 366(4):580–599.

Pakhotin, P. and Bracci, E. (2007). Cholinergic Interneurons Control the Excitatory Input to the Striatum. *The Journal of Neuroscience*, 27(2):391–400.

Parent, A., Charara, A., and Pinault, D. (1995). Single striatofugal axons arborizing in both pallidal segments and in the substantia nigra in primates. *Brain Research*, 698(1):280–284.

Parent, A. and Hazrati, L.-N. (1995a). Functional anatomy of the basal ganglia. I. The cortico-basal ganglia-thalamo-cortical loop. *Brain Research Reviews*, 20(1):91–127.

Parent, A. and Hazrati, L.-N. (1995b). Functional anatomy of the basal ganglia. II. The place of subthalamic nucleus and external pallidium in basal ganglia circuitry. *Brain Research Reviews*, 20(1):128–154.

Parthasarathy, H. B. and Graybiel, A. M. (1997). Cortically Driven Immediate-Early Gene Expression Reflects Modular Influence of Sensorimotor Cortex on Identified Striatal Neurons in the Squirrel Monkey. *Journal of Neuroscience*, 17(7):2477–2491.

Paulsen, J. S., Zhao, H., Stout, J. C., Brinkman, R. R., Guttman, M., Ross, C. A., Como, P., Manning, C., Hayden, M. R., Shoulson, I., and others (2001). Clinical markers of early disease in persons near onset of Huntington's disease. *Neurology*, 57(4):658–662.

Pennartz, C. M. A., Boeijinga, P. H., and Lopes da Silva, F. H. (1990). Locally evoked potentials in slices of the rat nucleus accumbens: NMDA and non-NMDA receptor mediated components and modulation by GABA. *Brain Research*, 529(1):30–41.

Planert, H., Szydlowski, S. N., Hjorth, J. J. J., Grillner, S., and Silberberg, G. (2010). Dynamics of Synaptic Transmission between Fast-Spiking Interneurons and Striatal Projection Neurons of the Direct and Indirect Pathways. *The Journal of Neuroscience*, 30(9):3499–3507.

Ponzi, A. and Wickens, J. (2010). Sequentially Switching Cell Assemblies in Random Inhibitory Networks of Spiking Neurons in the Striatum. *The Journal of Neuroscience*, 30(17):5894–5911.

Popescu, A. T., Popa, D., and Paré, D. (2009). Coherent gamma oscillations couple the amygdala and striatum during learning. *Nature Neuroscience*, 12(6):801.

Porter, A. J., Pillidge, K., Tsai, Y. C., Dudley, J. A., Hunt, S. P., Peirson, S. N., Brown, L. A., and Stanford, S. C. (2015). A lack of functional NK1 receptors explains most, but not all, abnormal behaviours of NK1r-/- mice. *Genes, Brain and Behavior*, 14(2):189–199.

Prescott, T. J. (2007). Forced Moves or Good Tricks in Design Space? Landmarks in the Evolution of Neural Mechanisms for Action Selection. *Adaptive Behavior*, 15(1):9–31.

Prescott, T. J., Montes González, F. M., Gurney, K., Humphries, M. D., and Redgrave, P. (2006). A robot model of the basal ganglia: Behavior and intrinsic processing. *Neural Networks*, 19(1):31–61.

Prescott, T. J., Redgrave, P., and Gurney, K. (1999). Layered Control Architectures in Robots and Vertebrates. *Adaptive Behavior*, 7(1):99–127.

Presti, M. F. and Lewis, M. H. (2005). Striatal opioid peptide content in an animal model of spontaneous stereotypic behavior. *Behavioural Brain Research*, 157(2):363–368.

Ramanathan, S., Hanley, J. J., Deniau, J.-M., and Bolam, J. P. (2002). Synaptic Convergence of Motor and Somatosensory Cortical Afferents onto GABAergic Interneurons in the Rat Striatum. *Journal of Neuroscience*, 22(18):8158–8169.

Ramkumar, P., Acuna, D. E., Berniker, M., Grafton, S. T., Turner, R. S., and Kording, K. P. (2016). Chunking as the result of an efficiency computation trade-off. *Nature Communications*, 7:12176.

Redgrave, P. (2007). Basal ganglia. *Scholarpedia*, 2(6):1825.

Redgrave, P. and Gurney, K. (2006). The short-latency dopamine signal: a role in discovering novel actions? *Nature Reviews Neuroscience*, 7(12):967–975.

Redgrave, P., Gurney, K., and Reynolds, J. (2008). What is reinforced by phasic dopamine signals? *Brain Research Reviews*, 58(2):322–339.

Redgrave, P., Prescott, T. J., and Gurney, K. (1999). The basal ganglia: a vertebrate solution to the selection problem? *Neuroscience*, 89(4):1009–1023.

Redgrave, P., Rodriguez, M., Smith, Y., Rodriguez-Oroz, M. C., Lehericy, S., Bergman, H., Agid, Y., DeLong, M. R., and Obeso, J. A. (2010). Goal-directed and habitual control in the basal ganglia: implications for Parkinson's disease. *Nature reviews. Neuroscience*, 11(11):760–772.

Redgrave, P., Vautrelle, N., and Reynolds, J. N. J. (2011). Functional properties of the basal ganglia's re-entrant loop architecture: selection and reinforcement. *Neuroscience*, 198:138–151.

Reiner, A. (2010). The Conservative Evolution of the Vertebrate Basal Ganglia. In Steiner, H. and Tseng, K. Y., editors, *Handbook of Basal Ganglia Structure and Function*, pages 29–62. Academic Press, Burlington, MA.

Reiner, A. and Anderson, K. D. (1990). The patterns of neurotransmitter and neuropeptide co-occurrence among striatal projection neurons: conclusions based on recent findings. *Brain Research Reviews*, 15(3):251–265.

Reiner, A., Medina, L., and Veenman, C. (1998). Structural and functional evolution of the basal ganglia in vertebrates. *Brain Research Reviews*, 28(3):235–285.

Richfield, E. K., Maguire-Zeiss, K. A., Vonkeman, H. E., and Voorn, P. (1995). Preferential loss of preproenkephalin versus preprotachykinin neurons from the striatum of Huntington's disease patients. *Annals of Neurology*, 38(6):852–861.

Richmond, P., Cope, A., Gurney, K., and Allerton, D. J. (2014). From Model Specification to Simulation of Biologically Constrained Networks of Spiking Neurons. *Neuroinformatics*, 12(2):307–323.

Romanelli, P., Esposito, V., Schaal, D. W., and Heit, G. (2005). Somatotopy in the basal ganglia: experimental and clinical evidence for segregated sensorimotor channels. *Brain Research Reviews*, 48(1):112–128.

Romanelli, P., Heit, G., Hill, B. C., Kraus, A., Hastie, T., and Brontë-Stewart, H. M. (2004). Microelectrode recording revealing a somatotopic body map in the subthalamic nucleus in humans with Parkinson disease. *Journal of Neurosurgery*, 100:611–618.

Rueda-Orozco, P. E. and Robbe, D. (2015). The striatum multiplexes contextual and kinematic information to constrain motor habits execution. *Nature Neuroscience*, 18(3):453–460.

Rymar, V. V., Sasseville, R., Luk, K. C., and Sadikot, A. F. (2004). Neurogenesis and stereological morphometry of calretinin-immunoreactive GABAergic interneurons of the neostriatum. *The Journal of Comparative Neurology*, 469(3):325–339.

Sano, H., Chiken, S., Hikida, T., Kobayashi, K., and Nambu, A. (2013). Signals through the Striatopallidal Indirect Pathway Stop Movements by Phasic Excitation in the Substantia Nigra. *Journal of Neuroscience*, 33(17):7583–7594.

Sano, H., Yasoshima, Y., Matsushita, N., Kaneko, T., Kohno, K., Pastan, I., and Kobayashi, K. (2003). Conditional Ablation of Striatal Neuronal Types Containing Dopamine D2 Receptor Disturbs Coordination of Basal Ganglia Function. *Journal of Neuroscience*, 23(27):9078–9088.

Savalia, T., Shukla, A., and Bapi, R. S. (2016). A Unified Theoretical Framework for Cognitive Sequencing. *Frontiers in Psychology*, 7.

Schmidt, R., Leventhal, D. K., Mallet, N., Chen, F., and Berke, J. D. (2013). Canceling actions involves a race between basal ganglia pathways. *Nature Neuroscience*, 16(8):1118–1124.

Schroll, H. and Hamker, F. H. (2013). Computational models of basal-ganglia pathway functions: focus on functional neuroanatomy. *Frontiers in Systems Neuroscience*, 7.

Schultz, W. (2007). Multiple Dopamine Functions at Different Time Courses. *Annual Review of Neuroscience*, 30(1):259–288.

Schultz, W. and Dickinson, A. (2000). Neuronal coding of prediction errors. *Annual review of neuroscience*, 23(1):473–500.

Searle, J. R. (1979). The intentionality of intention and action. *Inquiry*, 22(1-4):253–280.

Seeman, P. and Van Tol, H. H. M. (1994). Dopamine receptor pharmacology. *Trends in Pharmacological Sciences*, 15(7):264–270.

Sharott, A., Moll, C. K. E., Engler, G., Denker, M., Grün, S., and Engel, A. K. (2009). Different Subtypes of Striatal Neurons Are Selectively Modulated by Cortical Oscillations. *Journal of Neuroscience*, 29(14):4571–4585.

Shipp, S. (2017). The functional logic of corticostriatal connections. *Brain Structure and Function*, 222(2):669–706.

Sidibé, M. and Smith, Y. (1999). Thalamic inputs to striatal interneurons in monkeys: synaptic organization and co-localization of calcium binding proteins. *Neuroscience*, 89(4):1189–1208.

Sippy, T., Lapray, D., Crochet, S., and Petersen, C. C. H. (2015). Cell-Type-Specific Sensorimotor Processing in Striatal Projection Neurons during Goal-Directed Behavior. *Neuron*, 88(2):298–305.

Smith, J. B., Klug, J. R., Ross, D. L., Howard, C. D., Hollon, N. G., Ko, V. I., Hoffman, H., Callaway, E. M., Gerfen, C. R., and Jin, X. (2016). Genetic-Based Dissection Unveils the Inputs and Outputs of Striatal Patch and Matrix Compartments. *Neuron*, 91(5):1069–1084.

Smith, Y., Beyan, M. D., Shink, E., and Bolam, J. P. (1998). Microcircuitry of the direct and indirect pathways of the basal ganglia. *Neuroscience*, 86:353–388.

Smith-Roe, S. L. and Kelley, A. E. (2000). Coincident Activation of NMDA and Dopamine D1 Receptors within the Nucleus Accumbens Core Is Required for Appetitive Instrumental Learning. *Journal of Neuroscience*, 20(20):7737–7742.

Solbrig, M. V., Koob, G. F., and Lipkin, W. I. (2002). Key role for enkephalinergic tone in cortico–striatal–thalamic function. *European Journal of Neuroscience*, 16(9):1819–1822.

Solopchuk, O., Alamia, A., Olivier, E., and Zénon, A. (2016). Chunking improves symbolic sequence processing and relies on working memory gating mechanisms. *Learning & Memory*, 23(3):108–112.

Steiner, H. and Gerfen, C. R. (1998). Role of dynorphin and enkephalin in the regulation of striatal output pathways and behavior. *Experimental Brain Research*, 123(1-2):60–76.

Steno, N. (1669). *Discours de Monsieur Stenon sur l'anatomie du cerveau.* chez Robert de Ninuille, Paris.

Stephenson-Jones, M., Samuelsson, E., Ericsson, J., Robertson, B., and Grillner, S. (2011). Evolutionary Conservation of the Basal Ganglia as a Common Vertebrate Mechanism for Action Selection. *Current Biology*, 21(13):1081–1091.

Stephenson-Jones, M., Yu, K., Ahrens, S., Tucciarone, J. M., van Huijstee, A. N., Mejia, L. A., Penzo, M. A., Tai, L.-H., Wilbrecht, L., and Li, B. (2016). A basal ganglia circuit for evaluating action outcomes. *Nature*, 539(7628):289–293.

Stocco, A. (2012). Acetylcholine-Based Entropy in Response Selection: A Model of How Striatal Interneurons Modulate Exploration, Exploitation, and Response Variability in Decision-Making. *Frontiers in Neuroscience*, 6.

Stocco, A. and Lebiere, C. (2014). Inhibitory synapses between striatal projection neurons support efficient enhancement of cortical signals: A computational model. *Journal of Computational Neuroscience*, 37(1):65.

Surmeier, D. J., Song, W.-J., and Yan, Z. (1996). Coordinated Expression of Dopamine Receptors in Neostriatal Medium Spiny Neurons. *Journal of Neuroscience*, 16(20):6579–6591.

Takada, M., Tokuno, H., Nambu, A., and Inase, M. (1998). Corticostriatal projections from the somatic motor areas of the frontal cortex in the macaque monkey: segregation versus overlap of input zones from the primary motor cortex, the supplementary motor area, and the premotor cortex. *Experimental Brain Research*, 120(1):114–128.

Tan, C. O. and Bullock, D. (2008). A Dopamine–Acetylcholine Cascade: Simulating Learned and Lesion-Induced Behavior of Striatal Cholinergic Interneurons. *Journal of Neurophysiology*, 100(4):2409–2421.

Tanaka, S. C., Doya, K., Okada, G., Ueda, K., Okamoto, Y., and Yamawaki, S. (2004). Prediction of immediate and future rewards differentially recruits cortico-basal ganglia loops. *Nature Neuroscience*, 7(8):887–893.

Tateno, T., Harsch, A., and Robinson, H. P. C. (2004). Threshold Firing Frequency–Current Relationships of Neurons in Rat Somatosensory Cortex: Type 1 and Type 2 Dynamics. *Journal of Neurophysiology*, 92(4):2283–2294.

Taverna, S., Ilijic, E., and Surmeier, D. J. (2008). Recurrent Collateral Connections of Striatal Medium Spiny Neurons Are Disrupted in Models of Parkinson's Disease. *Journal of Neuroscience*, 28(21):5504–5512.

Tepper, J. M. and Bolam, J. P. (2004). Functional diversity and specificity of neostriatal interneurons. *Current Opinion in Neurobiology*, 14(6):685–692.

Tepper, J. M., Wilson, C. J., and Koós, T. (2008). Feedforward and feedback inhibition in neostriatal GABAergic spiny neurons. *Brain Research Reviews*, 58(2):272–281.

Tewari, A., Jog, R., and Jog, M. S. (2016). The Striatum and Subthalamic Nucleus as Independent and Collaborative Structures in Motor Control. *Frontiers in Systems Neuroscience*, 10.

Thomas, T. M., Smith, Y., Levey, A. I., and Hersch, S. M. (2000). Cortical inputs to m2-immunoreactive striatal interneurons in rat and monkey. *Synapse*, 37(4):252–261.

Thompson, J. C., Snowden, J. S., Craufurd, D., and Neary, D. (2002). Behavior in Huntington's disease: dissociating cognition-based and mood-based changes. *The Journal of neuropsychiatry and clinical neurosciences*, 14(1):37–43.

Thornton, E. and Vink, R. (2012). Treatment with a Substance P Receptor Antagonist Is Neuroprotective in the Intrastriatal 6-Hydroxydopamine Model of Early Parkinson's Disease. *PLoS ONE*, 7(4):e34138.

Threlfell, S. and Cragg, S. J. (2011). Dopamine Signaling in Dorsal Versus Ventral Striatum: The Dynamic Role of Cholinergic Interneurons. *Frontiers in Systems Neuroscience*, 5.

Tomkins, A., Humphries, M., Beste, C., Vasilaki, E., and Gurney, K. (2012). How Degrading Networks Can Increase Cognitive Functions. In Villa, A. E. P., Duch, W., Érdi, P., Masulli, F., and Palm, G., editors, *Artificial Neural Networks and Machine Learning – ICANN 2012*, number 7552 in Lecture Notes in Computer Science, pages 185–192. Springer Berlin Heidelberg.

Tomkins, A., Vasilaki, E., Beste, C., Gurney, K., and Humphries, M. D. (2014). Transient and steady-state selection in the striatal microcircuit. *Frontiers in Computational Neuroscience*, 7.

Tort, A. B. L., Kramer, M. A., Thorn, C., Gibson, D. J., Kubota, Y., Graybiel, A. M., and Kopell, N. J. (2008). Dynamic cross-frequency couplings of local field potential oscillations in rat striatum and hippocampus during performance of a T-maze task. *Proceedings of the National Academy of Sciences*, 105(51):20517–20522.

Traub, R. D., Kopell, N., Bibbig, A., Buhl, E. H., LeBeau, F. E. N., and Whittington, M. A. (2001). Gap Junctions between Interneuron Dendrites Can Enhance Synchrony of Gamma Oscillations in Distributed Networks. *Journal of Neuroscience*, 21(23):9478–9486.

Tremblay, P.-L., Bedard, M.-A., Langlois, D., Blanchet, P. J., Lemay, M., and Parent, M. (2010). Movement chunking during sequence learning is a dopamine-dependant process: a study conducted in Parkinson's disease. *Experimental Brain Research*, 205(3):375–385.

Tunstall, M. J., Oorschot, D. E., Kean, A., and Wickens, J. R. (2002). Inhibitory interactions between spiny projection neurons in the rat striatum. *Journal of Neurophysiology*, 88(3):1263–1269.

Turjanski, N., Weeks, R., Dolan, R., Harding, A. E., and Brooks, D. J. (1995). Striatal D1 and D2 receptor binding in patients with Huntington's disease and other choreas A PET study. *Brain*, 118(3):689–696.

van der Meer, M. A. A., Kalenscher, T., Lansink, C. S., Pennartz, C., Berke, J. D., and Redish, D. A. (2010). Integrating Early Results on Ventral Striatal Gamma Oscillations in the Rat. *Frontiers in Neuroscience*, 4.

van der Meer, M. A. A. and Redish, A. D. (2009). Low and High Gamma Oscillations in Rat Ventral Striatum have Distinct Relationships to Behavior, Reward, and Spiking Activity on a Learned Spatial Decision Task. *Frontiers in Integrative Neuroscience*, 3.

Van Wimersma Greidanus, T. B. and Maigret, C. (1988). Grooming behavior induced by substance P. *European Journal of Pharmacology*, 154(2):217–220.

Verstynen, T. D., Badre, D., Jarbo, K., and Schneider, W. (2012). Microstructural organizational patterns in the human corticostriatal system. *Journal of Neurophysiology*, 107(11):2984–2995.

Verwey, W. B. (2001). Concatenating familiar movement sequences: the versatile cognitive processor. *Acta Psychologica*, 106(1):69–95.

Verwey, W. B., Abrahamse, E. L., and de Kleine, E. (2010). Cognitive Processing in New and Practiced Discrete Keying Sequences. *Frontiers in Psychology*, 1.

von Euler, U. S. and Pernow, B. (1954). Effects of Intraventricular Administration of Substance P. *Nature*, 174(4421):184.

von Euler, U. S. and Pernow, B. (1956). Neurotropic Effects of Substance P. *Acta Physiologica Scandinavica*, 36(3):265–275.

Voorn, P., Vanderschuren, L. J. M. J., Groenewegen, H. J., Robbins, T. W., and Pennartz, C. M. A. (2004). Putting a spin on the dorsal–ventral divide of the striatum. *Trends in Neurosciences*, 27(8):468–474.

Wall, N. R., De La Parra, M., Callaway, E. M., and Kreitzer, A. C. (2013). Differential Innervation of Direct- and Indirect-Pathway Striatal Projection Neurons. *Neuron*, 79(2):347–360.

Wang, H.-B., Deng, Y.-P., and Reiner, A. (2007). In situ hybridization histochemical and immunohistochemical evidence that striatal projection neurons co-containing substance P and enkephalin are overrepresented in the striosomal compartment of striatum in rats. *Neuroscience Letters*, 425(3):195–199.

Weeks, R. A., Piccini, P., Harding, A. E., and Brooks, D. J. (1996). Striatal D1 and D2 dopamine receptor loss in asymptomatic mutation carriers of Huntington's disease. *Annals of Neurology*, 40(1):49–54.

West, M. O., Carelli, R. M., Pomerantz, M., Cohen, S. M., Gardner, J. P., Chapin, J. K., and Woodward, D. J. (1990). A region in the dorsolateral striatum of the rat exhibiting single-unit correlations with specific locomotor limb movements. *Journal of Neurophysiology*, 64(4):1233–1246.

Wickens, J. R., Kotter, R., and Alexander, M. E. (1995). Effects of local connectivity on striatal function: Simulation and analysis of a model. *Synapse*, 20(4):281–298.

Willingham, D. B. and Koroshetz, W. J. (1993). Evidence for dissociable motor skills in Huntington's disease patients. *Psychobiology*, 21(3):173–182.

Wilson, C. J. and Groves, P. M. (1980). Fine structure and synaptic connections of the common spiny neuron of the rat neostriatum: A study employing intracellular injection of horseradish peroxidase. *The Journal of Comparative Neurology*, 194(3):599–615.

Wilson, C. J. and Kawaguchi, Y. (1996). The origins of two-state spontaneous membrane potential fluctuations of neostriatal spiny neurons. *Journal of Neuroscience*, 16(7):2397–2410.

Wu, Y. and Parent, A. (2000). Striatal interneurons expressing calretinin, parvalbumin or NADPH-diaphorase: a comparative study in the rat, monkey and human. *Brain Research*, 863(1):182–191.

Wymbs, N. F., Bassett, D. S., Mucha, P. J., Porter, M. A., and Grafton, S. T. (2012). Differential Recruitment of the Sensorimotor Putamen and Frontoparietal Cortex during Motor Chunking in Humans. *Neuron*, 74(5):936–946.

Yamada, H., Matsumoto, N., and Kimura, M. (2004). Tonically Active Neurons in the Primate Caudate Nucleus and Putamen Differentially Encode Instructed Motivational Outcomes of Action. *Journal of Neuroscience*, 24(14):3500–3510.

Yan, T. C., Hunt, S. P., and Stanford, S. C. (2009). Behavioural and neurochemical abnormalities in mice lacking functional tachykinin-1 (NK1) receptors: A model of attention deficit hyperactivity disorder. *Neuropharmacology*, 57(7–8):627–635.

Yan, T. C., McQuillin, A., Thapar, A., Asherson, P., Hunt, S. P., Stanford, S. C., and Gurling, H. (2010). NK1 (TACR1) receptor gene 'knockout' mouse phenotype predicts genetic association with ADHD. *Journal of Psychopharmacology*, 24(1):27–38.

Yelnik, J. (2002). Functional anatomy of the basal ganglia. *Movement Disorders*, 17(S3):S15–S21.

Yeterian, E. H. and Van Hoesen, G. W. (1978). Cortico-striate projections in the rhesus monkey: The organization of certain cortico-caudate connections. *Brain Research*, 139(1):43–63.

Yin, H. H., Mulcare, S. P., Hilário, M. R. F., Clouse, E., Holloway, T., Davis, M. I., Hansson, A. C., Lovinger, D. M., and Costa, R. M. (2009). Dynamic reorganization of striatal circuits during the acquisition and consolidation of a skill. *Nature Neuroscience*, 12(3):333–341.

Yung, K. K. L., Smith, A. D., Levey, A. I., and Bolam, J. P. (1996). Synaptic Connections Between Spiny Neurons of the Direct and Indirect Pathways in the Neostriatum of the Rat: Evidence from Dopamine Receptor and Neuropeptide Immunostaining. *European Journal of Neuroscience*, 8(5):861–869.

Záborszky, L., Alheid, G. F., Beinfeld, M. C., Eiden, L. E., Heimer, L., and Palkovits, M. (1985). Cholecystokinin innervation of the ventral striatum: A morphological and radioimmunological study. *Neuroscience*, 14(2):427–453.

Zhuang, X., Oosting, R. S., Jones, S. R., Gainetdinov, R. R., Miller, G. W., Caron, M. G., and Hen, R. (2001). Hyperactivity and impaired response habituation in hyperdopaminergic mice. *Proceedings of the National Academy of Sciences*, 98(4):1982–1987.