Dopamine-modulated dynamic cell assemblies generated by the GABAergic striatal microcircuit

Mark D. Humphries*, Ric Wood, Kevin Gurney

Adaptive Behaviour Research Group, Department of Psychology, University of Sheffield, S10 2TN, UK

Abstract

The striatum, principal input structure of the basal ganglia, is crucial to both motor control and learning. It receives convergent input from all over neocortex, hippocampal formation, amygdala and thalamus, and is the primary recipient of dopamine in the brain. Within the striatum is a GABAergic microcircuit that acts upon these inputs, formed by the dominant medium-spiny projection neurons (MSNs) and fast-spiking interneurons (FSIs). There has been little progress in understanding the computations it performs, hampered by the non-laminar structure that prevents identification of a repeating canonical microcircuit. We here begin the identification of potential dynamically-defined computational elements within the striatum. We construct a new three-dimensional model of the striatal microcircuit's connectivity, and instantiate this with our dopamine-modulated neuron models of the MSNs and FSIs. A new model of gap junctions between the FSIs is introduced and tuned to experimental data. We introduce a novel multiple spike-train analysis method, and apply this to the outputs of the model to find groups of synchronised neurons at multiple time-scales. We find that, with realistic *in vivo* background input, small assemblies of synchronised MSNs spontaneously appear, consistent with experimental observations, and that the number of assemblies and the time-scale of synchronisation is strongly dependent on the simulated concentration of dopamine. We also show that feed-forward inhibition from the FSIs counter-intuitively increases the firing rate of the MSNs. Such small cell assemblies forming spontaneously only in the absence of dopamine may contribute to motor control problems seen in humans and animals following loss of dopamine cells.

Key words: striatum, Parkinson's disease, action selection, large-scale models, Izhikevich neuron model

1. Introduction

The striatum is a large subcortical nucleus that forms the principal input structure of the basal ganglia. Diseases that directly affect the striatum or its primary afferents – such as Huntington's or Parkinson's disease – lead to profound deficits in motor control. In particular, loss of dopamine cells in Parkinson's disease and its animal models leads to motor symptoms of rigidity, akinesia, and tremor (Schwarting & Huston, 1996; Kirik et al., 1998; Ferro et al., 2005), and the striatum is the main locus of dopamine's action, containing the highest density of dopamine receptors in the vertebrate brain (Dawson et al., 1986; Richtand et al., 1995; Hurd et al., 2001). Moreover, an intact dopamine system also seems to be critical for many forms of learning (Whishaw & Dunnett, 1985; Ferro et al., 2005), consistent with reported correlations between dopamine cell firing and the prediction error required by reinforcement learning theories (Redgrave & Gurney, 2006; Schultz, 2007). An intact striatum is similarly required for successful acquisition of many instrumental conditioning tasks (Yin & Knowlton, 2006). An understanding of the striatum's computational operation would thus shed light on a fundamental contributor to both motor control and learning.

Within the striatum lies a complex, predominantly GABAergic, microcircuit (Bolam et al., 2006). Medium spiny projection neurons (MSNs) are the only output neurons and comprise up to 95% of the cell population in rat, with GABAergic and cholinergic interneurons forming most of the remaining cell population. Despite their comparatively small number, the GABAergic fast-spiking interneurons (FSIs) in particular exert a very strong influence on the MSNs (Koos & Tepper, 1999), receive input from similar sources, and are interconnected by both chemical synapses and gap junctions. Dopamine has multiple effects on these neuron types, via multiple receptor types: indeed, the exact effects of dopamine receptor activation on the MSN have been much debated (Surmeier et al., 2007). Seemingly ideal for underpinning its multiple functional roles, the striatum receives massive convergent input from the neocortex, thalamus, hippocampal formation, and amygdala (McGeorge & Faull, 1989; Groenewegen et al., 1999; Glynn & Ahmad, 2002; Smith et al., 2004), and dopamine modulates striatal neurons' responses to them.

Despite, or perhaps due to, this complexity of structure and input, there are few well-quantified theories of the striatum's computational role. Many theories of striatal-specific or global basal ganglia function draw explicit attention to the role of the inhibitory local MSN collaterals as a substrate for competitive dynamics (e.g. Wickens et al., 1991; Pennartz et al., 1994; Beiser & Houk, 1998; Frank, 2005), whether that compe-

*Corresponding author

Email address: m.d.humphries@sheffield.ac.uk (Mark D. Humphries)

tition be labelled 'decision making', 'motor program selection' or 'pattern classification'. Wickens and colleagues' domain hypothesis is the most developed, and proposes that the basic computational element of the striatum is the set – or "domain" – of all MSNs that are mutual inhibitory (see e.g. Wickens et al., 1991; Alexander & Wickens, 1993; Wickens et al., 1995). In simulation, they have shown that winner-takes-all like competition occurs within a single domain, while winners-share-all dynamics (multiple active neurons) occur in networks composed of multiple overlapping domains (Wickens et al., 1991; Alexander & Wickens, 1993). Similar results have been obtained in analytical studies of general mutually inhibitory neural networks (Fukai & Tanaka, 1997).

All such theories of competitive dynamics are faced by the problems that the inhibition provided by the local MSN collaterals is weak (Jaeger et al., 1994; Tunstall et al., 2002; Czubayko & Plenz, 2002; Koos et al., 2004; Taverna et al., 2004), that a single MSN is only contacted by between 12-18% of MSNs in its dendritic field (Tepper et al., 2004), and that *mutual* inhibition is the exception rather than rule (Tunstall et al., 2002; Tepper et al., 2004).

Some theories do predict such weak connections. Bar-Gad et al. (2003) have proposed that the striatum compresses information relayed to it from cortex, transmitting back the compressed version via the basal ganglia output nuclei. They noted that the two layer network formed by the striatum and the output nuclei can be mapped to standard neural network implementations of principal components analysis, and that these require weak correlation in a layer corresponding to the striatum. While an interesting hypothesis, this mapping does not account for the microcircuit of the striatum, or the effects of the numerous neuromodulators within it. Other models of the whole basal ganglia circuit do not rely on the local collaterals within striatum for their computations, rather proposing that the striatum is both integrator of diverse cortical information and filter on weak cortical inputs, as the first stage of an input selection mechanism implemented by the whole basal ganglia (as opposed to just the striatum), (Gurney et al., 2001; Humphries et al., 2006) – but these models too do not account for the striatal microcircuit.

Our aim is to find out what computations can be supported by the intrinsic circuitry of the striatum, what – if any – "basic computational elements" exist, and develop computational theories of function on this basis. In particular, we wish to understand the role of the dominant GABAergic circuits of the striatum: the rare, but powerful, FSIs, and the weak, asymmetrical, but comparatively plentiful MSN local collaterals. Understanding the contribution of all the striatum's elements ideally requires large-scale models (Djurfeldt et al., 2008) that replicate the neuron types, numbers, and connectivity at one-to-one scale. Such models can give deep insight into the role of each neuron class in local circuit dynamics.

The purpose of this paper is twofold. First, we draw together for the first time a series of techniques we have developed for leveraging anatomical and physiological constraint data, some of which promise general applicability (beyond the striatum) in microcircuit construction: (1) a powerful computational neuroanatomy method for extracting the best connectivity statistics from impoverished data; (2) the development of reduced models for dopamine modulation of striatal neurons, which replicate the output of detailed compartmental models; and (3) a rigorous method for spike generation which allows good approximation to cortical input. We add to these here by introducing: (1) a gap junction model tunable to known membrane properties; (2) a principled method for parameterising the spike generation tool based on anatomical and physiological data; and (3) a novel method for detecting patterns in multi-unit activity at multiple time-scales, with general applicability to simulation or experimental data.

Second, we begin the identification of computational elements within the striatum, and examine how these might support hypotheses for competitive dynamics underpinned by the GABAergic neurons of the striatum. Specifically, we construct a three-dimensional model of the striatal microcircuit's connectivity, and instantiate this with our dopamine-modulated neuron models of the MSNs and FSIs. We apply our multiple spike-train analysis to the outputs of this model to find groups of synchronised neurons at multiple time-scales. We then show that, with realistic in vivo background input, small assemblies of synchronised MSNs spontaneously appear, consistent with experimental observations (Carrillo-Reid et al., 2008), and that the number of assemblies and the time-scale of synchronisation is strongly dependent on the simulated concentration of dopamine.

2. Creating the striatal microcircuit

Building large-scale models at up to 1:1 scale, neuron for neuron, is an ambitious aim. In particular, as recognised by the Blue Brain Project (Markram, 2006), these models are severely limited by the need for accurate connectivity. There is a wealth of studies showing how the structure of a network is a strong determinant of its dynamics (see e.g. Nishikawa et al., 2003; Kwok et al., 2007; Galan, 2008), and that the typical fall-back of completely regular or random networks give false impressions about both synchronisation and stability (see especially Watts & Strogatz, 1998; Lago-Fernandez et al., 2000). It is thus imperative that we begin from as accurate a network structure as possible.

2.1. The striatal microcircuit

First, we specify the GABAergic microcircuit of the striatum (Tepper et al., 2004). Figure 1 shows its complete set of connections and neuron types; these are intermingled throughout the non-laminar structure of the striatum. The MSNs number around 2,790,000 in the rat, with a (shrinkage-corrected) density of 85,000 per mm^3 (Oorschot, 1996). Various estimates have placed this total as anything up to 95% of all neurons in the striatum (Gerfen & Wilson, 1996), though a figure of 90% is more commonly quoted (Kawaguchi et al., 1995). The MSNs can be split into two populations on the basis of their dominant expression of either the D1 or D2 dopamine receptor, and these populations are of roughly equal size. In addition to their long axonal projections to targets in the pallidum (D2 MSNs) and

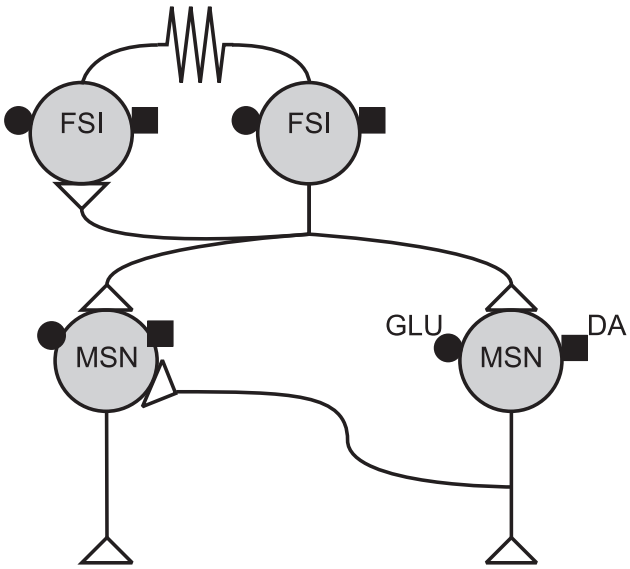

Figure 1: The striatal GABAergic microcircuit studied in this paper. Primary input to the striatum comes from glutamatergic (GLU: •) fibres originating in the neocortex, thalamus, hippocampal formation and amygdala, and dopaminergic (DA: ■) fibres originating in the hindbrain dopamine cell bands. All striatal axo-dendritic connections (△) are GABAergic and hence inhibitory. The fast-spiking interneurons (FSIs) can form dendro-dendritic gap junctions between them. The medium spiny neuron (MSN) population can be divided in two on the basis of the dominant dopamine receptor (D1 or D2) they express.

substantia nigra pars reticulata (D1 MSNs) both types have extensive local axon collaterals, which ramify in approximately the same volume as the parent neurons' dendrites.

The physiological class of FSIs seem to correspond to the class of parvalbumin-immunoreactive interneurons (Kawaguchi, 1993), and these comprise around 3-5% of the striatal cell population in the rat (Kawaguchi et al., 1995). Their axons remain wholly in the striatum, and target both MSNs (Koos & Tepper, 1999) and other FSIs (Kita et al., 1990). In addition, there are dendro-dendritic gap junctions between FSIs (Koos & Tepper, 1999). Both MSNs and FSIs receive glutamatergic input from cortical and thalamic sources, and dopaminergic input from the hindbrain dopamine cell bands.

We focus on this microcircuit as the neuron types are the best characterised (Tepper et al., 2004), but hence omit at least two other physiological classes of interneuron found in the striatum. The long-lasting hyperpolarisation class correspond to the large aspiny cholinergic interneurons (Kawaguchi, 1993). We are focusing here on the short time-scale dynamics in striatum, which are thought to be dominated by the GABAergic elements (Tepper et al., 2004; Mallet et al., 2005). Future work on this circuit will incorporate the cholinergic interneurons, as they may play role in setting the dynamic state of the striatal network (Wickens et al., 1991) and their regulation of dopamine release affects plasticity at cortico-striatal synapses (Zhou et al., 2002; Wang et al., 2006). The low-threshold spiking class correspond to the interneurons that co-express nitric oxide, somatostatin, and neuropeptide Y (Kawaguchi, 1993; Kawaguchi et al., 1995); this class may also express GABA (Kubota & Kawaguchi, 2000). The FSIs probably dominate MSN behaviour, as they form far more synapses on somas (Kubota & Kawaguchi, 2000), whereas the low-threshold spiking neurons may form an inhibitory network between the cholinergic interneurons (Sullivan et al., 2008).

2.2. The neuroanatomical model

We developed a novel computational neuroanatomy method to build a three-dimensional model of the striatum that is as accurate as possible given current neuroanatomical data (Wood et al., 2007). The strength of this method is that it can be updated and re-run each time new relevant data becomes available. We review the outline of the method and the results essential for reconstructing the network used here.

Our approach builds on the underlying assumption that the probability of connection between a given pair of neurons is proportional to the distance between the cell bodies, and the overlap of their neurites at that distance. For a standard axo-dendritic synapse, the probability of connection is thus proportional to the joint volume occupied by both the axonal field of the source neuron and the dendritic field of the target neuron. However, like much neural tissue, detailed data on the dendrites, axons, and their three-dimensional structure were not available for the MSNs and FSIs.

We thus developed the method outlined in Figure 2. This method relies on developing stochastic growth models for the dendrites and axons of both MSNs and FSIs. For the dendritic trees, we used an existing growth algorithm (Burke et al., 1992) and found its parameters using a genetic algorithm search of a fitness space defined by known parameters (e.g. number of branch points) of the neuron type's dendritic tree. For the axon, which has a simpler structure, we created our own growth algorithm based on known properties of MSN and FSI axons. By creating models for the dendrite and axon structure, we had a full set of data on the dendritic branches and axons at each distance from the soma, including their approximate volume. Using the growth algorithms, we produced a large number of dendritic trees and axons to estimate expected neurite volume.

Based on this, we could then compute the expected volume of a sphere that was occupied by dendrite (or axon) at a given distance from the cell body. Both MSNs (Wilson & Groves, 1980; Zheng & Wilson, 2002) and FSIs (Kawaguchi, 1993; Koos & Tepper, 1999) have approximately spherical dendritic and axonal fields, and so we could compute the expected amount of neurite in all directions – effectively modelling a mean-field dendrite or axon. Then, in turn, we could compute the expected volume of overlap between the spherical fields given the distance between cell bodies for each connection type. For every $1\mu\text{m}^3$ voxel in this overlapping volume, we computed the probability of its occupancy by both neurites (axon and dendrite or dendrite and dendrite, depending on the connection type) and thus the probability of a potential contact. Summing over all voxels in the overlapping volume thus gave us the expected number of contacts for each distance between cell bodies.

2.3. Construction of the network

We found that the expected number of contacts between two neurons, as a function of the distance d_s between the two somas,

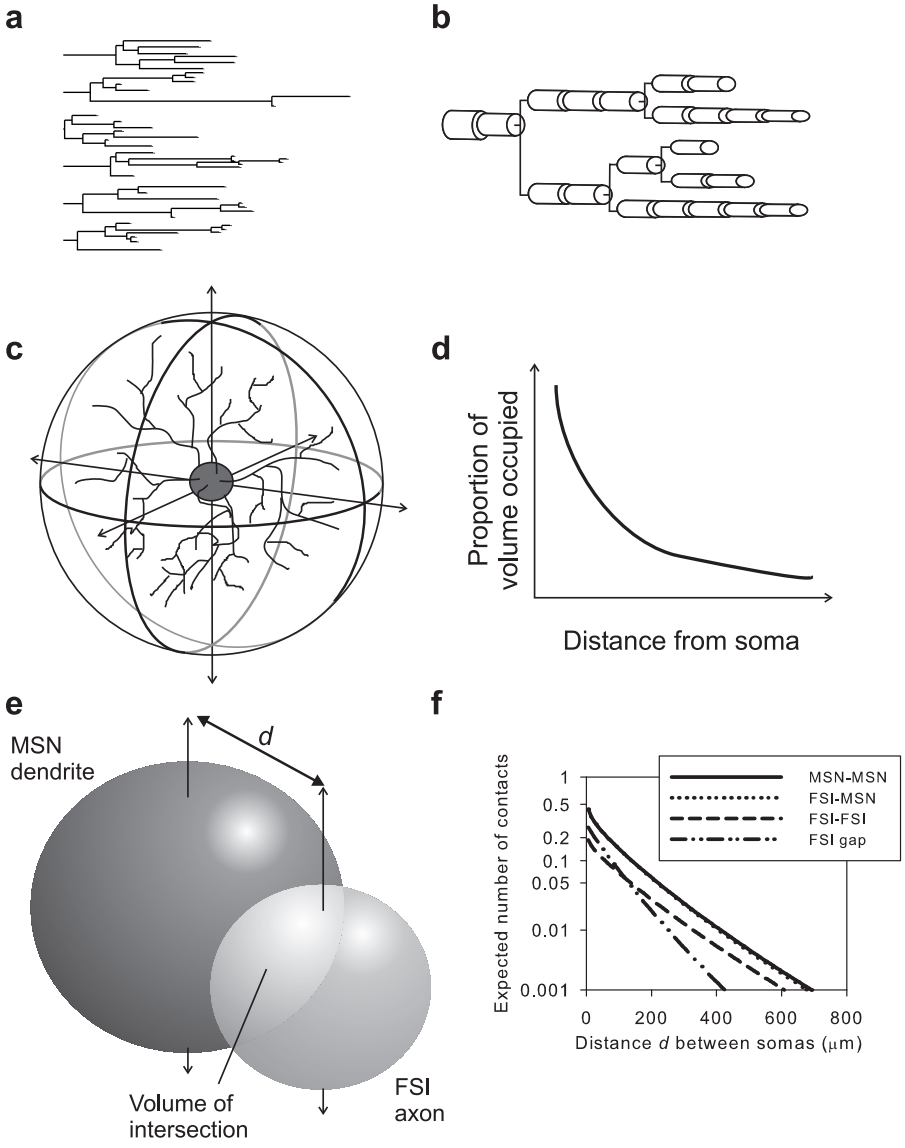

Figure 2: Anatomy model construction. (a) We create complete dendrograms using stochastic algorithms, bounded by known properties of the dendrites and axons. This example shows all six dendrites of the complete dendrogram for one MSN. (b) Each segment of each branch is modelled as a cylinder, whose diameter tapers with distance from the soma – summing over all branches gives the total volume of dendrite (or axon) at each distance from soma. (c) We then compute the proportion of spherical volume occupied by dendrite (or axon) at each distance from the soma. (d) Expected values for occupied volume are computed over many repetitions of the growth algorithm. The result is a continuous function of volume occupancy for each dendrite and axon type. (e) Volume of intersection of all dendrite and axon fields found for each distance between somas; volume discretised into 1 μm^3 voxels. (f) For each voxel, given its distance from the respective somas, we compute probability of intersection between fields (dendrite-axon or dendrite-dendrite) from volume occupancy functions (in panel d). We then sum over all probabilities to get expected number of contacts between neuron pairs as a function of distance between their somas. These are all functions of the form (1), with parameters given in Table 1; we use these functions to construct our network.

was well fit by the truncated power law

$$E_c(d_s) = \alpha d_s^{-\beta} e^{-d_s \gamma}, \quad (1)$$

for every connection type. Table 1 gives the specific parameter values for each of the four connection types in the striatal GABAergic microcircuit: between MSNs formed by the local axon collaterals synapsing on MSN dendritic trees; FSI axonal connections on MSN dendritic trees; FSI axonal connections on FSI dendritic trees; and gap junctions between FSI dendritic trees. Figure 2f shows the four resulting functions.

Table 1: Parameters for the expected number of contacts between neuron pairs.

Connection	α	β	γ
MSN-MSN	0.5567	0.1212	0.008
FSI-MSN	0.5528	0.1184	0.0082
FSI-FSI	0.2216	0.083	0.008
FSI gap	0.2892	0.0099	0.0132

We use these functions to construct our striatal network. First, we specify the three dimensions of our simulated region of striatum. The resulting volume V mm^3 defines the number of neurons (see section 2.1): given the 85,000 MSNs per mm^3 (Oorschot, 1996) we get $V \times 85000$ MSNs, and 3% of this is added as FSIs (Kawaguchi et al., 1995). All neurons are then randomly assigned a three-dimensional position within the defined volume, with a minimum distance of 10μm enforced.

For all pairs of neurons with potential connections we then apply (1) with the appropriate parameters from Table 1 for the connection type (MSN-MSN local collaterals, FSI-MSN axo-dendritic, FSI-FSI axo-dendritic, FSI-FSI dendro-dendritic gap junctions). As shown in Figure 2f, the expected number of connections was always much less than one, and so we used these functions as giving the *probability* of connection given the distance between somas – then the total number of such connections in a sufficiently large network would yield the same expected connection function. We have successfully used this to build and run models up to 1mm^3, though the models we use here are kept small so that a thorough analysis of the outputs remains tractable.

3. Model neurons

The model striatal network forms the basis for our study of its dynamics. If we are to build at such scales, we require individual neuron models that are simple enough to be computationally tractable, but sufficiently complex to capture key membrane properties that contribute to the characteristic behaviour of a neuron species. Our neuron model of choice is the recent canonical spiking model of Izhikevich (2007), which has been employed in some notably large-scale models (Izhikevich et al., 2004).

We previously extended these model neurons by incorporating dopaminergic modulation of intrinsic and synaptic ion-channels, which we review below. In this paper we extend the model further by introducing a model of gap junctions between FSIs and tune parameter values to data from gap-junction coupled cortical FSIs.

3.1. Reduced models of striatal neurons

In his recent book, Izhikevich (2007) gives a biophysical form of his canonical model for spike generation. Given that v is the membrane potential, and u is the contribution of the neuron class's dominant slow current, we have

$$C\dot{v} = k(v - v_r)(v - v_t) - u + I \quad (2)$$

$$\dot{u} = a\left[b(v - v_r) - u\right], \quad (3)$$

with reset condition

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

where C is capacitance, v_r and v_t are the resting and threshold potentials, I is a current source, a is a time constant, and c is the reset potential (i.e. the value of the membrane potential immediately after an action potential is fired). Parameters k and b are derived from the I-V curve of the neuron and d is tuned to achieve the desired spiking behaviour. We solve all neuron models using the forward Euler method with a time-step of 0.01 ms – this small time-step is necessary because of the fast dynamics of the FSI (Humphries & Gurney, 2007).

3.1.1. Dopamine-modulated MSNs

Izhikevich (2007) provided parameter values that modelled a MSN response to current injection. We introduced a framework for reformulating and extending this model to replicate the output of a detailed dopamine-modulated multi-compartment model (Moyer et al., 2007) – see (Humphries et al, submitted) for details. The MSN population is split in two by the expression of the dominant dopamine receptor type (D1 or D2). These receptors have different affects on both intrinsic and synaptic ion channels (see Surmeier et al., 2007, for review). We express the relative level of dopamine receptor occupancy by the parameters ϕ_1 (for D1) and ϕ_2 (for D2), normalised to the interval [0,1]. We add dopaminergic modulation of intrinsic ion channels in D1 MSNs by extending (2) to

$$C\dot{v}_{D1} = k(v_{D1} - v_r)(v_{D1} - v_t) - u + I + \phi_1 g_{DA}(v_{D1} - E_{DA}), \quad (4)$$

where the term $\phi_1 g_{DA}(v_{D1} - E_{DA})$ is sufficient to simulate the hyperpolarising effect of D1 activation when at an already hyperpolarised membrane potential, and the depolarising effect of D1 activation when at an already depolarised membrane potential (Surmeier et al., 2007).

For the D2 MSNs, we add dopaminergic modulation of intrinsic ion channels by extending (2) to

$$C\dot{v}_{D2} = k(1 - \alpha\phi_2)(v_{D2} - v_r)(v_{D2} - v_t) - u + I, \quad (5)$$

where we only decrease k by a factor of $(1 - \alpha\phi_2)$, which is sufficient to model the increased sensitivity to injection current following D2 activation (Moyer et al., 2007).

We model synaptic input to all the MSNs as

$$I = I_{ampa} + B(v)I_{nmda} + I_{gaba\text{-}fs} + I_{gaba\text{-}ms}, \quad (6)$$

where both I_{ampa} and I_{nmda} are derived from cortical input, $I_{\text{gaba-fs}}$ from FSI input, and $I_{\text{gaba-ms}}$ from local MSN collaterals. Each synaptic input of type z (ampa, nmda, gaba-fs,gaba-ms) is modelled by

$$I_z = \bar{g}_z h_z (E_z - v), \tag{7}$$

where $\bar{g}_z$ is the maximum conductance and E_z is the reversal potential. We use the standard single exponential model of post-synaptic currents

$$\dot{h}_z = \frac{-h_z(t)}{\tau_z}, \quad \text{and } h_z(t) \leftarrow h_z(t) + S(t)/\tau_z, \tag{8}$$

where τ_z is the appropriate synaptic time constant, and $S(t)$ is the number of pre-synaptic spikes arriving at all the neuron's receptors of type z at time t. Finally, we have the term $B(v)$ that models the voltage-dependent magnesium plug in the NMDA receptors (Moyer et al., 2007)

$$B(v) = \frac{1}{1 + \frac{[\text{Mg}^{2+}]_0}{3.57} \exp(-v \cdot 0.062)}, \tag{9}$$

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

We add D1 receptor dependent enhancement of NMDA-evoked EPSPs (Moyer et al., 2007) by

$$I^{\text{D1}}_{\text{nmda}} = I_{\text{nmda}}(1 + \beta_1 \phi_1), \tag{10}$$

and we add D2 receptor dependent attenuation of AMPA-evoked EPSPs (Moyer et al., 2007) by

$$I^{\text{D2}}_{\text{ampa}} = I_{\text{ampa}}(1 - \beta_2 \phi_2), \tag{11}$$

where β_1 and β_2 are scaling coefficients determining the relationship between dopamine receptor occupancy and the effect magnitude. All parameter values are given in Table 2.

3.1.2. Dopamine-modulated FSIs

The FSIs only express the D1-family of receptors on their membranes (Centonze et al., 2003). We add D1-receptor modulation by extending (2) to

$$C\dot{v}_{\text{fs}} = k[v_{\text{fs}} - v_{\text{r}}(1 - \eta\phi_1)](v_{\text{fs}} - v_{\text{t}}) - u_{\text{fs}} + I, \tag{12}$$

where we increase the nominal resting potential v_{r} by a factor of $(1 - \eta\phi_1)$, following experimental data from (Bracci et al., 2002; Centonze et al., 2003).

Following Izhikevich (2007), we use a nonlinear u term

$$\dot{u_{\text{fs}}} = \begin{cases} -a u_{\text{fs}}, & \text{if } v_{\text{fs}} < v_{\text{b}}, \\ a\left[b(v_{\text{fs}} - v_{\text{b}})^3 - u_{\text{fs}}\right], & \text{if } v_{\text{fs}} \geq v_{\text{b}}, \end{cases} \tag{13}$$

that enables the FSI model to show Type 2 dynamics, particularly a non-linear step at the start of its current-frequency curve from 0 to around 15-20 spikes/s.

Synaptic input to the striatal FSIs predominantly activates GABAa or AMPA receptors (Blackwell et al., 2003), NMDA receptors are rare. The dendrodendritic gap junctions provide

Table 2: Intrinsic and synaptic parameters for the medium spiny neuron model. Dimensions are given where applicable.

Parameter	Value	Source
C	50 pF	Izhikevich (2007)
b	-20	”
c	-55 mV	”
v_{r}	-80 mV	”
v_{peak}	40 mV	”
k	1.14	Humphries et al, submitted
v_{t}	-33.8 mV	”
a	0.05	”
d	377	”
α	0.03	”
g_{DA}	22.7 nS	”
E_{DA}	-68.4 mV	”
β_1	3.75	”
β_2	0.156	”
E_{ampa},E_{nmda}	0 mV	Moyer et al. (2007)
$E_{\text{gaba-fs}}$, $E_{\text{gaba-ms}}$	-60 mV	”
τ_{ampa}	6 ms	”
τ_{nmda}	160 ms	”
$\tau_{\text{gaba-fs}}$, $\tau_{\text{gaba-ms}}$	4 ms	”
g_{ampa}	6.1 nS	Humphries et al, submitted
g_{nmda}	3.05 nS	”
$g_{\text{gaba-ms}}$	4.36 nS	”
$g_{\text{gaba-fs}}$	21.8 nS	$\sim$ 5× MSN conductance Koos et al. (2004)
$[\text{Mg}^{2+}]_0$	1 mM	Jahr & Stevens (1990)

Table 3: Intrinsic and synaptic parameters for the fast spiking interneuron model. Dimensions are given where applicable. n.d.: no data.

Parameter	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	”
η	0.1	fitted to Bracci & Panzeri (2006)
ϵ	0.625	fitted to Gorelova et al. (2002)
E_{ampa},E_{nmda}	0 mV	n.d.; set as for MSNs
$E_{\text{gaba-fs}}$, $E_{\text{gaba-ms}}$	-60 mV	”
τ_{ampa}	6 ms	”
$\tau_{\text{gaba-fs}}$	4 ms	”
g_{ampa}	61 nS	n.d.; tuned to achieve realistic firing rates (section 5.2)
$g_{\text{gaba-fs}}$	20 nS	n.d.; assumes equivalent effect of FSI-FSI contacts as FSI-MSN contacts
g	30 nS	section 3.1.3
τ	11 ms	”

a further source of “synaptic” current (Koos & Tepper, 1999). Thus the synaptic current contributions are

$$I = I_{\text{ampa}} + I^*_{\text{gaba}} + I_{\text{gap}}, \tag{14}$$

where we add D2-receptor dependent modulation of GABAergic input (Bracci et al., 2002; Centonze et al., 2003) by

$$I^*_{\text{gaba}} = I_{\text{gaba}}(1 - \epsilon\phi_2), \tag{15}$$

where I_{gaba} is derived from FSI input. All parameter values are given in Table 3.

3.1.3. Tuning FSI gap junctions

A gap junction between FSIs i and j is modelled as a compartment with voltage v^*_{ij}, which has dynamics

$$\tau \dot{v}^*_{ij} = (v_i - v^*_{ij})(v_j - v^*_{ij}), \tag{16}$$

where τ is a time constant for voltage decay, and v_i and v_j are the membrane potentials of the FSI pair. The current introduced by that cable to the FSI pair is then

$$I^*_{\text{gap}}(i) = g(v^*_{ij} - v_i) \qquad I^*_{\text{gap}}(j) = g(v^*_{ij} - v_j), \tag{17}$$

where g is the effective conductance of the gap junction. The total gap junction input I_{gap} to an FSI is then the sum over all contributions I^*_{gap}.

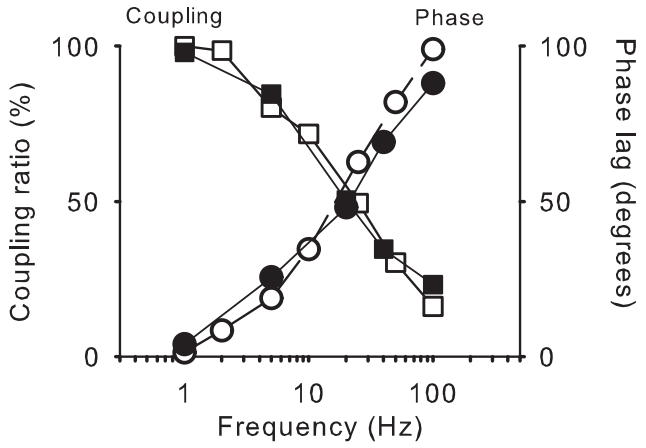

Figure 3: Tuning the gap junction model. Galarreta & Hestrin (1999) injected sinusoidal current into a cortical FSI at various frequencies, and recorded from another connected to it by a gap junction. They computed both the coupling ratio (■) and phase lag (•) of the second neuron’s membrane potential with respect to the injected neuron. We similarly connected a pair of model FSIs with our gap junction model, injected sinusoidal current into one, and hand-tuned the gap junction parameters (g and τ) to fit the data. A qualitatively good match was achieved by the model for both coupling ratio (□) and phase lag (∘).

We hand-tuned τ and g using a pair of FSI models connected by a gap junction. Our target data came from a study by Galarreta & Hestrin (1999) in which sinusoidal current at different frequencies was injected into one of a gap-junction coupled pair of cortical FSIs, and membrane voltages recorded from both: this data is ideal as it provides both voltage coupling strength and voltage phase-lag, which are affected by both g and τ. We injected sinusoidal current I into one FSI with an amplitude of 400 pA at different frequencies and computed the coupling coefficient (ratio of maximum amplitudes in the membrane voltages of the two neurons) and the phase-lag (voltage-peak offset as a function of the injection current frequency). Figure 3 shows we achieved a good qualitative match to both coupling coefficient and phase-lag from the experimental data, with $\tau = 11$ ms and $g = 150$ nS.

While this data-set was the most appropriate for tuning the gap junction model, we cannot immediately use the value for g. Two caveats have to be accounted for: first, that there is an unknown number of other FSIs connected by gap junctions to the studied pair; second, that the study was done in tissue from juvenile rats, and so would over-express gap junctions (Belluardo et al., 2000). Both of these would contribute to the decay of the coupling coefficient. Thus, we find we need to re-scale g to account for the approximate reduction in gap junctions in adult tissue and to account for other connections. In further simulations we explored fully-connected gap-junction networks of 3, 4 or 5 FSIs, as might be found in juvenile tissue. We found that repeating the same paired recording protocol in these networks did indeed predict a dramatic reduction in g: the multiple gap junctions acted to reinforce the effects of the injection current on the un-injected neuron. A five-fold reduction to $g = 30$ nS

produced an equivalent fit to the data in Figure 3 for all three networks, and so we used that figure here. This is also consistent with the comparatively weak coupling coefficients of 3% and 20% that have been reported for the few gap junction coupled striatal FSIs recorded to date (Koos & Tepper, 1999).

3.2. Input to network

In addition to its synaptic connections defined by our network model, each neuron received external input representing its cortical afferents. In many spiking neuron models, afferent input is generated by a set of Poisson processes. However, for large-scale models where each neuron receives hundreds or thousands of afferent inputs, this becomes unfeasible because of the memory requirements. Recently we have developed a series of tools addressing just this problem, using a method that collapses many afferent trains into an single equivalent spike-event count.

Each spike-event generator directly produces the spike-events that occur across N afferents to the neuron. At each time-step Δt, and given a mean spike rate r, we compute the probability of a spike per afferent as $p(s) = r\Delta t$. The total number of spike-events S at each time-step is then just drawn from a binomial distribution $S = B(N, p(s))$. The resulting time-series of spike-events is equivalent to the pooling of N spike trains modelled as independent renewal processes, the superset that includes Poisson processes.

We define N and r for the striatal network for the tonic background *in vivo* state, by combining data from anatomy and electrophysiology:

1. In a recent organotypic cortico-striatal-nigral co-culture study, Blackwell et al. (2003) reported that a striatal MSN receives an average of around 800 synaptic events per second during its depolarised ("up") state, but they could not distinguish excitatory and inhibitory potentials.
2. The ratio of asymmetric (putative excitatory):symmetric (all others) synapses in rat striatum is $\sim 3.9 : 1$ (Ingham et al., 1998).
3. If we conservatively assume that half the asymmetric synapses are cortical in origin, then we have a ratio of 2:1 potentially active synapses in the co-culture
4. Assuming this corresponds (roughly) to the proportion of glutamate:GABA activity, then cortical activity accounts for ~ 530 synaptic events per second.
5. Given the estimate of 4250 cortical inputs per MSN (Zheng & Wilson, 2002), the average firing rate of those cortical neurons is therefore ~ 0.12 spikes/s.
6. From *in vivo* extracellular recordings, we know that dedicated cortico-striatal neurons tonically fire a maximum of 5 spikes/s and pyramidal tract neurons with striatal collaterals tonically fire around 15 spikes/s (Bauswein et al., 1989; Turner & DeLong, 2000). The former dominate in number over the latter (Bauswein et al., 1989; Zheng & Wilson, 2002), suggesting an overall mean rate around 2-5 spikes/s.
7. Taking the lower mean single neuron rate of 2 spikes/s, and the estimate of 530 synaptic events per second, we see that just 265 active cortico-striatal neurons are required to achieve this – or just 6 % of the total afferent cortical population.

Overall then, tonic cortico-striatal activity sufficient to drive MSN firing requires just $N \simeq 250$ trains, at a rate of $r \simeq 2$ spikes/s. We hence use $N = 250$ and $r = 1.9$ spikes/s for the MSN input throughout our simulations. In addition, we use the same N, r in the spike-event generators for the FSIs, as there is no data on cortical input to these neurons.

4. Detecting groups of synchronised cells in multi-unit data

We sought to identify potential candidates for the basic computational elements of the striatum from the dynamics of our large-scale models under background input. For our present purposes, we wanted to find groups of co-active or mutually antagonistic MSNs that could form the basis for competitive dynamics within the striatum. In addition, we studied this input regime to see if the reported striatal cell clusters spontaneously formed in vitro (Carrillo-Reid et al., 2008) could be identified in our model. However, analysis methods suitable for exploratory analysis of such large spike-train data-sets are lacking (Brown et al., 2004). We therefore developed a new algorithm for finding synchronised groups at multiple time-scales within a multiple spike-train data-set.

At its most general, our algorithm follows a two-step procedure. First, some measure of correlation between each pair (or more) of neurons is computed, resulting in a correlation matrix. Second, some method acts on this matrix to identify "strong" spike-train correlations within *groups* of neurons, thereby grouping the data-set into sets of neurons whose output is more related to each other than with the remaining neurons. A group is thus 3 or more neurons that are co-correlated. With this in mind, we detail our specific algorithm (our present choices for these two steps are specified in the Appendix):

1. For all N spike trains, correlate all $N(N-1)/2$ unique pairs, resulting in a correlation matrix $\mathbf{C}$ with entry $C_{ij} = C_{ji}$ being the correlation measure between spike-trains i and j. We use here the normalised Hamming distance (see Appendix)
2. Threshold each correlation matrix with threshold θ, to create a graph. In general the graph's adjacency matrix $\mathbf{A}$ is generated by some function $A_{ij} = f(C_{ij}, \theta) \in \{0, 1\}$, where a 1 denotes a link between nodes i and j on the graph, and 0 denotes the absence of a link. Here we use

$$A_{ij} = \begin{cases} 1 & \text{if } C_{ij} < \theta \\ 0 & otherwise. \end{cases} \tag{18}$$

Note that the sign of the comparison is dependent on the correlation method chosen: this direction (<) is suitable for our choice of Hamming distance, but would be reversed if, for example, the absolute correlation coefficient was used instead. Each retained link thus indicates the presence of a "strong" correlation between the respective neurons, where correlation strength is parameterised by θ

– the lower θ, the stronger the correlations that make up the graph.

3. Remove all nodes with less than 2 links from $\mathbf{A}$, resulting in reduced graph $\mathbf{A}^*$, which has n^* nodes and m^* links. We do this because any neuron with fewer than two links is not participating in any synchronised *group* at the specified threshold.
4. For simplicity, we proceed only if (a) $n^* > 5$, so that multiple groups are possible (given that each group is at least 3 neurons), and (b) $m^* > \ln(n^*)$ so that the graph should have a single *giant component* (Watts & Strogatz, 1998), that is, most nodes should be reachable from any other node along the links. These conditions could be relaxed in more complex uses of the current algorithm.
5. If we proceed, then we run a graph modularity algorithm (Newman, 2006a,b) on $\mathbf{A}^*$ (see Appendix). This method does not require prior specification of the number of groups, and is hence more suitable for exploratory data analysis than the family of k-means clustering methods or classic graph-partitioning algorithms. The result is a vector of length n^*, with the ith entry an integer $1 \ldots M$ identifying the membership of node i in the M identified groups.

4.1. Identifying structure in the dynamics of the cell groups

The above method results in a set of M groups for each combination of tested binsize δt and threshold θ, for each model that we simulate. Given this potentially vast data-set, how do we summarise the outcome of that combination, and identify which combinations have sufficiently rich dynamics to analyse further? Our present search is for basic computational elements that could be formed by antagonistic striatal cell assemblies. We thus seek a scalar metric β that encapsulates our current criteria: maxmising the number of groups M found; maximising the number of neurons contributing to groups, so that the ratio $n^*/N \to 1$; and clear evidence for both synchronisation and anti-correlation within the correlation matrix.

For the latter, consider a particular choice of $(\delta t, \theta)$ and the resulting correlation matrix $\mathbf{C}$. Let $\tilde{h}$ be the median non-zero Hamming distance for $\mathbf{C}$, h^* be the minimum non-zero Hamming distance for $\mathbf{C}$, and $\Delta = \tilde{h} - h^*$. The greater Δ, the greater the likely existence of both strongly correlated and anti-correlated neurons. Thus, we express our full set of criteria as

$$\beta = M \frac{n^*}{N} \Delta. \tag{19}$$

5. Results

We now have the necessary tools – models of anatomy, neurons, and input, and suitable analysis methods – to begin addressing the problem of identifying the computational elements of the striatum. We use in this paper a small striatal region of 250 μm^3, which gives us 1400 neurons, 1359 MSNs and 41 FSIs. This made a thorough analysis of both the network itself and all its outputs computationally tractable, and we keep this size throughout for consistency. We randomly split the MSNs into two equal sized populations, and assigned one set as the D1 MSNs and the other as the D2 MSNs. The resulting network was used as the basis for every simulation detailed below.

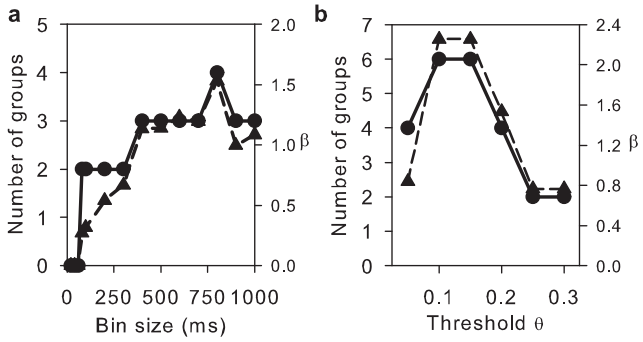

Figure 4: Spontaneous dynamic clusters form at a characteristic time-scale in the absence of dopamine. (a) Both (the number of groups (•) and the metric β (▲) have a unique maximum at $\delta t = 800$ ms, suggesting this is the characteristic time-scale of the dominant dynamics. (b) For that time-scale, changing the threshold θ for correlation reveals further structure. Both the number of groups and β are maximum at $\theta = 0.1, 0.15$. The reduction in both number of groups and β at the strong correlation threshold $\theta = 0.05$ illustrates that cell assemblies can exist at multiple strengths of correlation.

As noted above, we here assess the spontaneous dynamics of the network, simulating spontaneously firing cortical input. Every simulation was run for 10 seconds of network time, and every MSN data-set analysed with the cell group detection algorithm detailed above, using the set of binsizes $\delta t \in \{20, 40, 60, 80, 100, 200, 300, 400, 500, 600, 700, 800, 900, 1000\}$ ms. For all simulations, we set $\theta = 0.2$ as this was a conservative threshold for identifying potential structure. For some simulations, we use this threshold to identify a particular time-scale of interest, as defined by the δt that maximises (19), and re-analyse the corresponding correlation matrix $\mathbf{C}$ for a range of θ to pick out the detailed dynamical structure.

5.1. Spontaneous clusters of synchronised MSNs emerge without dopamine

Spontaneously formed striatal cell assemblies were observed using calcium imaging by (Carrillo-Reid et al., 2008), from an *in vitro* preparation that was excited by bath application of an NMDA agonist. To see if we could observe equivalent dynamics in our model, we began with a model without dopamine, setting $\phi_1 = \phi_2 = 0$, so that the model was close to this *in vitro* state.

We found that this basic model indeed supported multiple groups of synchronised cells. Figure 4 shows that our algorithm found groups at many time-scales, with a unique maximum in both the number of groups and β at $\delta t = 800$ ms. The majority of the MSNs were retained by the algorithm when groups were detected (range 1286-1354).

We took the correlation matrix $\mathbf{C}$ for $\delta t = 800$ ms and re-ran the cluster detection algorithm with a range of $\theta \in \{ 0.05, 0.1, 0.15, 0.2, 0.25, 0.3 \}$, to look in more detail at the dynamical structure. Both thresholds of $\theta = 0.1$ and $\theta = 0.15$ maximised the number of groups and β, producing an identical set of 6 groups (Figure 4b).

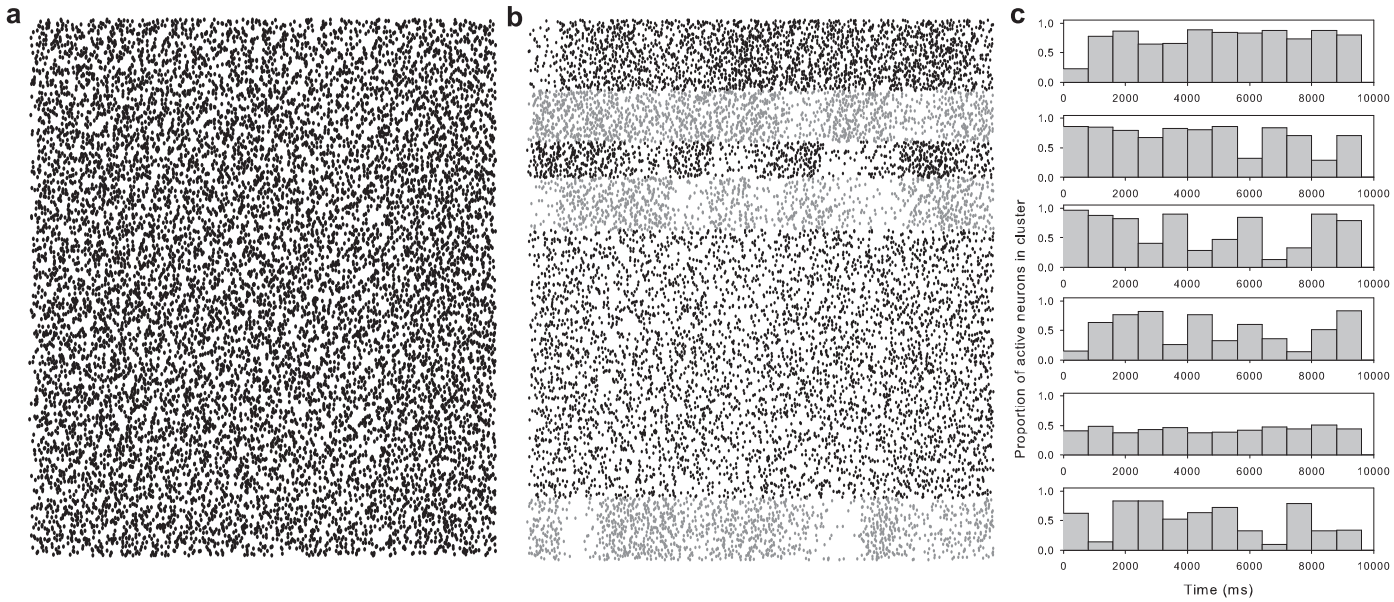

Figure 5: Synchronised MSN clusters at long time-scales in the absence of dopamine. (a) Unsorted raster plot for all MSNs retained for grouping by algorithm – spike trains are arranged in rows, each dot signifies the occurrence of a spike, and each train is shown for the full 10 seconds. (b) Raster plot of the same set of MSNs, sorted by groups detected by our algorithm – it found 6 groups with $\theta = 0.1$. All the spike trains belonging to each group are colour coded either black or grey, alternating by group. (c) Histograms of the proportion of active cells in each group for each time bin, illustrating the different structure of the dynamics in each group. Note, for example, the comparatively high firing rate of MSNs in group 1 (groups are numbered top-down). The total of 1328 MSNs were divided into groups sized: 174, 126, 88, 129, 659 and 152.

Figure 5 shows the striatal cell assemblies so identified. We see that, faced with the initial set of 1328 MSNs that contributed to graph $\mathbf{A}^*$, the algorithm could successfully find cell assemblies of different types and strength. Groups 2-4 and 6 each had at least two long (> 500 ms) periods of near-silence, but at different times for each group, and with different strengths of correlation between the individual spike trains – compare, for example, the two periods of near-silence for group 6 with the noisier but clearly identifiable periods for group 2 later in the simulation. By contrast, groups 1 and 5 are distinguished from the others by not having multiple synchronised periods, but are clearly distinguishable from each other on the basis of the different firing rates of their constituent neurons, and the initial period of silence shared by neurons in group 1.

5.2. *Increasing dopamine concentration simplifies spontaneous dynamics*

We proceeded to study the changes in the model's dynamics under changes in dopamine concentration, to search for potential correlates of known dopamine-related effects on MSN dynamics and, in turn, the striatum's role in motor control (see Introduction). We tested five levels of dopamine with equal activation applied to both D1 and D2 receptors, giving the set $\phi_1, \phi_2 = \phi \in \{\ 0.05, 0.1, 0.2, 0.5, 0.8\ \}$.

We found that increasing dopamine simplified the dynamical structure of the MSN output, as determined by our algorithm and metric β. Figure 6a shows that increasing dopamine generally increased the left-shift and decreased the height of the curve relating correlation binsize δt to the metric β. Indeed, Figure 6b shows that the maximum β found at any time-scale reduced monotonically with increasing dopamine. The time-scale δt at which this maximum β occurred also decreased with increasing dopamine.

The loss of synchronised cell groups was apparent even at low dopamine levels. For $\phi = 0.1$ maximum β was found for $\delta t = 400$ ms. When we evaluated the grouping algorithm for a range of θ as before, we found that $\theta = 0.2$ maximised β. For this threshold, only 3 groups were found. Figure 6c shows the lack of synchronised cell group dynamics in the the raster plots for all MSNs in these groups, emphasising the dramatic simplification of the structure. Nonetheless, the neurons grouped by the algorithm did share common properties: groups 1 and 2 were divided by firing rate (median rates 3.5 and 1.4 spikes/s, respectively), and group 3 were clustered by the period of silence at the start of the simulation.

This grouping by firing rate continued at higher levels of dopamine. For $\phi = 0.5, 0.8$ only 2 groups were found, clearly separated by median firing rates, and the time-scale δt that maximised β for both these levels approximately coincided with the average MSN firing rate across the whole network.

We found that the firing rate characteristics of our striatal model were commensurate with experimental data on *in vivo* activity under baseline conditions. The background (tonic) state of striatal dopamine we expect to be between 0.05 and 0.2 in our model, as the tonic concentration can be up to an order of magnitude lower than the concentration during phasic release (Gonon, 1997; Venton et al., 2003), represented by the upper limit of our normalised dopamine parameters ϕ_1, ϕ_2. In this approximate tonic dopamine range, Figure 7 shows the firing statistics are consistent with known properties of striatal

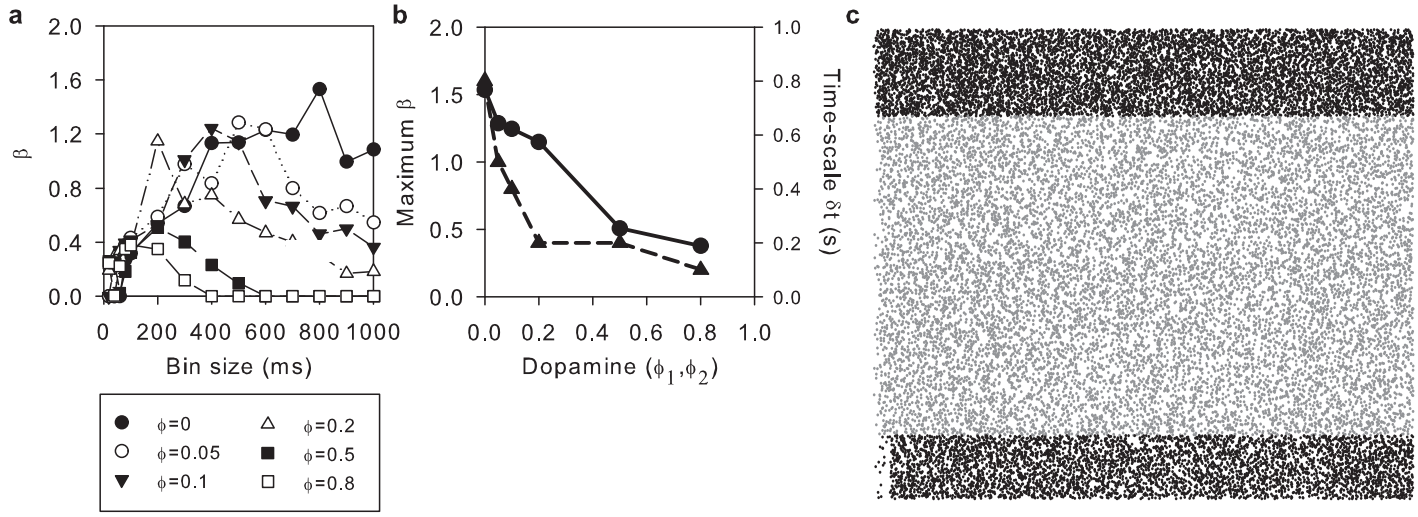

Figure 6: Increasing dopamine level reduced dynamical structure in the MSN output. (a) The relationship between correlation bin-size δt and metric β for all tested dopamine levels. (b) Both maximum β (•) and corresponding time-scale δt (▲) that maximised β decreased with increasing dopamine. (c) Raster plot for all MSNs grouped for dopamine level $\phi = 0.1$ at the threshold ($\theta = 0.2$) that maximised β. No synchronised cell group dynamics are evident. Nonetheless, the grouping algorithm does find neurons with common properties: groups 1 and 2 are split by firing rate, group 3 clustered by the brief silence at the start (groups are numbered top-down). The total of 1333 MSNs were divided into groups sized: 246, 907, 180.

cells. That is, MSNs are relatively quiet, while some FSIs fire strongly (Sandstrom & Rebec, 2003; Mallet et al., 2005). Figure 7 also shows that our model predicts a linear change in the firing rates of both neuron types with increasing dopamine. By contrast, MSNs showed a monotonic decrease, and FSIs no change, in the median inter-spike interval coefficient of variation (ISI CV) with increasing dopamine. That is, the MSN spike trains became increasingly regular, but there was no global change in FSI spike patterning despite their decrease in firing rate.

The FSI firing statistics showed two striking properties. First, the decrease in FSI firing rate with increased dopamine levels runs counter to the effect of dopamine on the single neuron model (equations 12 and 15). We attribute this to the action of the gap junctions: without them, the FSI median firing rate increases with increasing dopamine (Figure 7a). Second, previous studies have reported spike-to-spike synchrony in model fast-spiking interneuron networks coupled by both synaptic contacts and gap junctions (Traub et al., 2001). However, despite similar connectivity, we see no network-wide synchrony between our FSIs, irrespective of the level of dopamine (Figure 7c). Our results are nonetheless consistent with the reported uncorrelated behaviour of striatal FSIs *in vivo* (Berke, 2008).

5.3. Dissecting the network: contribution of microcircuit elements

To begin the task of understanding how the GABAergic striatal microcircuit produces these dynamics, we look at the contributions of the two elements that have dominated recent discussions (Tepper et al., 2004): "weak" feedback inhibition by the local MSN collaterals, and "strong" feedforward inhibition of MSNs by the FSIs.

5.3.1. An MSN-only network shows cell assembly sequencing

We first created an MSN-only network by removing all FSI connections, and repeated the simulations at all previously tested levels of dopamine. Figure 8 shows that, compared to the full model, the MSN-only network could produce spike train dynamics of greater complexity (as measured by β), and did not have a monotonic relationship between dopamine level and β. By comparing the spike trains rasters in Figure 8b and Figure 8c we can see that the metric β captures the visual impression of greater complexity in the structure of the neural dynamics.

With dopamine absent or at low levels, the grouped rasters also clearly show the presence of multiple groups that fire in patterns of silence and high activity. Their absence in the intact model clearly suggests that the FSIs desynchronise the MSN network. The ordering of these patterns is also suggestive of inhibition-based competition between a few of these groups, particularly in the absence of dopamine (Figure 8b). With a low level of dopamine, Figure 8c clearly shows a sequence of silent periods involving all 9 groups identified by our analysis method, with some groups showing further alternating periods of silence and activity.

To quantify these interactions at the scale of whole groups, for every group we compute a vector of the proportion of active neurons per time-bin (as in Figure 5c). When we compute the correlation coefficient between all vectors, we see in Figure 9 multiple negative correlations between cell groups in the MSN-only networks with absent or low dopamine (though we note that these are limited to the time-scale picked by algorithm). By contrast, the correlation coefficient distribution for the intact

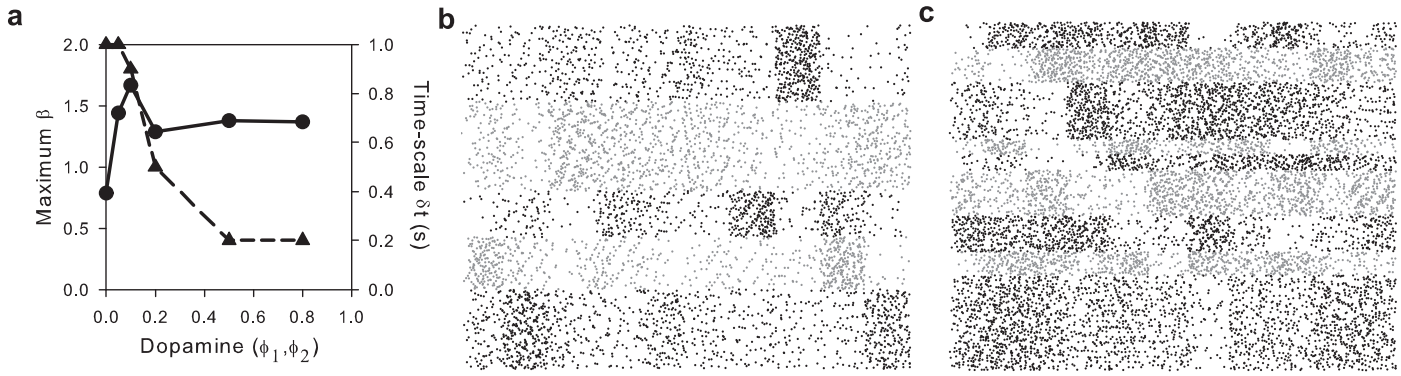

Figure 8: An MSN-only network generates a range of complex spike-train groupings, showing sequences of group firing. (a) The relationship between maximum β (•) and increasing dopamine level is no longer monotonically decreasing, as it is for the intact model, suggesting an increased complexity of spike-train structure for intermediate dopamine levels. However, the time-scale at which maximum β occurs (▲) does still generally decrease with increasing dopamine. (b) With no dopamine, the grouped raster plot shows multiple clear groups. Groups 1 and 3 in particular have roughly anti-correlated sequences of silence and activity, suggestive of competition between them. (c) At low levels of dopamine (here $\phi_1, \phi_2 = 0.1$), we see more groups at a wider range of sizes, with more complex dynamics: we have ordered the 9 groups to show a sequence of silent periods occurring across them all. Note that this is only one of many possible orderings – group 8, for example, is silent at the beginning of the simulation, just like group 1; groups 4 and 5 are simultaneously silent for approximately the same length of time, but only group 4 has a further period of complete silence. Groups in (b) and (c) shown for the threshold that maximised β (both $\theta = 0.1$), as before; groups numbered top-down.

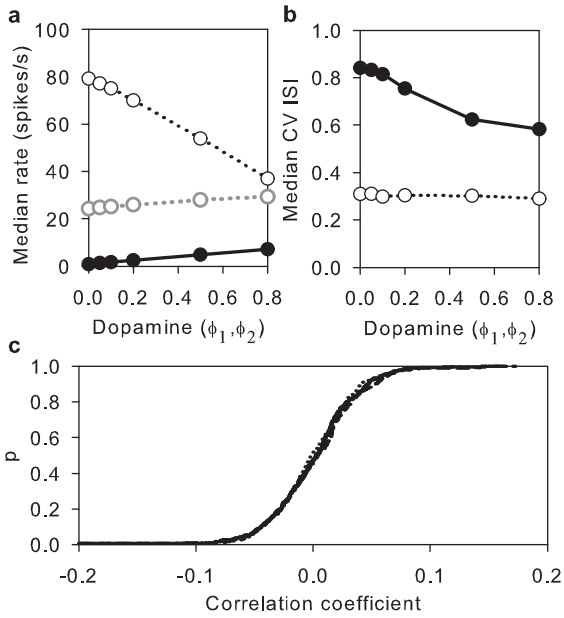

Figure 7: Effect of dopamine level on population firing statistics. (a) Both MSN (•) and FSI (black ○) populations had linear changes in median firing rate with increasing dopamine. Removing the FSI gap junctions reduced their median firing rate, but reversed the effect of changing dopamine levels (grey ○). (b) Increasing dopamine monotonically decreased the MSN population's median inter-spike interval (ISI) coefficient of variation (CV), but had no effect on the FSIs' median ISI CV. The model thus predicts increased regularity of MSN spike trains under constant elevated dopamine levels and background synaptic barrage. (c) The FSIs showed no evidence of network-wide synchrony, for any level of dopamine. The distributions of pair-wise correlation coefficients for all active FSIs are shown for $\phi = 0$ (solid), $\phi = 0.1$ (dashed), and $\phi = 0.8$ (dotted): network synchrony would be evident as a positively-biased asymmetric distribution. (Pair-wise correlation coefficients computed using spike trains binned at 1/min{mean spikes/s,median spikes/s}, so that the bins matched the firing rates for that dopamine level. Hence, we looked for spike-to-spike correlations on the rough time-scale of the inter-spike interval).

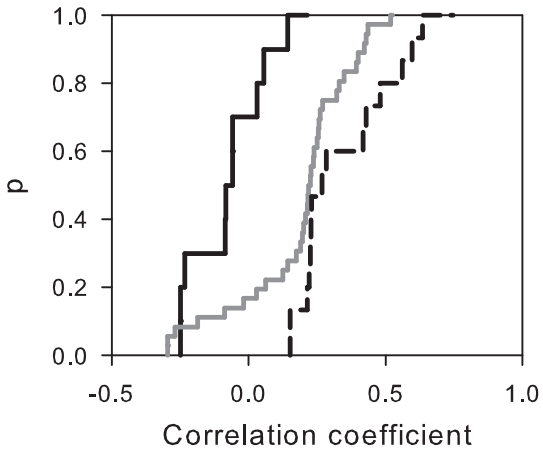

Figure 9: Detected cell assemblies show alternating firing sequences in MSN-only networks. We computed correlation coefficients between the activity vectors of every pair of cell groups. The empirical cumulative distributions of these for the MSN-only networks show a range of negative correlations between groups; showing $\phi = 0$ (black solid line) and $\phi = 0.1$ (grey solid line). By contrast, the distribution for the intact model with $\phi = 0$ (black dashed line) had no negative correlations, indicating the absence of distinct sequences of firing and hence no evidence for competitive interactions between groups – activity vectors for this model are shown in Figure 5c.

model's activity vectors is entirely positive.

The MSN-only network firing rate statistics were also distinct from the intact model. Figure 10a shows that, although the median MSN firing rate did monotonically increase with increasing dopamine, the rate was always less than that of the intact model. We thus see a counter-intuitive attenuation of firing rate following the removal of inhibition provided by the FSIs. The MSN spike-trains were also roughly consistent in their irregularity (as measured by ISI CV), independent of the level of dopamine and of the changes in firing rate (Figure 10b).

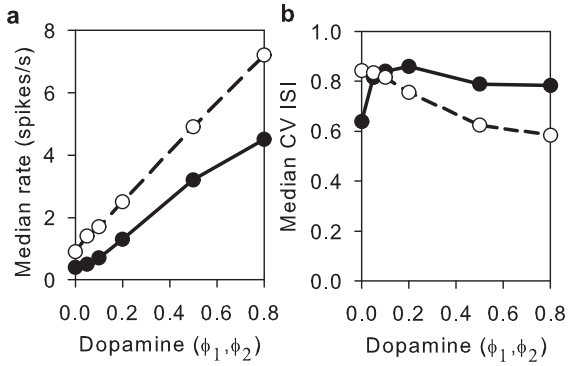

Figure 10: Effect of dopamine level on MSN population firing statistics in the MSN-only network. (a) The MSN-only network firing rate (•) increased with increasing dopamine, but the median rate was always lower than for the intact model (○). We infer that inhibitory input from FSIs has a facilitatory effect on MSN activity. (b) In a further departure from the intact model, the MSN-only network did not have a monotonic relationship between dopamine and the median ISI CV: rather, the spike trains seemed consistently irregular, independent of the firing rate changes.

5.3.2. The FSI network dominates the MSN network

We removed all local MSN collaterals from the full network to assess the impact of the FSI feed-forward inhibition. Figure 11 shows that, similar to the full model, both maximum β and corresponding time-scale generally decreased with increasing dopamine. Without dopamine, Figure 11b shows that 5 groups in all were detected, similar in distribution of size and dynamics to the full model in Figure 5b. The notably greater β at low levels of dopamine, compared to the full model, corresponds to the detection of more groups that could be distinguished by firing rate, as shown in Figure 11c. Yet, similar to the full model, there was a loss of within-group structure in the presence of dopamine. Comparing the groupings in Figure 11 to those of the MSN-only model in Figure 8 and the intact model in Figure 6, both emphasises the previous conclusion that the FSIs desynchronise the MSN-only network, and shows that they dominant the dynamics of the intact model.

5.3.3. Relating dynamics to network structure

Having found a wide variety of complexity in the spike-trains from our various models, we sought to relate the striatal network structure to the dynamic cell assemblies and the relationships between them. We focussed on the output from the MSN-only model with no and low dopamine ($\phi_1 = \phi_2 = 0.1$), shown in Figure 8b-c, as these were both using the simplest network and had the clearest set of discrete cell assemblies.

Standard graph metrics (Watts & Strogatz, 1998; Humphries & Gurney, 2008) were computed for the whole MSN network, each network formed by the neurons within a dynamic group, and each network formed between a pair of groups. These metrics were: the mean shortest path length L between each pair of neurons in the network; the clustering coefficients C^{ws} (the mean density of interconnections between all immediate neighbours of each neurons) and C^{Δ} (the number of closed 3-neuron feed-forward loops in the network); and the small-world-ness coefficients $S^{\text{ws}}, S^{\Delta}$ (Humphries & Gurney, 2008), corresponding to the two forms of clustering coefficient. $S^{\text{ws}}, S^{\Delta} > 1$ implies that the network is tending to a small-world regime of short path length but high clustering. We attempted to relate these to the cell assemblies and their sequencing, using the values for the whole network as baselines for comparison.

The MSN-only network formed a classic random graph, rather than a small-world, with small-world-ness values of $S^{\text{ws}}, S^{\Delta} \simeq 0.87$. The mean path length was $L = 1.92$ (maximum path length 3), and clustering coefficients were $C^{\text{ws}} = 0.092$ and $C^{\Delta} = 0.093$. All MSNs were reachable from all others on the network, and so no dynamic group was defined by its physical isolation from the network. We also used the modularity-based graph-cutting method (see Appendix) directly on the graph formed by the connections in the MSN-only network: this gave just 2 groups of approximately equal size, showing that the groups found in the dynamics were not related to equivalent groups in the physical network.

We hypothesised that each dynamic group was formed by neurons with fewer local collaterals between them, and so were less able to inhibit each other, as indicated by either or both of (a) a longer path length and (b) lower clustering within the network formed by those neurons. However, the networks formed by each dynamic group could not be distinguished by their graph metrics: path lengths and clustering coefficients did not systematically depart from those for the whole network, and the small-world-ness coefficients fell in the range expected for such small networks. We also found no difference in path lengths, in either direction, between all pairs of groups with negatively correlated activity vectors. Thus, we could not relate the seemingly antagonistic arrangement of cell firing to corresponding arrangements of local collateral inhibition, though we note further work is needed here to more precisely define the sequence of firing groups.

Finally, we found that the groups were not defined by their physical positions within the 250 μm^3 region of simulated striatum. All groups had approximately the same median distance from the centre of the region, and all had approximately the same distributions of distances as well. Thus, no group was seemingly defined by edge effects on the network.

6. Discussion

To study the striatal GABAergic microcircuit, we have brought together for the first time a detailed model of striatal anatomy, models of its main neurons, their modulation by dopamine, and connection by gap junctions, and models of cortical input. Further, we proposed a new algorithm for finding structure in the multiple spike train data-sets resulting from the striatum model. We used this method to gain unique insight into the computations of the microcircuit, and identify potential "basic computational elements" for further study. These techniques have general application to the study of microcircuits.

6.1. Detecting structure in multiple spike-trains

Our goal was to find co-active and antagonistic MSNs within the striatum. Naturally, this meant facing the problem of iden-

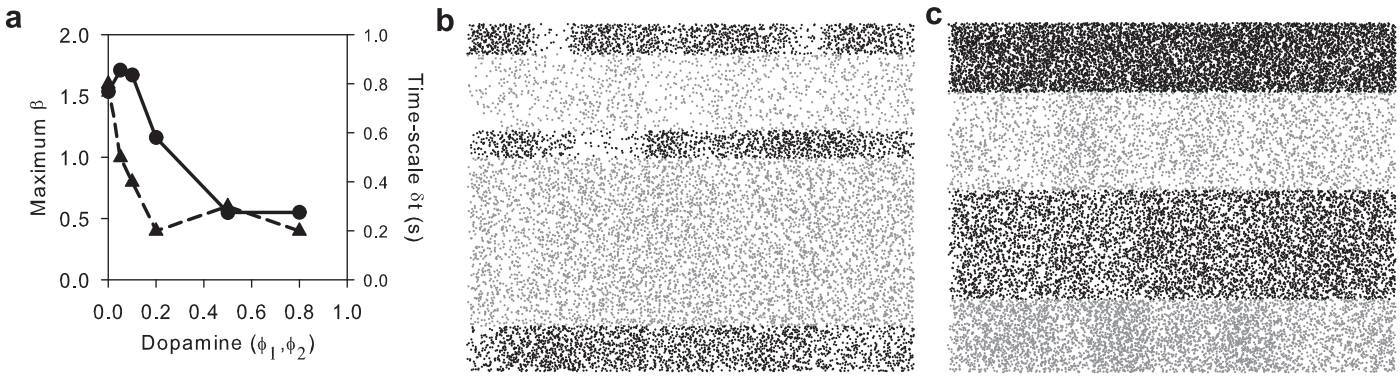

Figure 11: Feed-forward inhibition from the FSIs dominates and desynchronises the MSN network. (a) Increasing dopamine generally decreased both maximum β (•) and the corresponding time-scale at which it occurred (▲), similar to the full model. However, the values of β at low dopamine levels were larger, reflecting the presence of more spike train groups that could be distinguished by firing rates. (b) Without dopamine, the grouped raster shows 5 groups, similar in distribution of size and within-group dynamics to the full model. (c) The introduction of a low level of dopamine (here $\phi_1, \phi_2 = 0.1$) does not cause the same drop-off in β as it did for the full model, corresponding to the presence of more spike-train groups that could be distinguished by firing rate. Nonetheless, the within-group structure to the dynamics has disappeared in the presence of dopamine. Groups in (b) and (c) shown for the threshold that maximised β (b, $\theta = 0.1$; c, $\theta = 0.2$); groups numbered top-down.

tifying meaningful groups from many hundreds or thousands of simultaneously recorded neurons. Such problems are at the cutting-edge of current neuroscience research. Recent multi-electrode experimental techniques (Buzsaki, 2004; Miller & Wilson, 2008) and calcium imaging (Carrillo-Reid et al., 2008) are already identifying hundreds of neurons, and there is a clear need for suitable multi-unit methods (Brown et al., 2004), as their absence places a fundamental limit on the use of such large-scale data.

Our graph-cut algorithm, based on a Hamming distance correlation metric, is an excellent candidate method for such exploratory analyses, as it does not require the user to pre-suppose the size or number of the neuron groups within the data. We have shown here that the algorithm can successfully find a wide range of groups that differ in number, size, discreteness of firing, and time-scale of activation. Extending this approach to determine neurons shared between more than one cell assembly (Carrillo-Reid et al., 2008), and to improve detection of groupings at multiple time-scales is the subject of ongoing work.

6.2. Microcircuit dynamics

We found no evidence for strongly synchronised neuron groups or for competitive dynamics between groups, under conditions we considered to best approximate an *in vivo* state with tonic dopamine. Of course, we cannot rule out the possibility that such competing neurons groups are on the scale of the size of this network, but then only a few such groups could be sustained in the striatum as a whole. Certainly, such a group would be much larger than that proposed in the "domain" theory (Wickens et al., 1991). Having prototyped all the appropriate methods here, we can now use larger scale networks and search for the existence of larger cell assemblies.

Without dopamine, noisy but synchronised groups of MSNs were observed. This is consistent with the increased synchrony of rat MSNs recorded under anaesthetic following dopamine depletion (Tseng et al., 2001). Moreover, it provides an explanation for the perplexing finding that dopamine-depletion causes a fragmentation of striatal cell clusters corresponding to sensorimotor stimulation of particular body parts (Cho et al., 2002). Our model clearly suggests that such fragmentation of striatal cells into smaller groups after dopamine loss occurs as a natural consequence of the microcircuit's dynamics.

Taken together, observations of the model with and without dopamine suggest that the dominance of asynchronous firing with dopamine is beneficial: then the response to subsequent structured input would not be corrupted by ongoing synchronised dynamics in the striatum, whereas part of the effect of dopamine-loss is to hinder striatal computations by having small groups that respond incorrectly to structured input.

6.3. Dissection of the network

The MSN-only, dopamine-free network showed the clearest set of cell groups that were consistent with some form of competitive dynamic within the striatal network. With low-levels of dopamine, the same network showed a structured sequence of silent periods across all groups. These results show why MSN-only, dopamine-free models could mislead theories on striatal function to focussing on competition within it.

The MSN-only network also provided evidence that the FSIs have a counter-intuitive facilitatory effect on MSN firing. This is consistent with experimental evidence for an excitatory effect of GABAergic input to MSNs (Bracci & Panzeri, 2006), but goes beyond this by showing a detectable affect on firing rate. In the model, we see that this is simply due to the disparity between the reversal potential for chloride and the nominal resting potential of the MSN. GABAergic input acts to drive the membrane potential toward this reversal potential; yet the MSN's isolated resting potential is around -80 to -90 mV, due to an inwardly-rectifying potassium current. From this starting point, GABAergic input to the MSN can depolarise the neuron, and could make it easier to fire, not less. However, the facilitatory effect must depend on a balance of glutamatergic and GABAergic input: an excess of GABAergic input would simply clamp the membrane potential to the reversal potential for

chloride. Our model shows that, beyond just a depolarising effect (Plenz, 2003), realistically parameterised GABAergic input to the MSNs can produce detectable changes in firing rate.

The MSN-collateral lesioned models revealed that the comparatively simple, noisy structure of the full model's spike train groups was seemingly dominated by the FSI input, supporting views that the sparse FSIs play a major role in the striatum (Tepper et al., 2004). They appeared to obscure or desynchronise the groups formed by the MSN-only network. This points to a different hypothesis for the "basic computational element" of the striatum: that the local MSN collaterals can support competitive dynamics within the striatum, but only when the damping influence of the GABAergic FSIs is removed. A possible mechanism for this is provided by the GABAergic control of FSIs by globus pallidus (Bevan et al., 1998). In this picture, a small region of striatum would support local competition following enhanced pallidal inhibition of a group of GABAergic interneurons. Firmer conclusions here require gathering of better data on striatal FSIs in particular.

The complex behaviour of the model FSI network emphasises the need for more data. We saw an unintuitive effect of gap junctions. Their presence caused the FSI population firing rate to fall with increasing dopamine, despite dopamine having an excitatory effect on the individual neuron. Removing the gap junctions reversed this trend. However, it also caused a large overall fall in population firing rate due to the homogenisation of the firing patterns. With gap junctions, we typically saw the firing rates spread over 100 spikes/s or more with peaks at both very low and high ends of this range; without them, the firing rates were within 3 spikes/s of each other. Presumably the FSI network transitions between multiple states between these two extremes, dependent on gap junction strength. A particularly interesting avenue for future work is the potential for this strength to be controlled dynamically by the nitric oxide releasing interneurons in the striatum (O'Donnell & Grace, 1997).

Further interesting extensions include the other omitted interneuron class (cholinergic) and the effects of GABAb receptors. The sustained activity of the FSIs could raise extrasynaptic GABA concentration sufficiently to activate extrasynaptic GABAb receptors by volume transmission. In the striatum these are predominantly located on axon terminals (Galvan et al., 2006). Their activation by baclofen attenuates both excitatory and inhibitory post-synaptic potentials in MSNs (Nisenbaum et al., 1993). This suggests they are located on both glutamatergic terminals, decoupling the MSNs from their input, and on GABAergic terminals, self-regulating inhibition throughout the striatum. A thorough study awaits the incorporation of three-dimensional volume transmission models into the striatal microcircuit studied here.

6.4. Conclusions

We have found some tantalising hints about the nature of the computations within the striatal microcircuit, and their relationship to dopamine-depleted disease states. The cell assemblies found here under a variety of conditions can now form the basis for a focussed study of the striatum's dynamics under structured input: its responses to slow-wave cortical activity under anaesthetic (Tseng et al., 2001), broad-scale synchrony during free behaviour (Berke et al., 2004), and its reorganisation during learning (Tang et al., 2007). More generally, in developing our model we have assembled an armoury of tools applicable to the study of neuronal microcircuits.

Acknowledgements. This work was funded by the EU Framework 6 Project IST-027819-IP, EPSRC Research Grant EP/C516303/1, and the EPSRC "CARMEN" e-Science Pilot Project. We thank Nathan Lepora for discussions and prior contributions to the single neuron models.

A. Specification for the spike-train clustering algorithm

We specify here the choices of correlation method and graph-modularity algorithm for use in our general algorithm for analysing multi-unit data (section 4).

A.1. Choice of correlation method: Hamming distance

There is a wide choice of correlation methods for spike train pairs, but few fulfill the multiple criteria we require: reducible to a scalar value, and suitable for repeated, exhaustive computation of all pairs at multiple time-scales. We propose here a novel (to the best of our knowledge) metric that fulfills these criteria. For each spike train, we divide time into bins of width δt, and in each bin record a 1 for the presence of any spikes, and a 0 for the absence. The resulting binary vector of length q thus records when the neuron is active or not. For each pair of binary vectors, we compute the normalised Hamming distance h, which is just the proportion of bins that differ between the two vectors: the smaller h, the closer the two vectors, and thus the more synchronised the two neurons at time-scale δt. Therefore we use $C_{ij} = C_{ji} = h$ to construct the correlation matrix.

This encoding and correlation method has two advantages for us. By not encoding the magnitude of activity we do not confound the co-activity of two neurons with scalar measurements of their respective magnitudes. More importantly, using the Hamming distance places equal weight on co-active and co-silence periods. This allows us to find neurons that are co-active, rather than the more limited set that are co-active at similar rates. Of course, in other applications of the general algorithm, a more detailed exploration of a smaller subset might be desirable: then the correlation matrix could be constructed by counting spikes in each bin and applying a continuous metric, such as the correlation coefficient.

A.2. Choice of grouping method: graph "modularity"

Standard data clustering techniques, such as k-means and its various extensions, can be used to find clusters within the matrix $\mathbf{A}^*$ or directly on the correlation matrix $\mathbf{C}$, but these are limited by the need for the user to define the number of clusters in advance. Similarly, standard graph partitioning algorithms applied to $\mathbf{A}^*$ require prior specification of the size of the resulting partitioned groups (Newman, 2006a). They are thus difficult to use for the kind of exploratory data analysis required here. We propose here a novel use of a new graph-partitioning method

that circumvents this problem by not requiring prior specification of the number of groups. We can provide only an outline of this method here – the reader is referred to (Newman, 2006a,b) for details.

The general problem of dividing networks into clusters has generated many techniques in network theory, generally under the label of "community detection". Central to many of them is an attempt to maximise a benefit function

$$\begin{aligned} Q = &[\text{number of edges in a community}] \\ &-[\text{expected number of such edges}], \end{aligned} \tag{20}$$

often called "modularity", over all possible subdivisions of the graph describing the network. That is, the division that maximises Q creates the clearest division of the graph into two sets, with more connections within them, but fewer connections between them, than expected. Key here is quantification of the "expected number of edges", which is encapsulated in the null model graph $\mathbf{P}$. We can define Q in matrix notation as

$$Q = \mathbf{s}^T \mathbf{B} \mathbf{s}, \tag{21}$$

where $\mathbf{s}$ is a vector denoting the group membership of each node (defined in (24)), and $\mathbf{B}$ is the "modularity" matrix whose entries

$$B_{ij} = A_{ij} - P_{ij} \tag{22}$$

denote the difference between the number of links A_{ij} connecting nodes i and j (allowing for graphs with multiple links between nodes), and the expected number of links P_{ij}. Following Newman (2006a), we use here the null model

$$P_{ij} = \frac{k_i k_j}{2m}, \tag{23}$$

where m is the number of unique links in the original graph (here m^*), and k_i, k_j are the degrees of (number of links made by) nodes i and j. This null model is closely related to the so-called "configuration" model: it essentially forms a random graph with the same expected degree sequence as the graph being analysed.

If we compute the eigenvalues β_i and the eigenvectors $\mathbf{u}_i$ of $\mathbf{B}$ for all n nodes, and order the eigenvalues so that $\beta_1 \geq \beta_2 \geq \ldots \geq \beta_n$, then it turns out that we can use the leading eigenvector $\mathbf{u}_1$ to partition the nodes into two groups:

$$s_i = \begin{cases} 1 & \text{if } u_i(1) \geq 0 \\ -1 & \text{if } u_i(1) < 0. \end{cases} \tag{24}$$

We can then use the resulting vector $\mathbf{s}$ to compute modularity Q from (21): if Q is positive, we retain the split; if it is negative, we do not split. We need only do this once, as the sum in (21) is maximised by the choice of (24).

To find more than two groups, we repeat the process on each subgraph defined by the nodes in a group. As each subgraph is smaller than the original graph, a correction is applied to (22) to account for the smaller number of nodes and edges

$$B_{ij}^{(g)} = A_{ij} - P_{ij} - \delta_{ij} \left[k_i^{(g)} - k_i \frac{d_g}{2m} \right], \tag{25}$$

where δ_{ij} is the Dirac delta functions, $k_i^{(g)}$ is the degree of the ith node in subgraph g, and d_g is the sum of all *total* degrees k_i (from the original graph) for the nodes in the subgraph. Then we replace $\mathbf{B} \leftarrow \mathbf{B}^{(g)}$ in (21), and repeat the computation of eigenvalues and eigenvectors. The subgraph is then split according to vector $\mathbf{s}$ if $Q > 0$ as before. This procedure is iterated for all subgraphs until no split yields $Q > 0$: then we have reached a state where no subgraph can increase its modularity by dividing further.

In our application of this method, the outcome is a set of node groups, corresponding to all found subgraphs. Each subgraph in turn corresponds to a set of neurons with sufficiently strong correlation between them to fulfill the definition of modularity (20).

References

Alexander, M. E., & Wickens, J. R. (1993). Analysis of striatal dynamics: the existence of two modes of behaviour. *J Theor Biol*, *163*, 413–438.

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

Bauswein, E., Fromm, C., & Preuss, A. (1989). Corticostriatal cells in comparison with pyramidal tract neurons: contrasting properties in the behaving monkey. *Brain Res*, *493*, 198–203.

Beiser, D. G., & Houk, J. C. (1998). Model of cortical-basal ganglionic processing: Encoding the serial order of sensory events. *J Neurophysiol*, *79*, 3168–3188.

Belluardo, N., Mudo, G., Trovato-Salinaro, A., Gurun, S. L., Charollais, A., Serre-Beinier, V., Amato, G., Haefliger, J. A., Meda, P., & Condorelli, D. F. (2000). Expression of connexin36 in the adult and developing rat brain. *Brain Res*, *865*, 121–138.

Berke, J. D. (2008). Uncoordinated firing rate changes of striatal fast-spiking interneurons during behavioral task performance. *J Neurosci*, *28*, 10075–10080.

Berke, J. D., Okatan, M., Skurski, J., & Eichenbaum, H. B. (2004). Oscillatory entrainment of striatal neurons in freely moving rats. *Neuron*, *43*, 883–896.

Bevan, M. D., Booth, P. A. C., Eaton, S. A., & Bolam, J. P. (1998). Selective innervation of neostriatal interneurons by a subclass of neuron in the globus pallidus of the rat. *J Neurosci*, *18*, 9438–9452.

Blackwell, K. T., Czubayko, U., & Plenz, D. (2003). Quantitative estimate of synaptic inputs to striatal neurons during up and down states in vitro. *J Neurosci*, *23*, 9123–9132.

Bolam, J. P., Bergman, H., Graybiel, A. M., Kimura, M., Plenz, D., Seung, H. S., Surmeier, D. J., & Wickens, J. R. (2006). Microcircuits in the striatum. In S. Grillner, & A. M. Graybiel (Eds.) *Microcircuits: the interface between neurons and global brain function*, (pp. 165–190). Cambridge, MA: MIT Press.

Bracci, E., Centonze, D., Bernardi, G., & Calabresi, P. (2002). Dopamine excites fast-spiking interneurons in the striatum. *J Neurophysiol*, *87*, 2190–2194.

Bracci, E., & Panzeri, S. (2006). Excitatory GABAergic effects in striatal projection neurons. *J Neurophysiol*, *95*, 1285–1290.

Brown, E. N., Kass, R. E., & Mitra, P. P. (2004). Multiple neural spike train data analysis: state-of-the-art and future challenges. *Nat Neurosci*, *7*, 456–461.

Burke, R. E., Marks, W. B., & Ulfhake, B. (1992). A parsimonious description of motoneuron dendritic morphology using computer simulation. *J Neurosci*, *12*, 2403–2416.

Buzsaki, G. (2004). Large-scale recording of neuronal ensembles. *Nat Neurosci*, *7*, 446–451.

Carrillo-Reid, L., Tecuapetla, F., Tapia, D., Hernandez-Cruz, A., Galarraga, E., Drucker-Colin, R., & Bargas, J. (2008). Encoding network states by striatal cell assemblies. *J Neurophysiol*, *99*, 1435–1450.

Centonze, D., Grande, C., Usiello, A., Gubellini, P., Erbs, E., Martin, A. B., Pisani, A., Tognazzi, N., Bernardi, G., Moratalla, R., Borrelli, E., & Calabresi, P. (2003). Receptor subtypes involved in the presynaptic and postsy-

naptic actions of dopamine on striatal interneurons. *J Neurosci*, *23*, 6245–6254.

Cho, J., Duke, D., Manzino, L., Sonsalla, P. K., & West, M. O. (2002). Dopamine depletion causes fragmented clustering of neurons in the sensorimotor striatum: evidence of lasting reorganization of corticostriatal input. *J Comp Neurol*, *452*, 24–37.

Czubayko, U., & Plenz, D. (2002). Fast synaptic transmission between striatal spiny projection neurons. *Proc Natl Acad Sci U S A*, *99*, 15764–15769.

Dawson, T. M., Gehlert, D. R., McCabe, R. T., Barnett, A., & Wamsley, J. K. (1986). D-1 dopamine receptors in the rat brain: a quantitative autoradiographic analysis. *J Neurosci*, *6*, 2352–2365.

Djurfeldt, M., Ekeberg, O., & Lansner, A. (2008). Large-scale modeling a tool for conquering the complexity of the brain. *Front Neuroinform*, *2*, 1.

Ferro, M. M., Bellissimo, M. I., Anselmo-Franci, J. A., Angellucci, M. E. M., Canteras, N. S., & Cunha, C. D. (2005). Comparison of bilaterally 6-OHDA- and MPTP-lesioned rats as models of the early phase of parkinson's disease: histological, neurochemical, motor and memory alterations. *J Neurosci Methods*, *148*, 78–87.

Frank, M. J. (2005). Dynamic dopamine modulation in the basal ganglia: a neurocomputational account of cognitive deficits in medicated and nonmedicated parkinsonism. *J Cogn Neurosci*, *17*, 51–72.

Fukai, T., & Tanaka, S. (1997). A simple neural network exhibiting selective activation of neuronal ensembles: from winner-take-all to winners-share-all. *Neural Comput*, *9*, 77–97.

Galan, R. F. (2008). On how network architecture determines the dominant patterns of spontaneous neural activity. *PloS One*, *3*, e2148.

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

Galvan, A., Kuwajima, M., & Smith, Y. (2006). Glutamate and GABA receptors and transporters in the basal ganglia: What does their subsynaptic localization reveal about their function? *Neuroscience*, *143*, 351–375.

Gerfen, C., & Wilson, C. (1996). The basal ganglia. In L. Swanson, A. Bjorklund, & T. Hokfelt (Eds.) *Handbook of chemical neuroanatomy. Vol 12: Integrated systems of the CNS. Part III*, (pp. 371–468). Amsterdam: Elsevier.

Glynn, G., & Ahmad, S. O. (2002). Three-dimensional electrophysiological topography of the rat corticostriatal system. *J Comp Physiol A*, *188*, 695–703.

Gonon, F. (1997). Prolonged and extrasynaptic excitatory action of dopamine mediated by d1 receptors in the rat striatum in vivo. *Journal of Neuroscience*, *17*, 5972–5978.

Gorelova, N., Seamans, J. K., & Yang, C. R. (2002). Mechanisms of dopamine activation of fast-spiking interneurons that exert inhibition in rat prefrontal cortex. *J Neurophysiol*, *88*, 3150–3166.

Groenewegen, H. J., Wright, C. I., Beijer, A. V., & Voorn, P. (1999). Convergence and segregation of ventral striatal inputs and outputs. *Ann N Y Acad Sci*, *877*, 49–63.

Gurney, K., Prescott, T. J., & Redgrave, P. (2001). A computational model of action selection in the basal ganglia II: Analysis and simulation of behaviour. *Biol Cybern*, *85*, 411–423.

Humphries, M. D., & Gurney, K. (2007). Solution methods for a new class of simple model neurons. *Neural Comp*, *19*, 3216–3225.

Humphries, M. D., & Gurney, K. (2008). Network 'small-world-ness': A quantitative method for determining canonical network equivalence. *PLoS One*, *3*, e0002051.

Humphries, M. D., Stewart, R. D., & Gurney, K. N. (2006). A physiologically plausible model of action selection and oscillatory activity in the basal ganglia. *J Neurosci*, *26*, 12921–12942.

Hurd, Y. L., Suzuki, M., & Sedvall, G. C. (2001). D1 and D2 dopamine receptor mRNA expression in whole hemisphere sections of the human brain. *J Chem Neuroanat*, *22*, 127–137.

Ingham, C. A., Hood, S. H., Taggart, P., & Arbuthnott, G. W. (1998). Plasticity of synapses in the rat neostriatum after unilateral lesion of the nigrostriatal dopaminergic pathway. *J Neurosci*, *18*, 4732–4743.

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

Izhikevich, E. M., Gally, J. A., & Edelman, G. M. (2004). Spike-timing dynamics of neuronal groups. *Cereb Cortex*, *14*, 933–44.

Jaeger, D., Hitoshi, K., & Wilson, C. J. (1994). Surround inhibition among projection neurons is weak or nonexistent in the rat neostriatum. *J Neurophysiol*, *72*, 2555–2558.

Jahr, C. E., & Stevens, C. F. (1990). A quantitative description of NMDA receptor-channel kinetic behavior. *J Neurosci*, *10*, 1830–1837.

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

Kawaguchi, Y., Wilson, C. J., Augood, S. J., & Emson, P. C. (1995). Striatal interneurones: Chemical, physiological and morphological characterization. *Trends Neurosci*, *18*, 527–535.

Kirik, D., Rosenblad, C., & Bjorklund, A. (1998). Characterization of behavioral and neurodegenerative changes following partial lesions of the nigrostriatal dopamine system induced by intrastriatal 6-hydroxydopamine in the rat. *Exp Neurol*, *152*, 259–277.

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

Koos, T., & Tepper, J. M. (1999). Inhibitory control of neostriatal projection neurons by GABAergic interneurons. *Nature Neurosci*, *2*, 467–472.

Koos, T., Tepper, J. M., & Wilson, C. J. (2004). Comparison of IPSCs evoked by spiny and fast-spiking neurons in the neostriatum. *J Neurosci*, *24*, 7916–7922.

Kubota, Y., & Kawaguchi, Y. (2000). Dependence of GABAergic synaptic areas on the interneuron type and target size. *J Neurosci*, *20*, 375–386.

Kwok, H. F., Jurica, P., Raffone, A., & van Leeuwen, C. (2007). Robust emergence of small-world structure in networks of spiking neurons. *Cogn Neurodyn*, *1*, 39–51.

Lago-Fernandez, L. F., Huerta, R., Corbacho, F., & Siguenza, J. A. (2000). Fast response and temporal coherent oscillations in small-world networks. *Phys Rev Lett*, *84*, 2758–2761.

Mallet, N., Le Moine, C., Charpier, S., & Gonon, F. (2005). Feedforward inhibition of projection neurons by fast-spiking GABA interneurons in the rat striatum in vivo. *J Neurosci*, *25*, 3857–3869.

Markram, H. (2006). The blue brain project. *Nat Rev Neurosci*, *7*, 153–160.

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

Miller, E. K., & Wilson, M. A. (2008). All my circuits: using multiple electrodes to understand functioning neural networks. *Neuron*, *60*, 483–488.

Moyer, J. T., Wolf, J. A., & Finkel, L. H. (2007). Effects of dopaminergic modulation on the integrative properties of the ventral striatal medium spiny neuron. *J Neurophysiol*, *98*, 3731–3748.

Newman, M. E. J. (2006a). Finding community structure in networks using the eigenvectors of matrices. *Phys Rev E*, *74*, 036104.

Newman, M. E. J. (2006b). Modularity and community structure in networks. *Proc Natl Acad Sci U S A*, *103*, 8577–8582.

Nisenbaum, E. S., Berger, T. W., & Grace, A. A. (1993). Depression of glutamatergic and GABAergic synaptic responses in striatal spiny neurons by stimulation of presynaptic GABAB receptors. *Synapse*, *14*, 221–242.

Nishikawa, T., Motter, A. E., Lai, Y.-C., & Hoppensteadt, F. C. (2003). Heterogeneity in oscillator networks: are smaller worlds easier to synchronize? *Phys Rev Lett*, *91*, 014101.

O'Donnell, P., & Grace, A. A. (1997). Cortical afferents modulate striatal gap junction permeability via nitric oxide. *Neuroscience*, *76*, 1–5.

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. *J Comp Neurol*, *366*, 580–99.

Pennartz, C. M., Groenewegen, H. J., & da Silva, F. H. L. (1994). The nucleus accumbens as a complex of functionally distinct neuronal ensembles: an integration of behavioural, electrophysiological and anatomical data. *Prog Neurobiol*, *42*, 719–761.

Plenz, D. (2003). When inhibition goes incognito: feedback interaction between spiny projection neurons in striatal function. *Trends Neurosci*, *26*, 436–443.

Redgrave, P., & Gurney, K. (2006). The short-latency dopamine signal: a role in discovering novel actions? *Nat Rev Neurosci*, *7*, 967–975.

Richtand, N. M., Kelsoe, J. R., Segal, D. S., & Kuczenski, R. (1995). Regional quantification of D1, D2, and D3 dopamine receptor mRNA in rat brain using a ribonuclease protection assay. *Brain Res Mol Brain Res*, *33*, 97–103.

Sandstrom, M. I., & Rebec, G. V. (2003). Characterization of striatal activity in conscious rats: contribution of NMDA and AMPA/kainate receptors to both spontaneous and glutamate-driven firing. *Synapse*, *47*, 91–100.

Schultz, W. (2007). Behavioral dopamine signals. *Trends Neurosci*, *30*, 203–

210.

Schwarting, R. K., & Huston, J. P. (1996). The unilateral 6-hydroxydopamine lesion model in behavioral brain research. analysis of functional deficits, recovery and treatments. *Prog Neurobiol*, *50*, 275–331.

Smith, Y., Raju, D. V., Pare, J. F., & Sidibe, M. (2004). The thalamostriatal system: a highly specific network of the basal ganglia circuitry. *Trends Neurosci*, *27*, 520–527.

Sullivan, M. A., Chen, H., & Morikawa, H. (2008). Recurrent inhibitory network among striatal cholinergic interneurons. *J Neurosci*, *28*, 8682–8690.

Surmeier, D. J., Ding, J., Day, M., Wang, Z., & Shen, W. (2007). D1 and D2 dopamine-receptor modulation of striatal glutamatergic signaling in striatal medium spiny neurons. *Trends Neurosci*, *30*, 228–235.

Tang, C., Pawlak, A. P., Prokopenko, V., & West, M. O. (2007). Changes in activity of the striatum during formation of a motor habit. *Eur J Neurosci*, *25*, 1212–1227.

Tateno, T., Harsch, A., & Robinson, H. P. (2004). Threshold firing frequency-current relationships of neurons in rat somatosensory cortex: type 1 and type 2 dynamics. *J Neurophysiol*, *92*, 2283–2294.

Taverna, S., van Dongen, Y. C., Groenewegen, H. J., & Pennartz, C. M. (2004). Direct physiological evidence for synaptic connectivity between medium-sized spiny neurons in rat nucleus accumbens in situ. *J Neurophysiol*, *91*, 1111–1121.

Tepper, J. M., Koos, T., & Wilson, C. J. (2004). GABAergic microcircuits in the neostriatum. *Trends Neurosci*, *27*, 662–669.

Traub, R. D., Kopell, N., Bibbig, A., Buhl, E. H., LeBeau, F. E., & Whittington, M. A. (2001). Gap junctions between interneuron dendrites can enhance synchrony of gamma oscillations in distributed networks. *J Neurosci*, *21*, 9478–9486.

Tseng, K. Y., Kasanetz, F., Kargieman, L., Riquelme, L. A., & Murer, M. G. (2001). Cortical slow oscillatory activity is reflected in the membrane potential and spike trains of striatal neurons in rats with chronic nigrostriatal lesions. *Journal of Neuroscience*, *21*, 6430–6439.

Tunstall, M. J., Oorschot, D. E., Kean, A., & Wickens, J. R. (2002). Inhibitory interactions between spiny projection neurons in the rat striatum. *J Neurophysiol*, *88*, 1263–1269.

Turner, R. S., & DeLong, M. R. (2000). Corticostriatal activity in primary motor cortex of the macaque. *J Neurosci*, *20*, 7096–7108.

Venton, B. J., Zhang, H., Garris, P. A., Phillips, P. E., Sulzer, D., & Wightman, R. M. (2003). Real-time decoding of dopamine concentration changes in the caudate-putamen during tonic and phasic firing. *J Neurochem*, *87*, 1284–1295.

Wang, Z., Kai, L., Day, M., Ronesi, J., Yin, H. H., Ding, J., Tkatch, T., Lovinger, D. M., & Surmeier, D. J. (2006). Dopaminergic control of corticostriatal long-term synaptic depression in medium spiny neurons is mediated by cholinergic interneurons. *Neuron*, (pp. 43–452).

Watts, D. J., & Strogatz, S. H. (1998). Collective dynamics of 'small-world' networks. *Nature*, *393*, 440–442.

Whishaw, I. Q., & Dunnett, S. B. (1985). Dopamine depletion, stimulation or blockade in the rat disrupts spatial navigation and locomotion dependent upon beacon or distal cues. *Behav Brain Res*, *18*, 11–29.

Wickens, J. R., Alexander, M. E., & Miller, R. (1991). Two dynamic modes of striatal function under dopaminergic-cholinergic control: simulation and analysis of a model. *Synapse*, *8*, 1–12.

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

Wilson, C. J., & Groves, P. M. (1980). Fine structure and synaptic connections of the common spiny neuron of the rat neostriatum: a study employing intracellular inject of horseradish peroxidase. *J Comp Neurol*, *194*, 599–615.

Wood, R., Humphries, M., & Gurney, K. (2007). A large scale biologically realistic model of the neostriatum. In *IBAGS IX*. IBAGS IX, Egmond aan Zee, The Netherlands.

Yin, H. H., & Knowlton, B. J. (2006). The role of the basal ganglia in habit formation. *Nat Rev Neurosci*, *7*, 464–476.

Zheng, T., & Wilson, C. J. (2002). Corticostriatal combinatorics: the implications of corticostriatal axonal arborizations. *J Neurophysiol*, *87*, 1007–1017.

Zhou, F.-M., Wilson, C. J., & Dani, J. A. (2002). Cholinergic interneuron characteristics and nicotinic properties in the striatum. *J Neurobiol*, *53*, 590–605.