

The Basal Ganglia and Cortex Implement Optimal Decision Making Between Alternative Actions

Rafal Bogacz

R.Bogacz@bristol.ac.uk

Department of Computer Science, University of Bristol, Bristol BS8 1UB, U.K.

Kevin Gurney

K.Gurney@shef.ac.uk

Department of Psychology, University of Sheffield, Sheffield S10 2TP, U.K.

Neurophysiological studies have identified a number of brain regions critically involved in solving the problem of action selection or decision making. In the case of highly practiced tasks, these regions include cortical areas hypothesized to integrate evidence supporting alternative actions and the basal ganglia, hypothesized to act as a central switch in gating behavioral requests. However, despite our relatively detailed knowledge of basal ganglia biology and its connectivity with the cortex and numerical simulation studies demonstrating selective function, no formal theoretical framework exists that supplies an algorithmic description of these circuits. This article shows how many aspects of the anatomy and physiology of the circuit involving the cortex and basal ganglia are exactly those required to implement the computation defined by an asymptotically optimal statistical test for decision making: the multihypothesis sequential probability ratio test (MSPRT). The resulting model of basal ganglia provides a new framework for understanding the computation in the basal ganglia during decision making in highly practiced tasks. The predictions of the theory concerning the properties of particular neuronal populations are validated in existing experimental data. Further, we show that this neurobiologically grounded implementation of MSPRT outperforms other candidates for neural decision making, that it is structurally and parametrically robust, and that it can accommodate cortical mechanisms for decision making in a way that complements those in basal ganglia.

1 Introduction ---

Recent experimental results have established that both the cortex and the basal ganglia are involved in decision making between alternative actions (Chevalier, Vacher, Deniau, & Desban, 1985; Deniau & Chevalier, 1985; Medina & Reiner, 1995; Redgrave, Prescott, & Gurney, 1999; Schall, 2001; Shadlen

& Newsome, 2001; Smith, Bevan, Shink, & Bolam, 1998). However, it is necessary to distinguish between two phases of developing task competence (Ashby & Spiering, 2004; Shadmehr & Holcomb, 1997). In the acquisition or learning phase, actions appropriate in a given behavioral state (i.e., a combination of ongoing behavior and stimulus) are being developed that are usually driven by external reward. In this phase, the main difficulty lies in finding the optimal policy: the mapping between states and actions that maximizes reward (Sutton & Barto, 1998). Consistent with this requirement, there is a great deal of evidence that the basal ganglia act as a substrate for reinforcement learning (O'Doherty et al., 2004; Samejima, Ueda, Doya, & Kimura, 2005; Schultz, Dayan, & Montague, 1997), and several models for this process have been proposed (Doya, 2000; Frank, Seeberger, & O'Reilly, 2004; Montague, Dayan, & Sejnowski, 1996).

In contrast, in the proficient phase, the mapping between stimulus and an appropriate response is well established, and the requirement is one of performing action selection or decision making. This consists of identifying the current behavioral state and executing the known appropriate action as soon as a certain level of confidence in the identification is reached (Gold & Shadlen, 2001, 2002). There is a great deal of evidence (reviewed below) for the involvement of the basal ganglia in action selection.

We view the hypotheses that basal ganglia perform reinforcement learning and that they perform action selection as complementary. Thus, at any particular time, the basal ganglia perform action selection proficiently with respect to a suite of alternatives that have already been learned. Subsequent learning phases will modify that suite of alternatives and shape the profile of selections that can be made. In this article, we focus on the neural mechanisms underlying decision making in the proficient phase. As such, this is the phase under discussion if no qualifier is specified. We return to the relation of our work to task acquisition and the learning phase in section 6.

1.1 Action Selection and Decision Making. Experimental data show that during the decision process in visual discrimination tasks, neurons in cortical areas representing alternative actions gradually increase their firing rate, thereby accumulating evidence supporting these alternatives (Schall, 2001; Shadlen & Newsome, 2001). Hence, the models of decision making based on neurophysiological data (Shadlen & Newsome, 2001; Wang, 2002) assume that there exist connections from neurons representing stimuli to the appropriate cortical neurons representing actions (these connections may develop during the months of training the animals undergo before these experiments). These cortical connections are assumed to encode the stimulus-response mapping.

However, even in simple, highly constrained laboratory tasks, there will be more than one possible response, and so there is a problem of action selection in which the representation for the correct response has to take control of the animal's motor plant. In natural, ethological settings, this problem is

exacerbated because, under these circumstances, there are usually multiple, complex sensory streams demanding a variety of behaviors.

The problem of action selection was recently addressed by Redgrave et al. (1999), who conceived of it as the resolution of conflict between command centers throughout the brain competing for behavioral expression. These authors examined the problem from a computational perspective and argued that competitions between brain centers vying for expression were best resolved by a central switch examining the urgency or salience of each action request, and that anatomical and physiological evidence pointed to the basal ganglia as the neural substrate for this switch. Thus, the basal ganglia receive widespread input from all over the brain (Parent & Hazrati, 1995a) and, in their quiescent state, send tonic inhibition to midbrain and brain stem targets implicated in executing motor actions, thus blocking cortical control over these actions (Chevalier et al., 1985; Deniau & Chevalier, 1985). Actions are supposed to be selected when neurons in the output nuclei have their activity reduced (under control of the rest of the basal ganglia), thereby disinhibiting their targets (Chevalier et al., 1985; Deniau & Chevalier, 1985).

The selection hypothesis for basal ganglia has been tested in biologically realistic computational models in a variety of anatomical contexts by several authors (Brown, Bullock, & Grossberg, 2004; Frank, 2005; Gurney, Prescott, & Redgrave, 2001a, 2001b; Humphries & Gurney, 2002).

In sum, the research reviewed above indicates that during decision making among alternative actions, cortical regions associated with the alternatives integrate evidence supporting each one and that the basal ganglia act as a central switch by evaluating this evidence and enabling those behavioral requests that are best supported (most salient).

1.2 Scope of the Letter. We have argued that between bouts of learning—that is, during proficient phases of activity—the primary computational role of the basal ganglia is to act as an action selection mechanism, mediating resolution of the action selection problem by gating behavioral requests. It is this mode of basal ganglia operation that we address in this article.

Biologically realistic network models (Gurney et al., 2001a, 2001b) have shown how the basal ganglia could perform the required selection computation. However, such models fail to elucidate possible analytic descriptions of the computation (i.e., selection) that provide it with a theoretical grounding. This letter provides an analytic description of function of a circuit involving cortex and basal ganglia by showing how an optimal abstract decision algorithm maps onto the anatomy and physiology of this circuit.

The main goal of this letter is to provide a new algorithmic framework for understanding the computation in the basal ganglia in the proficient phase. The algorithm relates to computations being performed at the systems level of description of the basal ganglia, that is, considering the circuit to be a

set of interacting neuronal populations described by their overall firing rate (Dayan, 2001). This does not preclude the possibility that computations may be performed at other levels of description (Gurney, Prescott, Wickens, & Redgrave, 2004) dealing with microcircuits, membranes, or molecular signalling pathways. However, as far as computation performed by virtue of the organization of the basal ganglia *in toto* is concerned, we argue that this will have an integrity apparent at the systems level, while acknowledging that it must ultimately be consistent with lower-level models. In view of these methodological considerations, we therefore deliberately do not attempt to incorporate the overwhelming amount of knowledge available for the basal ganglia at the microanatomical, physiological, and molecular levels of description.

The letter is organized as follows. Section 2 reviews relevant background material concerning the neurobiology and theory of decision making in cortex and basal ganglia. Section 3 deals with the central technical argument and proposes how the basal ganglia may perform action selection in an optimal way. Section 4 shows how the specific experimental predictions of the theory are verified by existing data. Section 5 compares the performance of the proposed model against other models of decision making. Section 6 discusses the relation of this work to other theories of action selection and published experimental data.

2 Review of the Neurobiology and Theory of Decision Making _____

This section reviews the material critical for an understanding of our model. The theory of optimal decision making in perceptual tasks has hitherto been grounded almost exclusively in cortical mechanisms. In contrast, we develop the theory of decision making in this article with respect to the neural circuit involving both the cortex and the basal ganglia. Thus, in proposing a neural mechanism for optimal decision making, we link two strands of research: that dealing with putative cortical decision mechanisms and that dealing with action selection in the basal ganglia. Elements from both areas are therefore required background material. First, we present the theory of optimal decision making, and then, we review those aspects of basal ganglia anatomy and physiology critical for the model.

2.1 Decision Making and Cortical Integration. The neural basis of decision making in cortex has been studied extensively using single-cell recordings (Britten, Shadlen, Newsome, & Movshon, 1993; Kim & Shadlen, 1999; Schall, 2001). Typically these studies have used a direction of motion discrimination task using fields of drifting random dots, with response via saccadic eye movements. During these experiments, the mapping between dot movement direction and required response was kept constant for many weeks of the training, so that these studies describe the proficient phase of task acquisition. After stimulus onset, neurons in cortical sensory areas (e.g.,

area MT in the visual motion task) respond if their receptive fields encounter the stimulus and are appropriately tuned to the overall direction of motion (Britten et al., 1993; Kim & Shadlen, 1999). However, the instantaneous firing rates in MT are noisy, probably reflecting the uncertainty inherent in the stimulus and its neural representation. Further, this noise is such that decisions based on the activity of MT neurons at a given moment in time would be inaccurate, because the largest firing rate does not always indicate the direction of coherent motion in the stimulus. Therefore, a statistical interpretation is required. An often used hypothesis (Gold & Shadlen, 2001, 2002) is that populations of neurons in MT encode evidence for a particular perceptual decision.

To formalize this, denote the evidence supporting decision i , (i is “left” or “right”) provided at time t , by $x_i(t)$. Then, under the neural encoding hypothesis, $x_i(t)$ corresponds to the total activity of MT neurons selective for direction i at time t . The decision-making process can be defined as one of finding which $x_i(t)$ has the highest mean (Gold & Shadlen, 2001, 2002). To solve it, it appears that subsequent cortical areas are invoked to accumulate evidence over time. Thus, in the motion discrimination task, neurons in the lateral intraparietal area (LIP) and frontal eye field (FEF) (which are implicated in the response via saccadic eye movements) gradually increase their firing rate (Schall, 2001; Shadlen & Newsome, 2001) and could therefore be computing

$$Y_i(T) = \sum_{t=1}^T x_i(t) \quad (2.1)$$

over the temporal interval $[1, T]$ (where we assume for simplicity a discrete representation of time). The accumulated evidence $Y_i(T)$ may now be used in making a decision about which $x_i(t)$ has the highest mean.

2.2 Modeling the Decision Criterion. The above description of cortical integration leaves open a central question: When should a neural mechanism stop the integration and execute the action with the highest accumulated evidence $Y_i(T)$? A simple solution to this problem is to execute an action as soon as any $Y_i(T)$ exceeds a certain decision threshold, yielding the so-called race model (Vickers, 1970). However, this model does not perform optimally. For example, in case of decision between two alternatives, it is more efficient to compute the difference between the accumulated evidence supporting the two alternatives and execute action as soon as this difference crosses a positive or a negative decision threshold. This procedure is known as a random walk (Laming, 1968; Stone, 1960) or a diffusion (Ratcliff, 1978) model, and it may be shown to implement a statistical decision test known as the sequential probability ratio test (SPRT) (Barnard, 1946; Wald, 1947). The SPRT is optimal in the following sense: among all decision methods

allowing a certain probability of error, it requires the shortest period of sampling the x_i —that is, it minimizes decision time (Wald & Wolfowitz, 1948).

2.3 The MSPRT. For more than two alternatives, there is no single optimal test in the sense that SPRT is optimal for two alternatives, but there are tests that are asymptotically optimal; that is, they minimize decision time for a fixed probability of error when this probability decreases to zero (Dragalin, Tertakovsky, & Veeravalli, 1999). These tests are the so-called multihypothesis SPRTs (MSPRT's) (Baum & Veeravalli, 1994; Dragalin et al., 1999), and for two alternatives, they simplify to the SPRT. While it has been shown that MSPRT may be performed in a two-layer connectionist network (McMillen & Holmes, 2006), the required complexities in this model mitigate against any obvious implementation in the brain (and, in particular, the cortex).

We now introduce the MSPRT (Baum & Veeravalli, 1994). A decision among N alternative actions can be formulated in the same way as for the case of the two alternatives described in section 2.1. That is, it amounts to finding which $x_i(t)$ has the highest mean (Gold & Shadlen, 2001, 2002). Let us define a set of N hypotheses H_i such that roughly speaking, each H_i corresponds to $x_i(t)$ having the highest mean. More precisely, we define H_i analogous to its definition for two alternatives (Gold & Shadlen, 2001, 2002); H_i is the hypothesis that $x_i(t)$ come from independent and identically distributed (i.i.d.) normal distributions with mean μ^+ and standard deviation σ , while $x_{j \neq i}(t)$ come from i.i.d. normal distributions with mean μ^- and standard deviation σ , where $\mu^+ > \mu^-$.

Bearing in mind that we are integrating evidence up until some time T , denote the entirety of sensory evidence available up to T by $input(T) = \{x_i(t) : 1 \leq i \leq N, 1 \leq t \leq T\}$. The MSPRT (Baum & Veeravalli, 1994) is equivalent to the following decision criterion at time T : for each alternative i , compute the conditional probability of hypothesis H_i given sensory inputs so far, $P_i(T) = P(H_i | input(T))$, and execute an action as soon as any of the $P_i(T)$ exceeds a certain decision threshold. Hence, sensory information is gathered until the estimated probability of one of the inputs having the highest mean exceeds the decision threshold.

Appendix A describes how $P_i(T)$ can be computed on the basis of sensory evidence. In particular, it shows that the logarithm of $P_i(T)$, which we denote by $L_i(T)$, is given by

$$L_i(T) = y_i(T) - \ln \sum_{k=1}^N \exp(y_k(T)) \quad (2.2)$$

where $y_i(T)$ is proportional to the accumulated evidence supporting action i . In particular, $y_i(T) = g^* Y_i(T)$, where $Y_i(T)$ is the accumulated evidence

supporting action i (see section 2.1) and g^* is a constant. We will refer to $y_i(T)$ as the salience of action i .

Thus, MSPRT implies that an action should be selected as soon as any of the $L_i(T)$ exceeds a fixed decision threshold. Equation 2.2 is the basis for mapping MSPRT onto the basal ganglia. However, before we proceed with this process, we describe some of the intuitive properties of MSPRT, as implemented in equation 2.2.

The right-hand side of equation 2.2 has two terms. The first term $y_i(T)$ is simply the salience and, on its own, describes a race model (Vickers, 1970). This term therefore incorporates information about the absolute size of the salience of the currently “winning” alternative. The second term in equation 2.2 occurs in subsequent analysis throughout the article, and we denote it by $S(T)$ where,

$$S(T) = \ln \sum_{k=1}^N \exp(y_k(T)). \quad (2.3)$$

The term $S(T)$ includes summation over all alternatives and does not depend on i . $S(T)$ therefore decreases the value of all $L_i(T)$ by the same amount, thereby increasing the minimum salience required for an action to be selected. It may therefore be thought of as representing response conflict, because its value is increased by more actions having high salience. In this way, $S(T)$ allows incorporation of information about the difference between the salience of the currently winning alternative and its competitors. The degree of scaling of the salience required for action selection implied by the particular form of $S(T)$ is critical for optimal decision making; it allows a much lower average decision time for fixed accuracy than when the scaling is not present (i.e., race model), as will be shown in section 5.

2.4 Basal Ganglia Connectivity. The basal ganglia connectivity used in our study contain the major pathways known to exist in basal ganglia anatomy and was based on that used in the model of Gurney et al. (2001a). Figure 1A shows this connectivity for rat in cartoon form (for reviews of basal ganglia anatomy, see Gerfen & Wilson, 1996; Mink, 1996; Smith et al., 1998). Cortex sends excitatory projections to the striatum (Nakano, Kayahara, Tsutsumi, & Ushiro, 2000) and subthalamic nucleus (STN) (Smith et al., 1998). The striatum is the largest basal ganglia nucleus and is divided into two populations of projection neurons differentiated, *inter alia*, by their anatomical targets and preferential dopamine receptor type (Gerfen & Young, 1988). The neurons in one striatal subpopulation send focused inhibitory projections to the basal ganglia output nuclei: the substantia nigra pars reticulata (SNr) and entopeduncular nucleus (EP) (the homologue of primate globus pallidus internal segment (GPi)). These striatal neurons are associated with D1-type dopamine receptors (Smith et al., 1998) and,

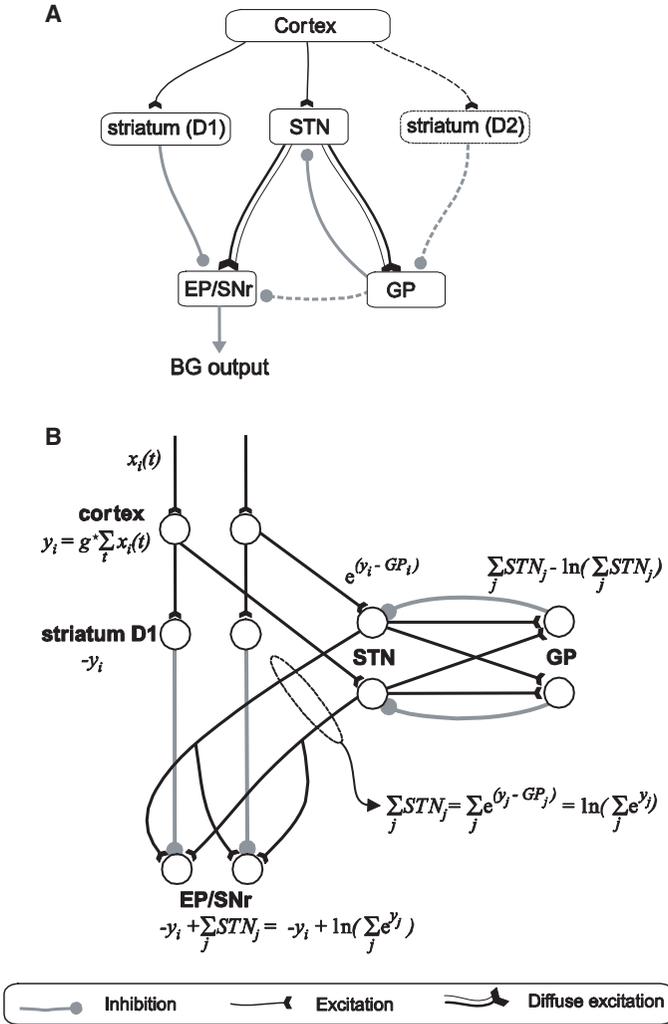


Figure 1: Comparison of connectivity of basal ganglia and a network implementing the multihypothesis sequential probability ratio Test (MSPRT). (A) Connectivity of basal ganglia nuclei and its cortical afferents in the rat (modified from Gurney et al., 2001a). Connections and nuclei denoted by dashed lines are not essential for the implementation of MSPRT. (B) Architecture of the network implementing MSPRT. The equations show expressions calculated by each layer of neurons.

together with their targets (SNr, EP), constitute the so-called “direct pathway” (in section 3 we associate this direct pathway with first term $y_i(T)$ in equation 2.2). Neurons in the other striatal population are also inhibitory, send focused projections to the globus pallidus (GP) (globus pallidus external segment or GPe in primate), and are associated with D2-type dopamine receptors (Smith et al., 1998). Neurons in the STN are glutamatergic and send diffuse excitatory projections to SNr/EP and GP (Parent & Hazrati, 1993, 1995a). The GP sends inhibitory connections to the output nuclei (Bevan, Smith, & Bolam, 1996). This second striatal population therefore gives rise to an indirect pathway to the output nuclei via GP and STN (in section 3, we propose that the nuclei traversed by the indirect pathway are involved in computation of term $S(T)$). The output nuclei send widespread inhibitory connections to the midbrain, brain stem (Faull & Mehler, 1978; Kha et al., 2001), and thalamus (Alexander, DeLong, & Strick, 1986).

2.5 Neuronal Selectivity in the Basal Ganglia. There is much evidence (reviewed below) pointing to a topographic representation of functionality within basal ganglia and its associated thalamocortical circuitry, leading to the hypothesis that these circuits support a range of discrete channels associated with different actions. This concept will be important for the model; in section 3, we associate each channel with a term $L_i(T)$.

At the largest scale of organization, Alexander et al. (1986) divided the loops from cortex, through basal ganglia, thalamus, and back to cortex, into five parallel, segregated circuits (cf. Nakano et al., 2000) associated with different functionality. Since this article treats the problem of action selection, we focus on motor and oculomotor circuits and return to consider information processing in the limbic and two prefrontal circuits in section 6.

Within the motor circuit, studies of awake primates in behavioral tasks have established that all basal ganglia nuclei have somatotopic organization. Thus, Crutcher and DeLong (1984a, 1984b) have shown that neurons selective for arm, leg, and face are located within different parts of the striatum. Further, within each body part, there are clusters of neurons responding selectively before and during movement of individual joints (often only in single direction) (Crutcher & DeLong, 1984a, 1984b). Similarly, Georgopoulos, DeLong, and Crutcher (1983) found neurons in other basal ganglia nuclei (STN, GP, EP) that were selective to the direction and speed of individual movements. These observations led Alexander et al. (1986) to propose that “the motor circuit may be composed of multiple, parallel sub-circuits or channels concerned with movement of individual body parts,” which traverse all nuclei of basal ganglia.

The notion of channels was incorporated into the computational model of Gurney et al. (2001a), who proposed that each action is associated anatomically with a discrete neural population within each nucleus. Channels are therefore defined at the input nuclei (striatum and STN) as populations

innervated by the cortical afferents associated with each action. Channel populations in other nuclei (GP, EP, SNr) are then defined by focused projections from corresponding striatal populations.

3 The Basal Ganglia Implements Selection Using MSPRT

We now show how the MSPRT test defined by equation 2.2 may be performed in a biologically constrained network model of the basal ganglia. For simplicity of explanation, we first show how equation 2.2 maps onto a model of basal ganglia including only a subset of the known anatomical connections (we exclude the connections marked by dotted lines in Figure 1A). Subsequently we demonstrate the mapping onto the model with the complete set of connectivity. The mapping between equation 2.2 and the network is shown graphically in Figure 1B. In our decomposition, each channel (see section 2.5) is associated with an action i and a term $L_i(T)$ in the MSPRT. Hence, we assume that there is a finite number N of available actions represented in a discrete (or “localist”) fashion (the topic of action representations is dealt with further in section 6).

We note first that the $L_i(T)$ are always negative or equal to 0, because $L_i(T) = \ln P_i(T)$, and $P_i(T)$ are probabilities. Thus, by definition, $P_i(T) \leq 1$, and so $\ln P_i(T) \leq \ln 1 = 0$. Therefore, the $L_i(T)$ themselves cannot be represented as firing rates in neuronal populations (since neurons cannot have negative firing rates). This may be overcome by assigning the network output OUT_i to $-L_i(T)$, that is,

$$OUT_i(T) = -y_i(T) + \ln \sum_{k=1}^N \exp(y_k(T)). \quad (3.1)$$

The decision is now made whenever any output decreases its activity below the threshold. Notice that this is consonant with the supposed action of basal ganglia outputs in performing selection by disinhibition of target structures (Chevalier et al., 1985; Deniau & Chevalier, 1985).

As described in section 1, we propose, along with others (Schall, 2001; Shadlen & Newsome, 2001), that quantities like $y_i(T)$, representing salience, are computed in cortical regions that project to basal ganglia. In the motion discrimination example (described in section 2.1), $y_i(T)$ would be computed in FEF, which is known to innervate the basal ganglia (Parthasarathy, Schall, & Graybiel, 1992). Since $y_i(T)$ is the product of the raw accumulated evidence $Y_i(T)$ and a scaling factor, g^* , we interpret g^* as the gain that cortex introduces in computing the salience (Brown et al., 2005). As shown in appendix A, the MSPRT algorithm specifies g^* exactly. Thus, there is an

optimal gain:

$$g^* = \frac{\mu^+ - \mu^-}{\sigma^2}, \quad (3.2)$$

where μ^+ , μ^- , σ parameterize the cortical inputs (see section 2.3). We return to the question of parametric robustness with respect to gain later.

Equation 3.1, describing the activity of the basal ganglia output nuclei, includes two terms, the first of which we propose is computed within the direct pathway, while the second term is within the pathway traversing STN and GP. The first term in equation 3.1, $-y_i(T)$, is an inhibitory component and cannot be supplied by cortex since its efferents are glutamatergic. We argue therefore that one function of the population of GABAergic striatal projection neurons with D1 receptors (see Figure 1A) is to provide an inhibitory copy of the salience signal to the output nuclei.

Turning to the second term in equation 3.1, this is $S(T)$ (defined in equation 2.3) which supplies an excitatory contribution to the output nuclei. A key aspect of $S(T)$ is that it involves summing over channels. The source of excitation in the basal ganglia is the STN, which sends diffuse projections to the basal ganglia output nuclei (Parent & Smith, 1987). Thus, each output neuron receives many afferents from widespread sources within STN, and so it is plausible that they are performing a summation over channels. In the network model, this is reflected in the fact that neurons in each channel i of the output nuclei compute the quantity $OUT_i(T) = -y_i(T) + \Sigma(T)$, where:

$$\Sigma(T) = \sum_{i=1}^N STN_i(T). \quad (3.3)$$

The model then implements MSPRT if $\Sigma(T) = S(T)$. We now show, first in outline and then more rigorously, how the form of $STN_i(T)$ required in order to ensure $\Sigma(T) = S(T)$ may be enabled by the interaction between STN and GP and the characteristic transfer functions of their neurons.

A first correspondence between equations 2.3 and 3.3 involves summation over channels. Second, since STN receives input $y_i(T)$ from the cortex, this suggests that the STN firing rate should be proportional to the exponent of its input. We also propose that the logarithm in equation 2.3 comes from interactions between STN and GP. The log transform may be thought of as a compression of the range of STN activity, plausibly derived from GP inhibition, since this is, in turn, under STN control. Thus, rather than supplying a fixed decrement in STN activity through a fixed level of inhibition, GP increases its inhibition in response to increased activity in STN.

We now formalize these requirements, resulting in quantitative predictions about the input-output relations of STN and GP neurons. First, we require that the firing rate of neurons in STN is proportional to an

exponential function of its inputs:

$$STN_i(T) = \exp(y_i(T) - GP_i(T)). \quad (3.4)$$

Since STN projects diffusely to GP (Parent & Hazrati, 1995b), we assume that the STN input to GP channel i is $\Sigma(T)$ rather than $STN_i(T)$. The required log transform is obtained by supposing that the firing rate of GP channel i , $GP_i(T)$, is given by

$$GP_i(T) = \Sigma(T) - \ln(\Sigma(T)), \quad (3.5)$$

since, substituting equation 3.5 into equation 3.4, summing over i , and solving for $\Sigma(T)$ then yields $\Sigma(T) = S(T)$.

In summary, an implementation of MSPRT defined by equations 3.1 to 3.5 may be realized by a subset of basal ganglia anatomy, defined in Figure 1B, if the behavior of neurons in STN and GP follows equations 3.4 and 3.5.

As described so far, the model lacks two known pathways within basal ganglia that were shown in Figure 1A. First, the GP functionality defined in equation 3.5 omits afferents from striatal projection neurons associated with D2-type dopamine receptors. Second, GP projections to the output nuclei have not been included in equation 3.1. It has been proposed that these pathways play a critical role in the learning phase, when they block actions that have been punished (Frank et al., 2004). This function is not included in our model, because we address only the computation in the proficient phase. Appendix B shows that incorporation of these pathways into an anatomically more complete scheme still admits a model of basal ganglia that supports MSPRT. Therefore, the model with all pathways shown in Figure 1A also achieves the optimal performance of the MSPRT.

4 Predicted Requirements for STN and GP Physiology Are Validated by Existing Data

In this section we compare the predictions of equations 3.4 and 3.5 concerning the firing rates of STN and GP neurons as a function of their input, with published experimental data. In order to make this comparison, model variables (e.g., $y_i(T)$, $STN_i(T)$, $GP_i(T)$) are assumed to be proportional to experimentally observed neuronal firing rates. Note, however, that proportionality constants are not uniquely specified by the model because a change in any such constant for a particular nucleus can be absorbed by rescaling the weights in projections from this nucleus to other areas. (The use of interpathway weights is illustrated in the anatomically more complete model described in appendix B.)

The forms for STN and GP functionality given in equations 3.4 and 3.5 were derived on the basis of the known anatomy of the basal ganglia and the

assumption that the network involving cortex and basal ganglia implements MSPRT. Since we did not use the physiological properties of STN and GP neurons in deriving equations 3.4 and 3.5, these equations represent predictions of the model for the physiological properties of STN and GP, thereby providing an independent means for testing the model. These predictions are very strong; in particular, the theory of section 3 implies that the firing rate of STN neurons should be proportional to the exponent of its input. Such a relation is highly unusual in most neural populations. Furthermore, such a relation is very different from the STN input-output relations assumed by other models: Gurney et al. (2001a) assume a piecewise linear relation, while Frank et al. (2005) assume a sigmoid relation.

The response properties of STN neurons have been studied extensively (Hallworth, Wilson, & Bevan, 2003; Overton & Greenfield, 1995; Wilson, Weyrick, Terman, Hallworth, & Bevan, 2004). Typically they have nonzero spontaneous firing and can achieve unusually high firing rates. Our proposed exponential form for firing rate as a function of input (see equation 3.4) explains these features since, in the absence of input, the model gives nonzero (unity) output and $\exp(\cdot)$ is a rapidly growing function yielding potentially high firing rates. In order to test the prediction of equation 3.4 quantitatively, we fitted exponential functions to firing rate data in the literature. Figure 2H shows the pooled results of this exercise based on two studies (Hallworth et al., 2003; Wilson et al., 2004). The fit to an exponential function is a good one, consistent with the prediction in equation 3.4.

Second, the theory makes predictions, defined by equation 3.5, concerning the firing rate of GP. First, we show that the function defined by equation 3.5 is roughly linear if we make the reasonable assumption that N (the number of channels or available actions) is large. Thus, since $y_i(T) > 0$, then from equation 2.3, $S(T)$ is bounded below by $\ln(N)$, so that $S(T)$ increases with N . Now, for large $S(T)$, $S(T) \gg \ln(S(T))$, so that the linear term in equation 3.5 dominates, and $GP_i(T)$ becomes an approximately linear function of its input $S(T)$.

We therefore predict that GP neurons display a roughly linear relation between input and firing rate, and two studies validate this. Nambu and Llinas (1994) have established that for those GP neurons that are most influential on the population firing rate, their firing rate is indeed well approximated by a linear function of the injected current (see Figure 2I), a result that is in agreement with an earlier study by Kita and Kitai (1991). In any case, a model in which GP neurons obey an exactly linear input firing rate relation departs little in performance from the model in which GP is described by equation 3.5 (see section 5.3).

Finally, it is intriguing to note that GP also includes two types of neurons whose input-output properties are logarithmic (Nambu & Llinas, 1994). Thus microcircuits within GP making use of intranucleus inhibitory collaterals (Nambu & Llinas, 1997) could also support the exact computation required by equation 3.5 for MSPRT.

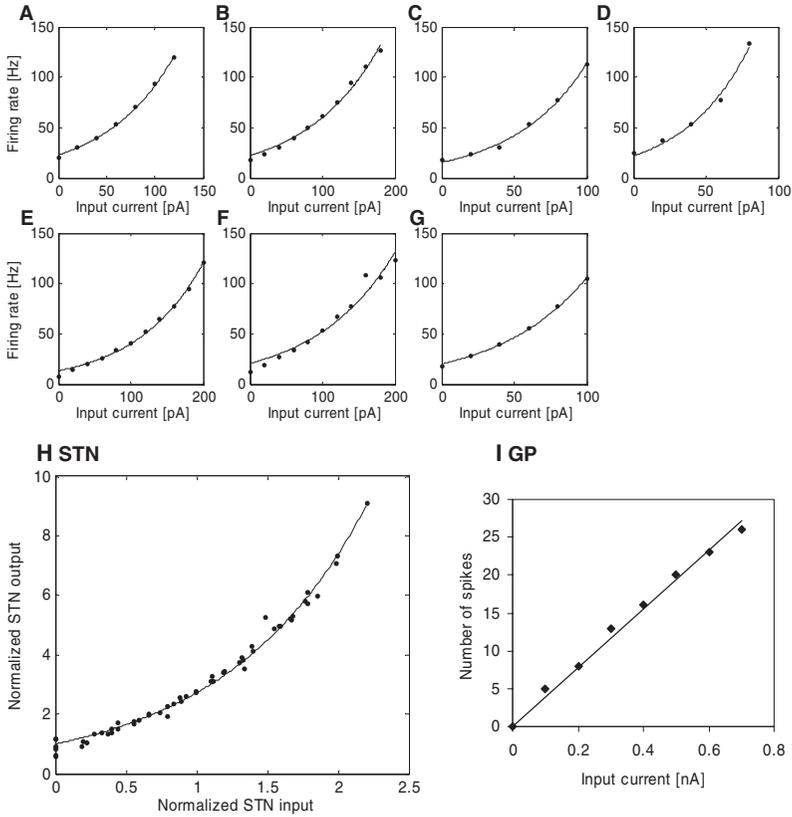


Figure 2: Firing rates f of STN and GP neurons as a function of input current I . (A–D) The panels replot data on the firing rate of STN neurons presented in Hallworth et al. (2003) in Figure 4b, 4f, 12d, and 13d respectively (control condition). (E–G) The panels replot the data from STN presented in Wilson et al. (2004) in Figures 1c, 2c, 2f respectively (control condition). Only firing rates below 135 Hz are shown. Lines show best fit of the function $f = a \exp(bI)$. (H) Scaled data from A–G ($f_j/a, bI_j$) plotted on the same axes for all neurons. (I) Number of spikes n produced by a GP neuron of type II (Nambu & Llinas, 1994) in a 242 ms stimulation interval using current injection I . (The data used in this figure, kindly provided by Atsushi Nambu, come from the same neuron analyzed in Figure 5g of Nambu and Llinas, 1994.)

5 Performance of MSPRT Model of the Basal Ganglia and Its Variants –

The performance of the algorithmically defined model described in section 3 was investigated in simulation.

5.1 Simulation Methods. In all numerical experiments described in this section, we simulated a decision process between N alternative actions, with plausible parameters describing sensory evidence. The evidence $x_i(t)$ was accumulated in integrators Y_i in time steps of $\delta t = 1\text{ms}$. For the correct alternative i , evidence $x_i(t)$ was generated from a normal distribution with mean $\mu^+ \delta t$ and variance $\sigma^2 \delta t$, while for other alternatives, $x_j(t)$ was generated from a normal distribution with mean $\mu^- \delta t$ and variance $\sigma^2 \delta t$. It transpires, in fact, that we require only the values of the signal $\mu^+ - \mu^-$, rather than individual means themselves (see appendix A). This was estimated from a sample participant in experiment 1 from the study of Bogacz, Brown, Moehlis, Holmes, and Cohen (2006), that is, $\mu^+ - \mu^- = 1.41$. An estimate of σ was taken from the same experiment to be 0.33. For each set of parameters, a decision threshold was found numerically that resulted in an error rate of $1\% \pm 0.2\%$ (SE); this search for threshold was repeated 10 times. For each of these 10 thresholds, the decision time was then found in simulation and their average used to construct the data points.

5.2 MSPRT in the Basal Ganglia Outperforms Alternative Decision Mechanisms. It is instructive to see quantitatively how the performance for the MSPRT model compares with that of two other standard models of decision making in the brain: the race model (Vickers, 1970) and a model proposed by Usher and McClelland (2001) (henceforth, the UM model). While the MSPRT has been shown to be asymptotically optimal as the error approaches zero, its performance with finitely large errors has to be evaluated numerically. To do this, we conducted simulations for differing numbers of competing inputs, N , for all three models, with a 1% error rate.

Figure 3A shows that the MSPRT consistently outperforms both the UM and race models (especially in the more realistic large N regime). This result is in agreement with recent work by McMillen and Holmes (2006) (who also showed another feature in Figure 3A—that as N increases, the performance of UM model asymptotically approaches that of the race model). For $N = 2$, the performance of the MSPRT and UM models is very similar since, in this instance, the latter approximates SPRT (Bogacz et al., 2006; Brown et al., 2005).

5.3 The Model is Parametrically Robust. As previously noted, MSPRT specifies a unique value of the cortical gain parameter g^* if MSPRT is to be faithfully implemented. We now analyze the performance of the model with different values of gain $g \neq g^*$ in a general cortical integrator relation $y_i(T) = gY_i(T)$ (instead of $y_i(T) = g^*Y_i(T)$). Equation 3.2 implies that the optimal value of gain g^* depends on the parameters of the inputs to the cortical integrators (μ^+ , μ^- , σ). These parameters are task specific and are therefore unlikely to be known to any neural decision system. It is therefore essential for any biologically realistic implementation of MSPRT

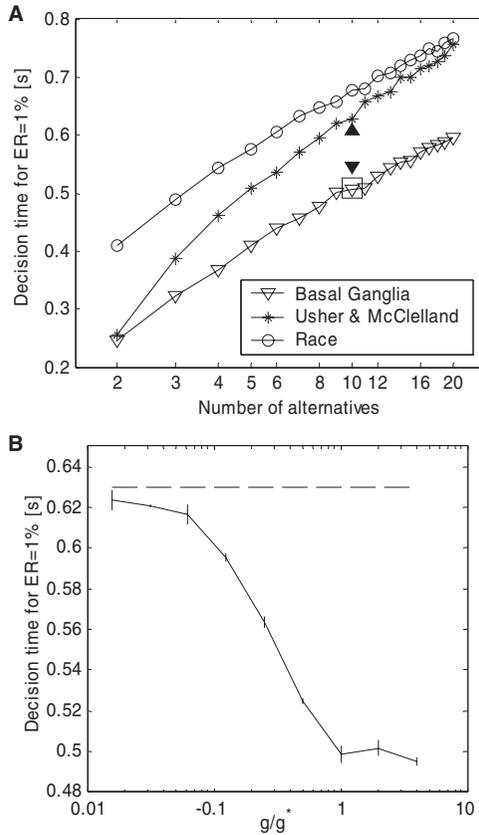


Figure 3: Comparison of decision times (DT) of various models described in the text. Simulation details for the MSPRT model are given in section 5.1. (A) Comparison of DT of MSPRT model, the Usher and McClelland (2001) (UM) model and race model (Vickers, 1970) for different numbers of alternative actions. The standard error of the mean decision time estimation (SEM) was always lower than 7.3 ms. The inhibition and decay parameters of the UM model were set to 100. The isolated filled triangles show DTs for the linearized GP model with two parameter pairs (g, a) (see section 5.3). The triangle pointing up shows DT for default values $g = g^*$ and $a = 1$, and the triangle pointing down for optimal values $g = 0.4g^*$ and $a = 0.84$. The isolated square shows the DT for a version of MSPRT model in which simple cortical integrators are replaced by a UM model with both decay and inhibition parameters equal to 10. (B) Robustness of MSPRT model under variation in the gain parameter. The solid line shows the dependence of DT for $N = 10$ alternatives on the value of parameter g (expressed via its ratio g^*). Error bars indicate SEM. The dashed line shows the decision time of the UM model.

that this mechanism does not significantly deviate from optimality under variation of the gain g from its optimal value, g^* .

Figure 3B shows that the model is indeed robust to changes in g . If $g > g^*$, the decision time does not increase, while if $g < g^*$, the decision time does increase, but it never exceeds that of the UM model (see appendix C). Hence, even if the parameters of the inputs to the cortical integrators are not known, the performance may be optimized by setting g as high as possible.

Turning to the functionality of GP, we evaluated the decrease in performance under an exact linearization of GP with respect to that shown in Figure 3A, for $N = 10$ alternatives. To do this, suppose the firing rate of a neuron in GP is a linear function of its input from STN with proportionality constant a :

$$GP_i(T) = a \sum_{j=1}^N STN_j(T) = a S(T). \quad (5.1)$$

Substituting equation 5.1 into equation 3.4 and summing over i yields

$$\ln(S(T)) + a S(T) = \ln\left(\sum_{i=1}^N \exp(y_i(T))\right). \quad (5.2)$$

equation 5.2 does not have a closed-form solution for $S(T)$, and so this is found by solving equation 5.2 numerically.

Decision times (DT) were contingent on the values of parameters g and a . With default values, $g = g^*$ and $a = 1$, $DT = 607$ ms (with SEM. = ± 3 ms), which is better than the UM model ($DT = 628$ ms) or race model ($DT = 676$ ms). A parameter search yielded optimal performance, with $DT = 545$ ms (± 3 SEM.), with $g = 0.4 g^*$ and $a = 0.84$.

5.4 Competition May Occur in Both Cortex and Basal Ganglia. The UM model (as well as the model of cortical decision making in area LIP by Wang, 2002) assumes that cortical integrators not only integrate evidence (as in the MSPRT model) but also actively compete with one another. Appendix D shows that if the cortex performs a computation equivalent to the UM model, then the activity levels of the basal ganglia output nuclei are exactly the same as in the original MSPRT model. As a consequence, the decision times remain the same, and the system as whole still achieves optimal performance. This is illustrated in Figure 3A, where an isolated open square symbol shows the DT for the model augmented with UM-based cortical processing; this DT is the same as for the original MSPRT model.

6 Discussion

6.1 Summary. Our main result is that a circuit involving cortex and the basal ganglia may be devoted to implementing a powerful (asymptotically optimal) decision mechanism (MSPRT) in a parametrically robust way. Further, our results suggest a division between a core functional anatomy (shown in Figure 1B) and additional pathways (incorporated in the anatomically complete model) that may serve other purposes (e.g., enhancement of robustness and learning) without compromising MSPRT. In addition, the MSPRT model was shown to outperform other decision mechanisms. While the UM and race models avail themselves of simple network implementations, the sophisticated architecture and neural functionality of the basal ganglia appear to have evolved to support the more powerful MSPRT, allowing the brain to make accurate decisions substantially faster than the simpler mechanisms. The model also made several predictions about the physiological properties of STN and GP neurons, which, while consistent with existing data, provide a challenge for further experimental studies to test them *in vitro* and *in vivo* with synaptic input.

6.2 Relationship to Other Models of Decision Making

6.2.1 Action Selection in Basal Ganglia. The model described in this article has exactly the same architecture (shown in Figure 1A) as a previous model of action selection in the basal ganglia in the proficient phase (Gurney et al., 2001a). This architecture has been shown before to exhibit appropriate action selection and switching properties in a computational model (Gurney et al., 2001b). The underlying architecture has also been shown to be functionally robust (from a selection perspective) in a variety of settings. Thus, it has also been shown to perform these functions within the anatomical context supplied by associated thalamocortical loops (Humphries & Gurney, 2002), under the addition of further circuitry intrinsic to the basal ganglia (Gurney, Prescott, et al., 2004), when embodied in a complete, behaving autonomous agent (Prescott, Montes-Gonzalez, Gurney, Humphries, & Redgrave, 2006) and in spiking neuron models (Humphries, Stewart, & Gurney, *in press*; Stewart, Gurney, & Humphries, 2005), which are constrained by significantly more physiological detail than their systems-level counterparts.

The model of Gurney et al. (2001a) and the MSPRT model are consistent in proposing similar functions of individual nuclei. In particular, we suppose here that the GP plays a crucial role in limiting STN activity (via a log transform). This function is similar to that proposed for GP in Gurney et al. (2001a), in which GP automatically limits the excitation of basal ganglia output nuclei in order to allow network mechanisms to perform selection. This work differs from Gurney et al. (2001a) in that it provides an analytic description for the computation performed during the selection, thereby

providing a new framework for understanding why the basal ganglia are organized in the way they are.

The function of STN in our model is similar to that posited by Frank (2005), who proposed that it “can dynamically modulate the threshold for executing responses depending on the degree of response conflict present.” The novelty of our work lies in specifying precisely how STN should modulate this threshold to optimize the performance.

6.2.2 Bayesian Decision Making. MSPRT can be viewed as a Bayesian method of decision making, as it is based on evaluating conditional probabilities using the Bayes theorem (see appendix A). Recently Yu and Dayan (2005) proposed a Bayesian model of attentional modulation of decision making. In this model, the final layer of an abstract neural network performs a computation equivalent to that accomplished by the outputs of the MSPRT model of basal ganglia presented here (i.e., each output unit computes exponents of $L_i(T)$). The novelty of our work lies in showing how this computation may be performed by an identified, biological network of neurons, the basal ganglia.

6.2.3 Cortical Decision Making. From the theoretical perspective, the use of integrated evidence leaves open the possibility that the cortex may operate as the first stage of a two-stage decision process, in which cortical mechanisms of the kind posited in the UM model (for example) make a first-pass filter for actions with small saliences, thereby preventing these requests from propagating to the basal ganglia for further processing. This possibility was confirmed by mathematical analysis and simulation, where identical results were obtained after incorporation of a first stage consisting of a UM network (rather than the simple integration implied in equation 2.1).

6.2.4 Integration of Evidence. The model presented here assumes that cortical neurons integrate evidence in support of alternative actions. Several mechanisms have been proposed for how this may occur; for example, Wang (2002) proposed that the integration occurs via excitatory connections in the cortex. However, if the basal ganglia model presented here is to be a universally applicable solution to the problem of action selection, then the appearance of accumulated evidence must be guaranteed at its inputs under all circumstances, irrespective of the specific action request and brain system generating it. Anatomically, the basal ganglia form a component of loops consisting of projections from cortex to basal ganglia, then to thalamus and back to cortex (Alexander & Crutcher, 1990). In a computational study of basal ganglia and cortex, Humphries and Gurney (2002) showed that cortical regions receiving thalamic input had their activity levels amplified beyond those of sensory areas that did not. This raises the intriguing

possibility that the feedback in these loops could serve to bootstrap evidence accumulation in cortex so that no separate mechanism is required.

6.2.5 Reinforcement Learning. As stated in section 1, this article focuses on the proficient phase of task acquisition. However, during learning, it is still necessary to identify the behavioral state (stimulus and ongoing behavior) and represent the stimulus response mapping in some way. Both processes are aspects of the decision-making process discussed in this article, and so it is useful to speculate how it might be possible to unify the accounts of decision making and learning into a single coherent framework.

To develop these ideas, we consider a version of the motion discrimination task described in section 2.1 in which the stimulus-response mapping is constantly modified and hence must be learned by the animal. In this case, it is unlikely that the stimulus-response mapping would be represented in the connections between MT and FEF, because these connections would have to be rapidly and continuously modified. Many models based on experimental data would assume that in this experiment, the stimulus-response mapping would be stored in much more plastic synapses in the prefrontal cortex or the striatum (Doya, 2000; Miller & Cohen, 2001; O'Doherty et al., 2004). This is also the kind of scheme considered by Ashby, Alfonso-Turken, and Waldron (1998) in the context of category learning with motor response. Here, early stimulus-response mappings are learned in the basal ganglia, while slower consolidation takes place in direct mappings between sensory and motor cortices.

Further, we note that behavioral state identification during learning could rely on integration of information, just as it does during the proficient phase in the theory presented here. However, in the case of learning, the integration may occur in regions different from those used during the proficient phase. For example, Gold and Shadlen (2003) have shown that in the version of the motion discrimination task in which the mapping between stimulus and direction of saccade is not known during the information integration, the integration does not occur in FEF.

If a unified framework encompassing action selection and learning could be developed along the lines outlined above, it would simultaneously allow predictions of the probabilities of taking alternative actions, as well as the probability distributions of onsets of action initiations (reaction times). However, before such a unified account is developed, a number of questions must be answered, in particular, where in the brain the evidence supporting alternative actions is computed in the learning phase. To address this question, we look forward to studies of neuronal responses in cortex and basal ganglia in the version of the motion discrimination task described above.

6.2.6 Working Memory. O'Reilly and Frank (2006) have proposed that the basal ganglia gate access to working memory and decides whether a newly presented stimulus should be stored in the working memory. According to

Alexander et al.'s (1986) multiple loop scheme, this kind of decision would be performed by the dorso-lateral prefrontal circuit. As noted in section 2.5, this letter focuses on motor and oculomotor circuits. It would therefore be interesting to investigate whether the basal ganglia implementing MSPRT could also optimize selection within working memory and, indeed, cognitive selection in general.

6.3 Relationship to Other Experimental Data

6.3.1 Psychological Data. A model of decision making must be consistent with the rich body of psychological data concerning reaction times (RT). Other psychological models are consistent with these data (Ratcliff, Van Zandt, & McKoon, 1999; Usher & McClelland, 2001), and consistency in this respect will therefore not distinguish our model in favor of these alternatives but, rather, is a necessary requirement for its psychological plausibility. Our model is indeed completely consonant with the account of RT data in two alternative choice paradigms given by the diffusion and SPRT models (e.g., Ratcliff et al., 1999), since, under these circumstances, MSPRT reduces to SPRT. For more than two alternatives, it is interesting to note that the decision time of the MSPRT model (shown in Figure 3A) is approximately proportional to the logarithm of the number of alternatives (cf. McMillen & Holmes, 2006), thus following the experimentally observed Hick's law (Teichner & Krebs, 1974) describing RT as a function of number of choices. Further support for the basal ganglia as a psychologically plausible response mechanism is provided in recent work by Stafford and Gurney (2004, in press).

6.3.2 The Neural Representation of Actions. In this letter, we assumed for simplicity that there was an anatomically separate channel for each possible action. However, this raises the question of what constitutes a separate action. For example, is moving one's hand 10 cm to the left a different action from moving it 15 cm to the left? If not, then how are these actions differentiated? If they are different actions (with respect to basal ganglia selection), then the number of actions is potentially infinite, and the basal ganglia are confronted with the seemingly impossible task of representing an infinite number of discrete channels.

Neurophysiological data provide a clue to a possible answer to these questions. Georgopoulos et al. (1983) studied neuronal responses in basal ganglia in a task (akin to the example above) in which different stimuli required an animal to move its hand in the same direction with three different amplitudes. They noticed that some hand-selective neurons had activity proportional to movement amplitude, while other neurons had activity inversely proportional to the amplitude (they responded most for short movements). This suggests that although movements of different joints may be represented by the separate neuronal populations (channels), the

fine tuning of the movement is represented in a distributed fashion within a channel.

It will therefore be of interest to extend the current theory to incorporate coding by distributed representations. The current MSPRT model describes activity of a neuronal population selective for each alternative within each nucleus by a single variable (corresponding to activity in a localist unit). Recently Bogacz (2006) has shown how to map linear localist decision networks into computationally equivalent distributed decision networks and derived the parameters of decision networks with distributed representations implementing SPRT. Although this network mapping process cannot be directly applied to the MSPRT model (because it includes nonlinear processing in STN), it is likely that a similar approach for particular types of nonlinearities present in the MSPRT model may be developed.

6.3.3 Responses of Striatal Neurons. As noted in section 1, the MSPRT model aims to provide a general framework for understanding computation in the basal ganglia as a whole during action selection in the proficient phase. Therefore, the model does not aim to incorporate all known data on basal ganglia neurons and, in particular, does not aim to explain data relating to the learning phase (during which the striatum is known to play a prominent role). However, it is of interest to see whether the behavior of the population of striatal projection neurons in the proficient phase is consistent with our ideas.

In the MSPRT model, we assume, in accordance with experimental observations (Crutcher & DeLong, 1984b; Georgopoulos et al., 1983), that during the proficient phase, the activity of striatal neurons encoding certain actions reflects the activity in corresponding cortical motor or oculomotor regions. Note that the above assumption does not prevent striatal neurons selective for particular actions to be modulated by expected reward in the proficient phase, as it has been shown that the cortical integrators are modulated by expected reward in the study by Platt and Glimcher (1999) in which the stimulus-response mapping was kept constant for many weeks of training and experiment.

We assumed that in the proficient phase, the neurons in striatum selective for actions should reflect the activity of cortical integrators. This predicts that in the motion discrimination task described in section 2.1, striatal neurons selective for alternative directions of eye movements should exhibit gradually increasing firing rates similar to those in the cortex. This prediction may seem to contradict the observation that striatal neurons are bistable (Wilson, 1995) with an active "up" state, and an inactive "down" state. However, Okamoto, Isomura, Takada, and Fukai (2005) have shown recently that even if neurons are bistable, and if the probability of the onset of the up state depends on the magnitude of input, these neurons may implement information integration, and their activity averaged across trials may be linearly increasing.

6.3.4 Dopaminergic Modulation. In this article, we do not analyze the influence of dopamine on basal ganglia computation. However, there are two types of dopamine release within the basal ganglia. First, phasic release (brief pulses of dopamine) is associated with salient or unexpected behavioral events. In particular, it has been proposed that the phasic dopamine signal represents a variable in temporal difference reinforcement learning, namely, the reward prediction error (i.e., the difference between the actual and the predicted levels of expected reward) (Montague et al., 1996; Schultz et al., 1997). Since we do not address learning here, we do not consider phasic release further.

A second type of dopamine release provides tonic or background levels, severe lowering of which can result in Parkinson's disease (Obeso et al., 2000). A recent modeling study of the basal ganglia circuit (Gurney, Humphries, Wood, Prescott, & Redgrave, 2004) has indicated that tonic levels of dopamine may influence the speed-accuracy trade-off in making responses. This is consistent with experimental results showing that the level of tonic dopamine influences RTs (Amalric & Koob, 1987; Amalric, Moukhles, Nieoullon, & Daszuta, 1995). It will be interesting to determine theoretically to what extent the tonic dopamine level can influence the speed-accuracy trade-off while still preserving the optimality of MSPRT and whether this mechanism can play a role in finding the speed-accuracy trade-off that maximizes the rate of reward acquisition in tasks including repeating sequences of choices (Bogacz et al., 2006; Simen, Cohen, & Holmes, 2006).

Appendix A: The MSPRT

This section supplies details of the derivation of equation 2.2. Noting the definition in the main text $P_i(T) = P(H_i | \text{input}(T))$, then, from Baye's theorem,

$$P_i(T) = \frac{P(\text{input}(T) | H_i) P(H_i)}{P(\text{input}(T))}. \quad (\text{A.1})$$

Notice that the hypotheses H_i are mutually exclusive (as x_i cannot simultaneously have mean μ^+ and μ^-). Further, we assume that the set of hypotheses H_i covers all possibilities concerning the distribution of sensory inputs (a standard assumption is statistical testing). Then the denominator of equation A.1 can be written as

$$P(\text{input}(T)) = \sum_{k=1}^N P(\text{input}(T) \wedge H_k), \quad (\text{A.2})$$

but from the definition of conditional probability, $P(\text{input}(T) \wedge H_k) = P(\text{input}(T)|H_k)P(H_k)$, so that equation A.1 can be written as

$$P_i(T) = \frac{P(\text{input}(T)|H_i)P(H_i)}{\sum_{k=1}^N P(\text{input}(T)|H_k)P(H_k)}. \tag{A.3}$$

We assume that we do not have any prior knowledge about which of the hypotheses is more likely, so all prior probabilities must be equal to each other, that is, $P(H_i) = 1/N$, and hence they cancel in equation A.3 and we obtain the original form of MSPRT given by Baum and Veeravalli (1994):

$$P_i(T) = \frac{P(\text{input}(T)|H_i)}{\sum_{k=1}^N P(\text{input}(T)|H_k)}. \tag{A.4}$$

We now compute the logarithm of P_i , defined by equation A.4:

$$\begin{aligned} L_i(T) &= \ln P_i(T) = \ln P(\text{input}(T)|H_i) \\ &\quad - \ln \sum_{k=1}^N \exp(\ln P(\text{input}(T)|H_k)). \end{aligned} \tag{A.5}$$

Equation A.5 already has a form similar to that in equation 2.2. We now show how to obtain equation 2.2 exactly. We first compute the term $\ln P(\text{input}(T)|H_i)$ that occurs in equation A.5:

$$\begin{aligned} \ln P(\text{input}(T)|H_i) &= \ln \prod_{t=1}^T \left(f_{(\mu^+, \sigma)}(x_i(t)) \prod_{\substack{j=1 \\ j \neq i}}^N f_{(\mu^-, \sigma)}(x_j(t)) \right) \\ &= \sum_{t=1}^T \ln f_{(\mu^+, \sigma)}(x_i(t)) + \sum_{t=1}^T \sum_{\substack{j=1 \\ j \neq i}}^N \ln f_{(\mu^-, \sigma)}(x_j(t)), \end{aligned}$$

where $f_{(\mu, \sigma)}$ denotes the probability density function of a normal distribution with mean μ and standard deviation σ . Therefore,

$$\begin{aligned} \ln P(\text{input}(T)|H_i) &= \sum_{t=1}^T \left(\ln \frac{1}{\sqrt{2\pi}\sigma} - \frac{(x_i(t) - \mu^+)^2}{2\sigma^2} \right) \\ &\quad + \sum_{\substack{j=1 \\ j \neq i}}^N \sum_{t=1}^T \left(\ln \frac{1}{\sqrt{2\pi}\sigma} - \frac{(x_j(t) - \mu^-)^2}{2\sigma^2} \right) \end{aligned}$$

$$\begin{aligned}
&= NT \ln \frac{1}{\sqrt{2\pi}\sigma} + \frac{1}{2\sigma^2} \left(-T((\mu^+)^2 + (N-1)(\mu^-)^2) \right. \\
&\quad \left. + \sum_{j=1}^N \sum_{t=1}^T (-x_j^2(t) + 2\mu^- x_j(t)) \right) + \frac{\mu^+ - \mu^-}{\sigma^2} \sum_{t=1}^T x_i(t).
\end{aligned}$$

The first two terms on the right-hand side do not depend on i , and so, denoting their sum by C , we have

$$\ln P(\text{input}(T) | H_i) = C + g^* Y_i(T), \quad (\text{A.6})$$

where $g^* = (\mu^+ - \mu^-)/\sigma^2$, and $Y_i(T)$ is defined in equation 2.1. Hence substituting equation A.6 into A.5, we obtain:

$$\begin{aligned}
L_i(T) &= C + g^* Y_i(T) - \ln \left(\exp(C) \sum_{k=1}^N \exp(g^* Y_k(T)) \right) \\
&= g^* Y_i(T) - \ln \sum_{k=1}^N \exp(g^* Y_k(T))
\end{aligned}$$

which gives equation 2.2.

Appendix B: Basal Ganglia Model, Including All Major Pathways _____

To accomplish inclusion of all pathways shown in Figure 1A, we introduce a model with weighted connection strengths. For simplicity, we assume that all the excitatory weights are equal to 1 and denote the inhibitory weight from nucleus A to nucleus B by $w_{A \rightarrow B}$. Denote the striatal projection neurons whose dopaminergic receptors are predominantly of D1 and D2 types by S1 and S2, respectively. The activity levels of basal ganglia nuclei in the network of Figure 1A are then given by:

$$GP_i = \sum_{j=1}^N STN_j - \ln \left(\sum_{j=1}^N STN_j \right) - w_{S2 \rightarrow GP} y_i \quad (\text{B.1})$$

$$STN_i = \exp(y_i - w_{GP \rightarrow STN} GP_i) \quad (\text{B.2})$$

$$OUT_i = -w_{S1 \rightarrow OUT} y_i + \sum_{j=1}^N STN_j - w_{GP \rightarrow OUT} GP_i. \quad (\text{B.3})$$

We now derive constraints on the weights that must be satisfied so $OUT_i(T) = -L_i(T)$ as defined by equation 2.2. Substituting equation B.1 into B.2,

$$STN_i = \exp \left(y_i (1 + w_{GP \rightarrow STN} w_{S2 \rightarrow GP}) - w_{GP \rightarrow STN} \sum_{j=1}^N STN_j + w_{GP \rightarrow STN} \ln \left(\sum_{j=1}^N STN_j \right) \right).$$

Summing over i and rearranging terms, we get

$$\begin{aligned} \sum_{i=1}^N STN_i &= \sum_{i=1}^N \exp(y_i (1 + w_{GP \rightarrow STN} w_{S2 \rightarrow GP})) \\ &\times \exp \left(w_{GP \rightarrow STN} \sum_{i=1}^N STN_i \right)^{-1} \cdot \left(\sum_{i=1}^N STN_i \right)^{w_{GP \rightarrow STN}}. \end{aligned}$$

Taking the logarithm of both sides and rearranging terms,

$$\begin{aligned} \ln \left(\sum_{i=1}^N \exp(y_i (1 + w_{GP \rightarrow STN} w_{S2 \rightarrow GP})) \right) &= (1 - w_{GP \rightarrow STN}) \ln \left(\sum_{i=1}^N STN_i \right) \\ &+ w_{GP \rightarrow STN} \sum_{i=1}^N STN_i. \end{aligned} \quad (B.4)$$

Now substitute equation B.1 into B.3:

$$\begin{aligned} OUT_i &= -y_i (w_{S1 \rightarrow OUT} - w_{GP \rightarrow OUT} w_{S2 \rightarrow GP}) \\ &+ (1 - w_{GP \rightarrow OUT}) \sum_{j=1}^N STN_j + w_{GP \rightarrow OUT} \ln \left(\sum_{j=1}^N STN_j \right). \end{aligned} \quad (B.5)$$

Comparing equations B.4 and B.5, we note that if the following condition is satisfied,

$$w_{GP \rightarrow OUT} = 1 - w_{GP \rightarrow STN}, \quad (B.6)$$

then we can substitute equation B.4 into B.5 and obtain

$$OUT_i = -y_i (w_{S1 \rightarrow OUT} - w_{GP \rightarrow OUT} w_{S2 \rightarrow GP}) + \ln \left(\sum_{j=1}^N \exp(y_j (1 + w_{GP \rightarrow STN} w_{S2 \rightarrow GP})) \right). \quad (B.7)$$

In order to satisfy $OUT_i(T) = -L_i(T)$, the two coefficients of y_i in equation B.7 must be equal, and the gain coefficient must be modified accordingly. Thus, the following must be satisfied:

$$w_{S1 \rightarrow OUT} - w_{GP \rightarrow OUT} w_{S2 \rightarrow GP} = 1 + w_{GP \rightarrow STN} w_{S2 \rightarrow GP}. \quad (B.8)$$

Substituting the constraint B.6 into B.8 gives

$$w_{S1 \rightarrow OUT} = 1 + w_{S2 \rightarrow GP}. \quad (B.9)$$

The optimal value of the gain must be modified and becomes

$$g^* = \frac{1}{1 + w_{GP \rightarrow STN} w_{S2 \rightarrow GP}} \frac{\mu^+ - \mu^-}{\sigma^2}.$$

Note that with the additional pathways, the optimal gain g^* is less than that required without these pathways. In summary, the network in Figure 1A implements MSPRT if the inhibitory weights satisfy constraints B.6 and B.9.

Appendix C: Performance for Nonoptimal Values of Parameter g _____

This section describes the performance of the basal ganglia model when parameter g has nonoptimal values. Section C.1 shows that if $g > g^*$, the decision time does not increase since, as $g \rightarrow \infty$, the model converges to the other asymptotically optimal test MSPRT_b. Section C.2 shows that if $g < g^*$, the decision time does increase, but it never exceeds that of the UM model (Usher & McClelland, 2001) because as $g \rightarrow 0$, the proposed model converges to an approximation of the UM model.

C.1 Overestimation of Gain. We first describe the other asymptotically optimal test MSPRT_b (Dragalin et al., 1999) and go on to show that as $g \rightarrow \infty$, the model approximates MSPRT_b. In this test, after each sample,

the following ratios are computed (Dragalin et al., 1999):

$$RB_i(T) = \frac{P(\text{input}(T) | H_i)}{\max_{\substack{1 \leq k \leq N \\ k \neq i}} P(\text{input}(T) | H_k)}.$$

The decision is made whenever one of the ratios exceeds a threshold. From equation A.6, the logarithms of the above ratios for the hypotheses defined in the main text become

$$LB_i(T) = g \left(Y_i(T) - \max_{\substack{1 \leq k \leq N \\ k \neq i}} Y_k(T) \right). \tag{C.1}$$

The decision is made whenever one of $LB_i(T)$ exceeds a threshold. Let $Y_{m1}(T), Y_{m2}(T), \dots, Y_{mi}(T), \dots, Y_{mN}(T)$ be the ordered sequence of $Y_i(T)$ with $Y_{m1}(T)$ the largest member. Then, according to equation C.1, the decision is made whenever $Y_{m1}(T) - Y_{m2}(T)$ exceeds a threshold. We now show that the basal ganglia model in the main text works in just this way, as $g \rightarrow \infty$. The log-ratio $L_i(T)$ in equation 2.2 of the main text is equal to

$$\begin{aligned} L_i(T) &= gY_i(T) - \ln \left(\exp(gY_{m1}(T)) + \exp(gY_{m2}(T)) + \sum_{j=3}^N \exp(gY_{mj}(T)) \right) \\ &= gY_i(T) - \ln \left(\exp(gY_{m1}(T)) \left(1 + \exp(g(Y_{m2}(T) - Y_{m1}(T))) \right. \right. \\ &\quad \left. \left. + \sum_{j=3}^N \exp(g(Y_{mj}(T) - Y_{m1}(T))) \right) \right) \\ &= gY_i(T) - gY_{m1}(T) - \ln \left(1 + \exp(g(Y_{m2}(T) - Y_{m1}(T))) \right. \\ &\quad \left. + \sum_{j=3}^N \exp(g(Y_{mj}(T) - Y_{m1}(T))) \right). \end{aligned}$$

A decision will be made when the largest among $L_i(T)$ exceeds a threshold z , that is, when

$$\begin{aligned} gY_{m1}(T) - gY_{m1}(T) - \ln \left(1 + \exp(g(Y_{m2}(T) - Y_{m1}(T))) \right. \\ \left. + \sum_{j=3}^N \exp(g(Y_{mj}(T) - Y_{m1}(T))) \right) > z. \end{aligned}$$

Two first terms cancel. Taking the negative, applying the $\exp(\cdot)$ function, and subtracting 1 gives

$$\exp(g(Y_{m2}(T) - Y_{m1}(T))) + \sum_{j=3}^N \exp(g(Y_{mj}(T) - Y_{m1}(T))) < z',$$

where z' is another threshold. After some manipulation, we have

$$\exp(g(Y_{m2}(T) - Y_{m1}(T))) \left[1 + \sum_{j=3}^N \exp(g(Y_{mj}(T) - Y_{m2}(T))) \right] < z'. \quad (\text{C.2})$$

Since $Y_i(T)$ are sums of samples from continuous normal distributions, then $Y_{m2}(T) > Y_{m3}(T)$ with probability 1 for $T > 0$. Therefore, as $g \rightarrow \infty$, the expressions $g(Y_{mj}(T) - Y_{m2}(T)) \rightarrow -\infty$, and the content of the square brackets in equation C.2 converges to 1. Thus, taking the logarithm and dividing by $-g$ gives

$$Y_{m1}(T) - Y_{m2}(T) > z''$$

(where z'' is also a threshold). Hence, as g goes to infinity, the proposed model makes a decision whenever the difference between the two largest $Y_i(T)$ exceeds a threshold, which is equivalent to MSPRT_b .

C.2 Underestimation of Gain. In this section we show that as $g \rightarrow 0$, the model converges to an approximation of the UM model. This model consists of N mutually inhibiting leaky integrators whose dynamics are described by

$$\dot{u}_i(T) = x_i(T) - ku_i(T) - w \sum_{\substack{j=1 \\ j \neq i}}^N u_j(T), \quad u_i(0) = 0, \quad (\text{C.3})$$

where k determines a leakage rate constant and w is the weight of inhibitory connections between the integrators. McMillen and Holmes (2005) showed that the performance of the UM model is optimized when $k = w$ and both go to infinity. In this case, putting $Y_i(T) = \int_0^T x_i(t)dt$, the decision is made whenever any of the following differences,

$$LUM_i(T) = Y_i(T) - \frac{1}{N} \sum_{j=1}^N Y_j(T),$$

exceeds a decision threshold (McMillen & Holmes, 2006). We now show that exactly the same criterion for a decision applies in the proposed basal ganglia model when $g \rightarrow 0$.

For small g , we may approximate the right-hand side of equation 2.2 of the main text by retaining only linear terms in Taylor expansions. Thus, working on the exponential

$$L_i(T) \rightarrow gY_i(T) - \ln \left(\sum_{k=1}^N (1 + gY_k(T)) \right) = gY_i(T) - \ln \left(N + g \sum_{k=1}^N Y_k(T) \right)$$

and then the $\ln()$ function,

$$L_i(T) \rightarrow gY_i(T) - \ln N - g \frac{1}{N} \sum_{k=1}^N Y_k(T).$$

A decision will be made whenever any $L_i(T)$ exceeds a threshold z , that is, when

$$gY_i(T) - \ln N - g \frac{1}{N} \sum_{k=1}^N Y_k(T) > z.$$

Subtracting $\ln N$ and dividing by g , we get the following condition for decision:

$$Y_i(T) - \frac{1}{N} \sum_{k=1}^N Y_k(T) > z'.$$

Hence, as $g \rightarrow 0$ the proposed basal ganglia model makes a decision under the same conditions as the UM with optimal values of inhibition and decay ($w = k$, and $w, k \rightarrow \infty$).

Appendix D: Model with Competing Cortical Integrators _____

Let us consider the UM model described in section C.2, in which decay is equal to inhibition ($w = k$, it is one of the conditions required for optimal performance of UM model—see above) and the integrators have optimal gain. Then equation C.3 becomes

$$\dot{u}_i(T) = g^* x_i(T) - w \sum_{j=1}^N u_j(T), \quad u_i(0) = 0. \quad (\text{D.1})$$

Note that in the above case, all the integrators receive exactly the same inhibition; hence, in the limit of intervals between samples going to 0, the relationship between variables y_i of the MSPRT model and u_i of the UM model becomes

$$\forall i \quad y_i(T) = u_i(T) + c(T),$$

that is, the variables differ by a term $c(T)$, which, although it differs over time (it will typically be negative), is the same for all integrators. However, we now show that if the same term $c(T)$ is added to the cortical firing rate of all integrators, the activity of output nuclei does not change. Thus,

$$\begin{aligned} OUT_i(T) &= -(y_i(T) + c(T)) + \ln \left(\sum_{j=1}^N \exp(y_j(T) + c(T)) \right) \\ &= -y_i(T) - c(T) + \ln \left(\exp(c(T)) \sum_{j=1}^N \exp(y_j(T)) \right) \\ &= -y_i(T) + \ln \left(\sum_{j=1}^N \exp(y_j(T)) \right). \end{aligned}$$

Since the activity of the output nuclei does not change, a version of MSPRT model in which simple cortical integration is replaced by UM model of equation D.1 achieves the same decision times (for any fixed error rate) as the original MSPRT model.

The above also explains a correspondence problem that arises in the MSPRT model because it assumes that at the beginning of the decision process, $y_i = 0$, whereas the baseline firing rate of cortical neurons is nonzero.

Acknowledgments

This work was supported by EPSRC grants: EP/C516303/1 and EP/C514416/1. We thank Atshushi Nambu for providing data used in Figure 2I. We thank Philip Holmes, Jonathan D. Cohen, Paul Overton, Sander Nieuwenhuis, Eric Shea-Brown, and Tobias Larsen for reading the earlier version of the manuscript and very useful comments, and Peter Redgrave for discussion.

References

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

- Alexander, G. E., DeLong, M. R., & Strick, P. L. (1986). Parallel organization of functionally segregated circuits linking basal ganglia and cortex. *Annu. Rev. Neurosci.*, 9, 357–381.
- Amalric, M., & Koob, G. F. (1987). Depletion of dopamine in the caudate nucleus but not in nucleus accumbens impairs reaction-time performance in rats. *J. Neurosci.*, 7(7), 2129–2134.
- Amalric, M., Moukhles, H., Nieoullon, A., & Daszuta, A. (1995). Complex deficits on reaction time performance following bilateral intrastriatal 6-OHDA infusion in the rat. *Eur. J. Neurosci.*, 7(5), 972–980.
- Ashby, F. G., Alfonso-Reese, L. A., Turken, A. U., & Waldron, E. M. (1998). A neuropsychological theory of multiple systems in category learning. *Psychol. Rev.*, 105, 442–481.
- Ashby, F. G., & Spiering, B. J. (2004). The neurobiology of category learning. *Behav. Cogn. Neurosci. Rev.*, 3(2), 101–113.
- Barnard, G. (1946). Sequential tests in industrial statistics. *Journal of Royal Statistical Society Supplement*, 8, 1–26.
- Baum, C. W., & Veeravalli, V. V. (1994). A sequential procedure for multihypothesis testing. *IEEE Transactions on Information Theory*, 40, 1996–2007.
- Bevan, M. D., Smith, A. D., & Bolam, J. P. (1996). The substantia nigra as a site of synaptic integration of functionally diverse information arising from the ventral pallidum and the globus pallidus in the rat. *Neuroscience*, 75(1), 5–12.
- Bogacz, R. (2006). *Optimal decision networks with distributed representation*. Manuscript submitted for publication.
- Bogacz, R., Brown, E., Moehlis, J., Holmes, P., & Cohen, J. D. (2006). The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced choice tasks. *Psychol. Rev.*, 113, 700–765.
- Britten, K. H., Shadlen, M. N., Newsome, W. T., & Movshon, J. A. (1993). Responses of neurons in macaque MT to stochastic motion signals. *Vis. Neurosci.*, 10(6), 1157–1169.
- Brown, E., Gao, J., Holmes, P., Bogacz, R., Gilzenrat, M., & Cohen, J. D. (2005). Simple networks that optimize decisions. *International Journal of Bifurcations and Chaos*, 15, 803–826.
- Brown, J. W., Bullock, D., & Grossberg, S. (2004). How laminar frontal cortex and basal ganglia circuits interact to control planned and reactive saccades. *Neural Netw.*, 17(4), 471–510.
- Chevalier, G., Vacher, S., Deniau, J. M., & Desban, M. (1985). Disinhibition as a basic process in the expression of striatal functions. I. The striato-nigral influence on tecto-spinal/tecto-diencephalic neurons. *Brain Res.*, 334(2), 215–226.
- Crutcher, M. D., & DeLong, M. R. (1984a). Single cell studies of the primate putamen. I. Functional organization. *Exp. Brain Res.*, 53(2), 233–243.
- Crutcher, M. D., & DeLong, M. R. (1984b). Single cell studies of the primate putamen. II. Relations to direction of movement and pattern of muscular activity. *Exp. Brain Res.*, 53(2), 244–258.
- Dayan, P. (2001). Levels of analysis in neural modeling. In R. A. Wilson & F. C. Keil (Eds.) *Encyclopedia of cognitive science*. London: Macmillan.
- Deniau, J. M., & Chevalier, G. (1985). Disinhibition as a basic process in the expression of striatal functions. II. The striato-nigral influence on

- thalamocortical cells of the ventromedial thalamic nucleus. *Brain Res.*, 334(2), 227–233.
- Doya, K. (2000). Complementary roles of basal ganglia and cerebellum in learning and motor control. *Curr. Opin. Neurobiol.*, 10(6), 732–739.
- Dragalin, V. P., Tertakovskiy, A. G., & Veeravalli, V. V. (1999). Multihypothesis sequential probability ratio tests—part I: asymptotic optimality. *IEEE Transactions on Information Theory*, 45, 2448–2461.
- Faull, R. L., & Mehler, W. R. (1978). The cells of origin of nigrotectal, nigrothalamic and nigrostriatal projections in the rat. *Neuroscience*, 3(11), 989–1002.
- Frank, M. J. (2005). *When and when not to use your subthalamic nucleus: Lessons from a computational model of the basal ganglia*. Paper presented at the International Workshop on Modelling Natural Action Selection, Edinburgh.
- Frank, M. J., Seeberger, L. C., & O'Reilly, R. C. (2004). By carrot or by stick: Cognitive reinforcement learning in Parkinsonism. *Science*, 306(5703), 1940–1943.
- Georgopoulos, A. P., DeLong, M. R., & Crutcher, M. D. (1983). Relations between parameters of step-tracking movements and single cell discharge in the globus pallidus and subthalamic nucleus of the behaving monkey. *J. Neurosci.*, 3(8), 1586–1598.
- Gerfen, C. R., & Wilson, C. J. (1996). The basal ganglia. In A. Bjorklund, T. Hokfelt, & L. Swanson (Eds.), *Handbook of chemical neuroanatomy* (Vol. 12, pp. 369–466). New York: Elsevier.
- Gerfen, C. R., & Young, W. S. III. (1988). Distribution of striatonigral and striatopallidal peptidergic neurons in both patch and matrix compartments: An in situ hybridization histochemistry and fluorescent retrograde tracing study. *Brain Res.*, 460(1), 161–167.
- Gold, J. I., & Shadlen, M. N. (2001). Neural computations that underlie decisions about sensory stimuli. *Trends Cogn. Sci.*, 5(1), 10–16.
- Gold, J. I., & Shadlen, M. N. (2002). Banburismus and the brain: Decoding the relationship between sensory stimuli, decisions, and reward. *Neuron*, 36(2), 299–308.
- Gold, J. I., & Shadlen, M. N. (2003). The influence of behavioral context on the representation of a perceptual decision in developing oculomotor commands. *J. Neurosci.*, 23(2), 632–651.
- Gurney, K., Humphries, M., Wood, R., Prescott, T. J., & Redgrave, P. (2004). Testing computational hypotheses of brain systems function: A case study with the basal ganglia. *Network*, 15(4), 263–290.
- Gurney, K., Prescott, T. J., & Redgrave, P. (2001a). A computational model of action selection in the basal ganglia. I. A new functional anatomy. *Biol. Cybern.*, 84(6), 401–410.
- Gurney, K., Prescott, T. J., & Redgrave, P. (2001b). A computational model of action selection in the basal ganglia. II. Analysis and simulation of behaviour. *Biol. Cybern.*, 84(6), 411–423.
- Gurney, K., Prescott, T. J., Wickens, J. R., & Redgrave, P. (2004). Computational models of the basal ganglia: From robots to membranes. *Trends Neurosci.*, 27(8), 453–459.
- Hallworth, N. E., Wilson, C. J., & Bevan, M. D. (2003). Apamin-sensitive small conductance calcium-activated potassium channels, through their selective coupling

- to voltage-gated calcium channels, are critical determinants of the precision, pace, and pattern of action potential generation in rat subthalamic nucleus neurons in vitro. *J. Neurosci.*, 23(20), 7525–7542.
- Humphries, M. D., & Gurney, K. N. (2002). The role of intra-thalamic and thalamo-cortical circuits in action selection. *Network*, 13(1), 131–156.
- Humphries, M. D., Stewart, R. D., & Gurney, K. N. (in press). A physiologically plausible model of action selection and oscillatory activity in the basal ganglia. *J. Neurosci.*
- Kha, H. T., Finkelstein, D. I., Tomas, D., Drago, J., Pow, D. V., & Horne, M. K. (2001). Projections from the substantia nigra pars reticulata to the motor thalamus of the rat: Single axon reconstructions and immunohistochemical study. *J. Comp. Neurol.*, 440(1), 20–30.
- Kim, J. N., & Shadlen, M. N. (1999). Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. *Nat. Neurosci.*, 2(2), 176–185.
- Kita, H., & Kitai, S. T. (1991). Intracellular study of rat globus pallidus neurons: Membrane properties and responses to neostriatal, subthalamic and nigral stimulation. *Brain Res.*, 564(2), 296–305.
- Laming, D. R. J. (1968). *Information theory of choice reaction time*. New York: Wiley.
- McMillen, T., & Holmes, P. (2006). The dynamics of choice among multiple alternatives. *Journal of Mathematical Psychology*, 50, 30–57.
- Medina, L., & Reiner, A. (1995). Neurotransmitter organization and connectivity of the basal ganglia in vertebrates: Implications for the evolution of basal ganglia. *Brain Behav. Evol.*, 46(4–5), 235–258.
- Miller, E. K., & Cohen, J. D. (2001). An integrative theory of prefrontal cortex function. *Annu. Rev. Neurosci.*, 24, 167–202.
- Mink, J. W. (1996). The basal ganglia: Focused selection and inhibition of competing motor programs. *Prog. Neurobiol.*, 50(4), 381–425.
- Montague, P. R., Dayan, P., & Sejnowski, T. J. (1996). A framework for mesencephalic dopamine systems based on predictive Hebbian learning. *J. Neurosci.*, 16(5), 1936–1947.
- Nakano, K., Kayahara, T., Tsutsumi, T., & Ushiro, H. (2000). Neural circuits and functional organization of the striatum. *J. Neurol.*, 247 (Suppl 5), V1–15.
- Nambu, A., & Llinas, R. (1994). Electrophysiology of globus pallidus neurons in vitro. *J. Neurophysiol.*, 72(3), 1127–1139.
- Nambu, A., & Llinas, R. (1997). Morphology of globus pallidus neurons: Its correlation with electrophysiology in guinea pig brain slices. *J. Comp. Neurol.*, 377(1), 85–94.
- Obeso, J. A., Rodriguez-Oroz, M. C., Rodriguez, M., Lanciego, J. L., Artieda, J., Gonzalo, N., & Olanow, C. W. (2000). Pathophysiology of the basal ganglia in Parkinson's disease. *Trends Neurosci.*, 23(10 Suppl), S8–19.
- O'Doherty, J., Dayan, P., Schultz, J., Deichmann, R., Friston, K., & Dolan, R. J. (2004). Dissociable roles of ventral and dorsal striatum in instrumental conditioning. *Science*, 304(5669), 452–454.
- Okamoto, H., Isomura, Y., Takada, M., & Fukai, T. (2005). *Temporal integration by stochastic dynamics of a recurrent network of bistable neurons*. Paper presented at the Computational Cognitive Neuroscience meeting, Washington, DC.

- O'Reilly, R. C., & Frank, M. J. (2006). Making working memory work: A computational model of learning in the frontal cortex and basal ganglia. *Neural Computation*, 18, 283–328.
- Overton, P. G., & Greenfield, S. A. (1995). Determinants of neuronal firing pattern in the guinea-pig subthalamic nucleus: An in vivo and in vitro comparison. *J. Neural Transm. Park. Dis. Dement. Sect.*, 10(1), 41–54.
- Parent, A., & Hazrati, L. N. (1993). Anatomical aspects of information processing in primate basal ganglia. *Trends Neurosci.*, 16(3), 111–116.
- Parent, A., & Hazrati, L. N. (1995a). Functional anatomy of the basal ganglia. I. The cortico-basal ganglia-thalamo-cortical loop. *Brain Res. Rev.*, 20(1), 91–127.
- Parent, A., & Hazrati, L. N. (1995b). Functional anatomy of the basal ganglia. II. The place of subthalamic nucleus and external pallidum in basal ganglia circuitry. *Brain Res. Rev.*, 20(1), 128–154.
- Parent, A., & Smith, Y. (1987). Organization of efferent projections of the subthalamic nucleus in the squirrel monkey as revealed by retrograde labeling methods. *Brain Res.*, 436(2), 296–310.
- Parthasarathy, H. B., Schall, J. D., & Graybiel, A. M. (1992). Distributed but convergent ordering of corticostriatal projections: analysis of the frontal eye field and the supplementary eye field in the macaque monkey. *J. Neurosci.*, 12(11), 4468–4488.
- Platt, M. L., & Glimcher, P. W. (1999). Neural correlates of decision variables in parietal cortex. *Nature*, 400(6741), 233–238.
- Prescott, T. J., Montes-Gonzalez, F. M., Gurney, K., Humphries, M. D., & Redgrave, P. (2006). A robot model of the basal ganglia: Behavior and intrinsic processing. *Neural Networks*, 19(1), 31–61.
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychol. Rev.*, 83, 59–108.
- Ratcliff, R., Van Zandt, T., & McKoon, G. (1999). Connectionist and diffusion models of reaction time. *Psychol. Rev.*, 106, 261–300.
- Redgrave, P., Prescott, T. J., & Gurney, K. (1999). The basal ganglia: A vertebrate solution to the selection problem? *Neuroscience*, 89(4), 1009–1023.
- Samejima, K., Ueda, Y., Doya, K., & Kimura, M. (2005). Representation of action-specific reward values in the striatum. *Science*, 310(5752), 1337–1340.
- Schall, J. D. (2001). Neural basis of deciding, choosing and acting. *Nat. Rev. Neurosci.*, 2(1), 33–42.
- Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science*, 275(5306), 1593–1599.
- Shadlen, M. N., & Newsome, W. T. (2001). Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *J. Neurophysiol.*, 86(4), 1916–1936.
- Shadmehr, R., & Holcomb, H. H. (1997). Neural correlates of motor memory consolidation. *Science*, 277(5327), 821–825.
- Simen, P. A., Cohen, J. D., & Holmes, P. (2006). Rapid decision threshold modulation by reward rate in a neural network. *Neural Networks*, 19(8), 1013–1026.
- Smith, Y., Bevan, M. D., Shink, E., & Bolam, J. P. (1998). Microcircuitry of the direct and indirect pathways of the basal ganglia. *Neuroscience*, 86(2), 353–387.
- Stafford, T., & Gurney, K. (2004). The role of response mechanisms in determining reaction time performance: Pieron's law revisited. *Psychon. Bull. Rev.*, 11, 975–987.

- Stafford, T., & Gurney, K. (in press). The basal ganglia as the selection mechanism in a cognitive task. *Philos. Trans. R. Soc. Lond. B Biol. Sci.*
- Stewart, R. D., Gurney, K., & Humphries, M. D. (2005). *A large-scale model of the sensorimotor basal ganglia: Evoked responses and selection properties*. Paper presented at the Society for Neuroscience meeting, Washington, DC.
- Stone, M. (1960). Models for choice reaction time. *Psychometrika*, *25*, 251–260.
- Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning*. Cambridge, MA: MIT Press.
- Teichner, W. H., & Krebs, M. J. (1974). Laws of visual choice reaction time. *Psychol. Rev.*, *81*, 75–98.
- Usher, M., & McClelland, J. L. (2001). The time course of perceptual choice: The leaky, competing accumulator model. *Psychol. Rev.*, *108*(3), 550–592.
- Vickers, D. (1970). Evidence for an accumulator model of psychophysical discrimination. *Ergonomics*, *13*, 37–58.
- Wald, A. (1947). *Sequential analysis*. New York: Wiley.
- Wald, A., & Wolfowitz, J. (1948). Optimum character of the sequential probability ratio test. *Annals of Mathematical Statistics*, *19*, 326–339.
- Wang, X. J. (2002). Probabilistic decision making by slow reverberation in cortical circuits. *Neuron*, *36*(5), 955–968.
- Wilson, C. J. (1995). The contribution of cortical neurons to firing patterns of striatal spiny neurons. In J. C. Houk, J. L. Davis, & D. G. Beiser (Eds.), *Models of information processing in the basal ganglia* (pp. 22–50). Cambridge, MA: MIT Press.
- Wilson, C. J., Weyrick, A., Terman, D., Hallworth, N. E., & Bevan, M. D. (2004). A model of reverse spike frequency adaptation and repetitive firing of subthalamic nucleus neurons. *J. Neurophysiol.*, *91*(5), 1963–1980.
- Yu, A. D., & Dayan, P. (2005). Inference, attention, and decision in a Bayesian neural architecture. In L. K. Saul, W. Yair, & L. Bottou (Eds.), *Advances in neural information processing systems*, *17* (pp. 1577–1584). Cambridge, MA: MIT Press.