

The Brain Implements Optimal Decision Making between Alternative Actions

Rafal Bogacz (R.Bogacz@bristol.ac.uk)

Department of Computer Science, University of Bristol
Bristol, BS8 1UB, UK

Kevin Gurney (K.Gurney@sheffield.ac.uk)

Department of Psychology, University of Sheffield
Sheffield S10 2TP, UK

Abstract

Neurophysiological studies have identified a number of brain regions critically involved in solving the problem of ‘action selection’ or ‘decision making’ which has been extensively studied in cognitive psychology. 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 behavioural 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, and that can fully explain *why* they are organized in the specific way they are. This paper addresses this question by showing 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 Multiple Sequential Probability Ratio Test (MSPRT). The resulting model of basal ganglia provides a rationale for their inter-nucleus connectivity and the idiosyncratic properties of particular neuronal populations. The model is consistent with data regarding the reaction times in choice tasks, and provides a mechanistic explanation of how they are generated in neural circuits.

Keywords: basal ganglia, decision making, action selection

Introduction

Deciding which action to select is a common and critical element of human mental life, and hence, has been extensively studied in cognitive psychology. Within the last half-century, on the basis of careful analyses of reaction time (RT) data in tasks involving a choice between two or more alternative actions, various mathematical models have been proposed which assume that, during the decision process, noisy evidence supporting the alternative actions is accumulated and an action is executed as soon as certain criteria are met (Laming, 1968; Ratcliff, 1978; Stone, 1960; Usher & McClelland, 2001; Vickers, 1970).

Recent neurophysiologic studies have shed light on how the brain can implement the processes of accumulation and criteria evaluation suggested by these psychological models. In the case of highly practiced tasks, it has been shown that, during the decision process, neurons in cortical areas representing alternative actions gradually increase their firing rate, thereby accumulating evidence supporting these alternatives (Schall, 2001; Shadlen & Newsome, 2001). It is

also known that the Basal Ganglia (BG) may play a critical role in the selection of alternative actions. In their quiescent state, BG output nuclei supply tonic inhibition to midbrain and brain stem targets (Deniau & Chevalier, 1985) implicated in executing motor actions, thus blocking cortical control over these actions. Actions are supposed to be selected when neurons in the output nuclei have their activity reduced (under control of the rest of BG) thereby disinhibiting their targets (Deniau & Chevalier, 1985).

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 BG act as a central ‘switch’ by evaluating this evidence and enabling the behavioural request which is best supported (most salient). However, despite extensive experimental and numerical simulation studies demonstrating selective function (Brown, Bullock, & Grossberg, 2004; M. J. Frank, Seeberger, & O'Reilly R, 2004; Gurney, Prescott, & Redgrave, 2001), no formal theoretical framework exists that can fully explain *why* the BG are organized in the way they are. Multiple mechanisms supporting a selective function may be organised in a variety of ways; is there a rationale for the *specific* pattern of connectivity observed in BG?

This paper addresses this question and provides an analytic description of function of a circuit involving cortex and BG, by showing how an optimal abstract decision algorithm ‘maps’ onto the anatomy and physiology of this circuit.

Review of the neurobiology and theory of decision making

Modelling cortical integration

The neural basis of decision making in cortex has been studied extensively using single-cell recordings (Britten, Shadlen, Newsome, & Movshon, 1993; 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. 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). 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 oft-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 alternative i , 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 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 LIP and 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

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

over the temporal interval $[1, T]$. The accumulated evidence $Y_i(T)$ may now be used in making a decision about which x_i has the highest mean.

Modelling 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 cumulated 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) (Wald, 1947). The SPRT is optimal in the following sense: among all decision methods allowing a certain probability of error, it minimizes the decision time.

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 which 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 Multiple SPRT's (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).

BG connectivity

The BG connectivity used in our study contains the major pathways known to exist in BG anatomy and was based on that used in the model of Gurney et al. (2001). Fig. 1a shows this connectivity for rat in cartoon form; for reviews of BG anatomy see (Alexander, DeLong, & Strick, 1986; Gerfen & Wilson, 1996; Smith, Bevan, Shink, & Bolam, 1998); the description below is based on these reviews. Cortex sends excitatory projections to the striatum and subthalamic nucleus (STN). The striatum is divided into two populations of projection neurons differentiated, inter alia,




Fig. 1: Comparison of connectivity of basal ganglia and a network implementing the Multiple 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.

by their anatomical targets and preferential dopamine receptor type. The neurons in one striatal sub-population (associated with D1-type dopamine receptors) send focused inhibitory projections to the BG output nuclei – the substantia nigra pars reticulata (SNr) and entopeduncular nucleus (EP). Neurons in the other striatal population (associated with D2-type dopamine receptors) are also inhibitory and send focused projections to the globus pallidus (GP). 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. The output nuclei send widespread inhibitory connections to the mid-brain, brainstem, and the thalamus.

Neuronal selectivity in the BG

Studies of awake primates in behavioural tasks have established that all BG nuclei have somatotopic organisation. Further, within each of the nuclei, there are clusters of neurons responding selectively before and during movement of individual joints (often only in single direction) (Crutcher & DeLong, 1984; Georgopoulos, DeLong, & Crutcher, 1983). These observations led Alexander et al. (1986) to propose that “the motor circuit may be composed of multiple, parallel subcircuits or *channels* concerned with movement of individual body parts”, which traverse all nuclei of BG.

The notion of channels was incorporated into the computational model of Gurney et al. (2001) who proposed that each action is associated anatomically with a discrete neural population within each nucleus.

The BG implements selection using MSPRT

In this Section we introduce the MSPRT (Baum & Veeravalli, 1994) and demonstrate how it maps onto a neural network possessing a number of striking similarities to the anatomy and physiology of the BG.

The MSPRT

Consider a decision between N alternative actions, and denote evidence supporting the alternatives at time t by $x_1(t), x_2(t), \dots, x_N(t)$. Let the hypothesis H_i correspond to x_i having the highest mean. More precisely, we define H_i analogously to its definition for two alternatives (Gold & Shadlen, 2001; 2002), namely; H_i is the hypothesis that $x_i(t)$ come from independent 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 $\mu^+ > \mu^-$.

One of the MSPRT (Baum & Veeravalli, 1994) requires the following computations. In each time step, N quantities $y_i(T) = g * Y_i(T)$ need to be computed, where $Y_i(T)$ is the accumulated evidence supporting action alternative i and g^* is a constant. We will refer to $y_i(T)$ as the *saliency* of action i . In MSPRT at each time T the N variables $L_i(T)$ need to be computed:

$$L_i(T) = y_i(T) - \ln \sum_{k=1}^N \exp(y_k(T)) \quad (2)$$

MSPRT implies that an action should be selected as soon as any of the $L_i(T)$ exceeds a fixed decision threshold. Eq. 2 is the basis for mapping MSPRT onto the BG. But, before we proceed with this process, we describe some of the intuitive properties of MSPRT, as implemented in Eq. 2.

The right hand side of Eq. 2 includes two terms. The first term $y_i(T)$ is simply the saliency and, on its own, describes a race model (Vickers, 1970). This term therefore incorporates information about the *absolute size* of the saliency of the currently ‘winning’ alternative, w . The second term in Eq. 2 occurs in subsequent analysis and we denote it by $S(T)$ where

$$S(T) = \ln \sum_{i=1}^N \exp(y_i(T)) \quad (3)$$

$S(T)$ decreases the value of all $L_i(T)$ by the same amount, thereby increasing the minimum saliency required for an action to be selected. Its value is increased by the presence of more actions, and by individual actions having higher saliency. It may therefore be thought of as representing *response conflict*

Model of action selection in the BG

We now show how the test defined by Eq. 2 may be performed in a biologically constrained network model of BG. For simplicity of explanation here we show how Eq. 2 maps onto a model of BG including only subset of the known anatomical connections. We exclude the connections marked by dotted lines in Fig. 1a. It has been proposed that these pathways play a role in reinforcement learning (M. J. Frank et al., 2004), a function which is not included in our model, because we address only action selection in highly practiced tasks. However, incorporation of these pathways into an anatomically more complete scheme still admits a model of BG which supports MSPRT (Bogacz & Gurney, 2006) and hence achieves the optimal performance.

The mapping between Eq. 2 and the network is shown graphically in Fig. 1b. In our decomposition, each channel is associated with an action i and with 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 (localist) way.

We note first that the $L_i(T)$ are always negative, since $S(T) \geq \ln(\exp(y_i(T))) = y_i(T)$. Hence, 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 $-L_i(T)$:

$$OUT_i(T) = -y_i(T) + \ln \sum_{k=1}^N \exp(y_k(T)) \quad (4)$$

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

As described in the Introduction, we propose, along with others (Schall, 2001; Shadlen & Newsome, 2001), that quantities like $y_i(T)$, representing saliency, are computed in cortical regions which project to BG. In the motion discrimination example, $y_i(T)$ would be computed in FEF

which is known to innervate BG (Parthasarathy, Schall, & Graybiel, 1992).

Eq. 4 implies that the salience signals $y_i(T)$ (or nonlinear combinations therein) have to be distributed to the output so as to yield both excitatory and inhibitory contributions. Eq. 4 includes two terms and below we propose that the first is computed within the direct pathway from striatum to the output nuclei, while the second term within the pathway traversing STN and GP. The first term in Eq. 4, $-y_i(T)$, is an inhibitory component and cannot be supplied by cortex since its efferents are excitatory. We argue, therefore, that one function of the population of inhibitory striatal neurons with D1 receptors (see Fig. 1a) is to provide an ‘inhibitory copy’ of the salience signal to the output nuclei.

Turning to the second term in Eq. 4, this is $S(T)$ (defined in Eq. 3) which supplies an excitatory contribution to the output nuclei. Now, a key aspect of $S(T)$ is that it involves summing over channels. The source of excitation in BG is the STN which sends *diffuse* projections to the BG 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:

$$\Sigma(T) = \sum_{i=1}^N STN_i(T) \quad (5)$$

The model then implements MSPRT if $\Sigma(T) = S(T)$. We now describe the input-output relations of STN and GP neurons which give $\Sigma(T) = S(T)$. First, we require that the firing rate of neurons in STN is proportional to an exponential function of its inputs

$$STN_i(T) = \exp(y_i(T) - GP_i(T)) \quad (6)$$

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 firing rate of GP

channel i , $GP_i(T)$, is given by

$$GP_i(T) = \Sigma(T) - \ln(\Sigma(T)) \quad (7)$$

since, substituting Eq. 7 into Eq. 6, summing over i , and solving for $\Sigma(T)$ then yields $\Sigma(T) = S(T)$.

In summary, an implementation of MSPRT defined by Eq. 4-7 may be realised by a subset of BG anatomy, defined in Fig. 1b, if the behaviour of neurons in STN and GP follows Eq. 6 and 7.

Predicted requirements for STN and GP physiology are validated by existing data

In this Section we compare the predictions of Eq. 6 and 7, 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 forms for STN and GP functionality given in Eqs. 6 and 7 were derived on the basis of: (i) known anatomy of basal nuclei and (ii) the assumption that the network involving cortex and BG implements MSPRT. Since we did not *use* the physiological properties of STN and GP neurons in deriving Eq. 6 and 7, 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 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 and so the satisfaction of this property is very unlikely to be a result of a chance, and thus would provide a strong confirmation of the theory.

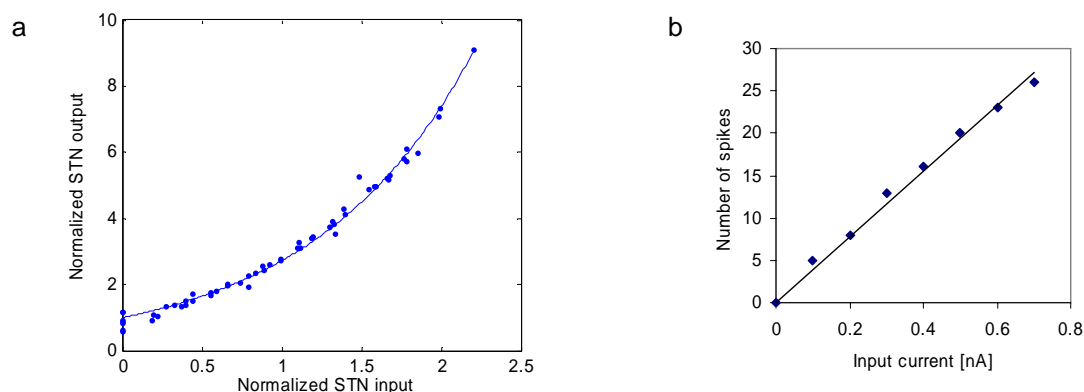


Fig. 2: Firing rates f of STN and GP neurons as a function of input current I . **a** STN neurons. The data are for seven neurons taken from studies (Hallworth et al.2003; Wilson et al.2004). For each neuron, a best fit function of the form $f = a \exp(b I)$ was found through data points (f_j, I_j) . The new data $(f_j/a, bI_j)$ were then plotted on the same axes for all neurons. **b** GP neurons. Number of spikes n produced by a GP neuron of type II (Nambu & Llinas, 1994) in a 242ms stimulation interval using current injection I . The data used in this figure were kindly provided by Atsushi Nambu, and they come from the same neuron which is analysed in Fig. 5g of (Nambu & Llinas, 1994).

The response properties of STN neurons have been studied extensively (Hallworth et al., 2003; Wilson et al., 2004). Typically they have, non-zero spontaneous firing, and can achieve unusually high firing rates. Our proposed exponential form for firing rate as a function of input (Eq. 6) explains these features since, in the absence of input, the model gives non-zero (unity) output and $\exp(\cdot)$ is a rapidly growing function yielding potentially high firing rates. In order to test the prediction of Eq. 6 quantitatively, we fitted exponential functions to firing rate data in the literature. Fig. 2a 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 Eq. 6.

Secondly, the theory makes predictions, defined by Eq. 7, concerning the firing rate of GP. First, we show that the function defined by Eq. 7 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 Eq. 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 Eq. 7 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 & 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 (Fig. 2b), a result which is in agreement with an earlier study by Kita & Kitai (1991).

Performance of MSPRT model of the BG

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). To do this, we conducted simulations for differing numbers of competing inputs, N , for all three models, with a 1% error rate (ER).

In the simulations, the evidence $x(t)$ was accumulated 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 $\sigma^2 \delta t$. Estimates of the parameters were taken from a sample participant in experiment 1 from the study of Bogacz et al. (submitted), i.e., $\mu^+ - \mu^- = 1.41$, $\sigma = 0.33$. For each set of parameters, a decision threshold was found numerically that resulted in ER of $1\% \pm 0.2\%$ (s.e.), and the decision time was then found for this threshold.

Fig. 3 shows that the MSPRT consistently achieves lower decision times than both the UM and race models (especially in the more realistic large N regime). This result is in agreement with recent work by McMillen & Holmes (in press) (who also showed another feature in Fig. 3 - 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




Fig. 3: Comparison of decision times of 3 models (see key) with thresholds giving error rate of 1% for different numbers of alternatives (shown on x-axis).

similar since, in this case, the latter approximates SPRT (Bogacz et al., submitted).

Discussion

Our main result is that a circuit involving cortex and the BG may be devoted to implementing a powerful (asymptotically optimal) decision mechanism (MSPRT). 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 BG appear to have evolved to support the more powerful MSPRT, allowing the brain to make accurate decision 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.

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 favour 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 Fig. 3) 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.

Acknowledgments

This work was supported by the EPSRC grants: EP/C516303/1 and EP/C514416/1.

References

- 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.
- Baum, C. W., & Veeravalli, V. V. (1994). A sequential procedure for multihypothesis testing. *IEEE Transactions on Information Theory*, 40, 1996-2007.
- Bogacz, R., Brown, E., Moehlis, J., Holmes, P., & Cohen, J. D. (submitted). The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced choice tasks. *Psychol Rev*.
- Bogacz, R., & Gurney, K. (2006). The basal ganglia and cortex implement optimal decision making between alternative actions. *Neural Computation*, submitted.
- 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, 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.
- Crutcher, M. D., & DeLong, M. R. (1984). 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.
- 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, 227-233.
- Dragalin, V. P., Tertakovsky, A. G., & Veeravalli, V. V. (1999). Multihypothesis sequential probability ratio tests – part I: asymptotic optimality. *IEEE Transactions on Information Theory*, 45, 2448-2461.
- 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).
- 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.
- Gurney, K., Prescott, T. J., & Redgrave, P. (2001). A computational model of action selection in the basal ganglia. *Biol Cybern*, 84(6), 401-410.
- 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.
- 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.
- Nambu, A., & Llinas, R. (1994). Electrophysiology of globus pallidus neurons in vitro. *J Neurophysiol*, 72(3), 1127-1139.
- 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 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 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.
- 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.
- Schall, J. D. (2001). Neural basis of deciding, choosing and acting. *Nat Rev Neurosci*, 2(1), 33-42.
- 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.
- 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.
- Stone, M. (1960). Models for choice reaction time. *Psychometrika*, 25, 251-260.
- 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.
- 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.