The basal ganglia as the selection mechanism in a cognitive task*

Tom Stafford & Kevin Gurney
Department of Psychology, University of Sheffield
Western Bank, Sheffield, S10 2TP, UK
t.stafford@shef.ac.uk

Abstract

This paper builds on our existing, biologically constrained, model of the basal ganglia, which was originally constructed under the premise that these subcortical structures perform action selection. Here we show how this same model, when used in conjunction with a connectionist model of processing in the Stroop task, can provide an improved account of human performance on that task. Our model accounts for a wide variety of phenomenon, and provides a framework for connecting Stroop processing with the neuroanatomical basis of action selection. This work validates modelling the basal ganglia as the vertebrate solution to the action selection problem and demonstrates the importance of action selection issues to understanding performance on cognitive tasks. Proposals are made concerning the desirable properties a selection mechanism must possess.

1 The basal ganglia as a vertebrate solution to the selection problem

> *'A selection problem arises whenever two or more competing systems seek simultaneous access to a restricted system.'* [Redgrave *et al.*, 1999]

It has been proposed that the basal ganglia is the vertebrate solution to the selection problem [Redgrave *et al.*, 1999]. In other words, that it resolves the competition between different neural command centres requesting behavioural control. This need for selection is most clear, and has been most thoroughly empirically explored, in terms of motor expression, but it is expected that similar functional architecture, and comparable functional requirements, underly selection in different domains.

The basal ganglia has external and internal connectivity that makes it suitable for performing the role of a selection mechanism. It receives inputs from virtually the entire cerebral cortex, limbic system structures such as the hippocampus and the amygdala, and, notably, the anterior cingulate cortex [Masterman and Cummings, 1997; Redgrave *et al.*, 1999].

*This work was supported in part by EPSRC grant EP/C516303/1 to Kevin Gurney.

The main input nucleus of the basal ganglia is the striatum, which provides the first processing of incoming signals from cortex. The projection neurons of the striatum (medium spiny neurons) are by default quiescent (in a down-state), and do not do not respond to low levels of input. Only after a substantial and coordinated excitatory input do they move to an up-state, in which they produce significant output which may subsequently be affected by smaller changes in input [Wilson, 1995].

Outputs from the basal ganglia project back to the cortex, via the thalamus, and to premotor areas of the brainstem. The output nucleus of the Basal ganglia is the globus pallidus, internal segment (GPi). Neurons here are are tonically active, inhibiting their target structures from enacting behaviour. Actions are enabled by the selective release of that inhibition. It is posited that signals from cortex indicate the 'salience' — i.e. the importance and urgency — of possible actions to the basal ganglia [Redgrave *et al.*, 1999]. Sufficiently large saliences result in the selective disinhibition of channels associated with that action, and thus the release of the action [Chevalier and Deniau, 1990].

1.1 Modelling confirms that the basal ganglia can perform action selection

We have constructed a neuronal network model of the basal ganglia, constrained by the known anatomy and physiology, and based on the selection hypothesis [Gurney *et al.*, 2001a; 2001b]. Analysis and simulation of this model [Gurney *et al.*, 2001b] shows that the basal ganglia display the properties of a good selection mechanism [Redgrave *et al.*, 1999]: the highest salience rapidly promotes appropriate channel selection; once selection has been made competitors do not distort that selection; however, significant changes in the salience inputs result in rapid and clean channel switching.

Embedding this model into its anatomical context provided by cortex and thalamic circuits, improves the selection behaviour and gives a more complete understanding of the functional role the different nuclei involved may be playing [Humphries and Gurney, 2002]. Further, using these models in robot controllers shows that the selection behaviour is of sufficient efficiency and sophistication to be behaviourally adequate in realistic environments [Girard *et al.*, 2003; Montes-Gonzalez *et al.*, 2000].

The simulation work presented below uses the basal gan-

glia model exactly as presented elsewhere [Gurney *et al.*, 2001b; Humphries and Gurney, 2002]. We focus on the benefits of using this biologically plausible mechanism, which has been demonstrated to possess ethologically realistic selection properties. The internal structure of the basal ganglia model is only discussed as far as is necessary to illustrate *why* it works as it does in the context of the current work.

1.2 Using the basal ganglia model in a cognitive task

This paper is concerned with a different extension of the model— into the domain of cognitive selection and performance, as measured by reaction times. In particular we consider a celebrated cognitive task that involves a selection conflict — the Stroop Task [Stroop, 1935] — and take as our starting point the most successful computational model of performance on this task to date [Cohen *et al.*, 1990] which we extend by integration with our existing model of basal ganglia function [Gurney *et al.*, 2001a; 2001b; Humphries and Gurney, 2002].

The purposes of this extension are threefold. Firstly, it allows an additional test of the basal ganglia model of action selection. The model was constructed using the known functional neuroanatomy and guided by the selection hypothesis. It was not explicitly designed to simulate reaction times, nor was it constrained by human cognitive performance. However, any action selection mechanism should also be able to act as a response selection mechanism in cognitive tasks. Therefore the performance of the model in this domain is a good test of its validity. Secondly, some aspects of human Stroop performance remain inadequately addresses by existing models, leaving the possibility open that a model containing new elements may improve the possible account and shed light on why previous models have not been so successful. Additionally, making connection to the possible underlying neurobiology enriches the account possible of Stroop processing. In particular, we anticipate that features of the basal ganglia model such as allowing arbitrary numbers of inputs and making provision for dopaminergic modulation of signal processing will provide opportunities for future experimental and modelling investigations. Thirdly, integrating cognitive and systems-neuroscience models sheds light on issues of selection from both levels of analysis. We will attempt to use our combined model to derive some general constraints on models of selection.

2 Modelling the Stroop Task

2.1 The Stroop Task

J. Ridley Stroop's famous task [Stroop, 1935] involves presenting words written in coloured inks. Participants must name the colour of the ink while trying to ignore the word, which can spell out the name of a colour. When the word-name is in contradiction to the ink-colour the task becomes effortful, slowed and error-prone. This is the interference effect, traditionally measured as the difference in reaction time (RT) or errors between the control condition (when the word-aspect of the stimuli is nominally neutral with respect to colour) and the conflict condition (when the word-aspect of the stimulus contradicts the color). There is a corresponding facilitation effect; when the word and the colour aspect match (the congruent condition) there is a speeding relative to the control condition. These two effects are asymmetrical; facilitation is typically far smaller than interference. The converse task — reading the word while ignoring the ink-colour — can also be assessed. Word-reading is faster than colour-naming, and is not affected by the colour-aspect of the stimulus (there is no interference or facilitation).

Traditionally the Stroop task has been discussed in terms of a conflict between automatic and controlled processes [MacLeod, 1991], and much progress has been made in using variations of the Stroop task to adumbrate the nature of 'automatic' processing [Besner and Stolz, 1999; Besner *et al.*, 1997; Dishon-Berkovits and Algom, 2000; Durgin, 2000]. But it is also apparent that the Stroop task involves a selection conflict and provides a thoroughly explored experimental framework for investigating cognitive aspects of selection.

2.2 A Model of the Stroop task

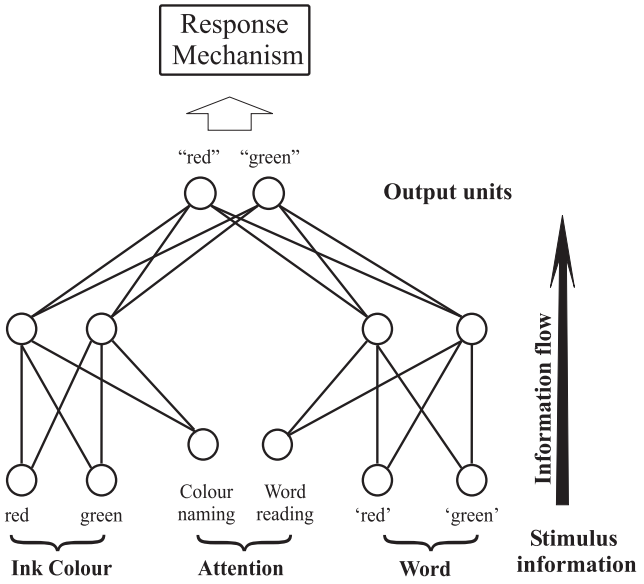

Figure 1: Architecture of the Cohen model

By far the most successful quantitative model of Stroop processing is that of Cohen *et al* [1990]. This simple connectionist model (hereafter 'The Cohen model') involves the translation of a localist input representation into a response representation, via a feed-forward two-layer network trained with backpropagation (Figure 1). The main features of this network are:

1. Differential training of the network: responses to word inputs are trained at ten times the frequency of responses to colour inputs. This results in a stronger weighting of signals representing this aspect of the stimulus.
2. Attentional sensitisation: the network implements attention as an additional input which off-sets a bias (in effect

a default inhibition) on all hidden units. This interacts with the sigmoidal output function of the units so that moderately sized signals do not result in a commensurate increase in output unless presented in combination with attentional input. Signals in the word-processing pathway, however, are large enough to partially overcome the default inhibition without the aid of attentional input.

3. Reaction times generated a by response mechanism that works on evidence accumulation: the output units of the network are taken to indicate, at each time point, the evidence favouring each response. This evidence is compared and accumulated until the total crosses a threshold – when a response is said to have been made. It is this feature feature of the model which is the concern of the current work.

Selection in Cohen *et al*'s (1990) model of the Stroop Task

The response mechanism of Cohen *et al*'s [1990] model is ignored in textbook treatments of the model [Ellis and Humphreys, 1999; Sharkey and Sharkey, 1995] and even overlooked in Cohen *et al*'s own analysis of the function of the model [Cohen *et al.*, 1990]. This reflects, we argue, a regrettable, but not untypical, neglect of the action selection problem in psychology. Reinforcing this view, we have recently, shown that, contrary to the original account of Cohen *et al*., it is the response mechanism, not the neuronal transfer function, which generates the important differences in reaction times between conditions [Stafford and Gurney, 2004; Stafford, 2003], and it is the response mechanism which explains the asymmetry in the magnitudes of the interference and facilitation effects in the Cohen model (a matter about which there has been some debate [MacLeod and MacDonald, 2000]). The response mechanism of the model is isomorphic to the diffusion model [Ratcliff, 1978; Ratcliff *et al.*, 1999], which has been shown to be an analytically tractable form of several connectionist models of decision, and an optimal decision algorithm for a two-choice decision situation [Bogacz *et al.*, submitted] where either desired accuracy or time-to-decision is specified (obviously these two mutually constrain each other). Further, potential neurobiological correspondences to the evidence accumulation processes of the diffusion model have been identified [Gold and Shadlen, 2000]. Thus our investigation of evidence accumulation as a mechanism of selection in this specific model may carry important lessons for theories of selection in general.

The response mechanism is also responsible for a major mismatch between model performance and human performance. A plausible alternative theory of Stroop processing — and of automatic processing in general — is that more automatic processes are simply faster. This theory would suggest that Stroop interference is due to the response evoked by the word aspect of the stimulus arriving at some response bottleneck earlier, creating slower selection of the opposite response when it arrives there. The experimental refutation of this theory involves presenting a coloured-ink patch next to a colour-word. If presented simultaneously the normal Stroop effect is found, but the spatial separation allows the asynchronous presentation of the colour and the word; a stimulus onset asynchrony (SOA) paradigm. If the colour appears sufficiently before the word then, according to the simple 'horse-race' theory, naming of the word should suffer interference from the colour information (a 'reverse Stroop effect'). This is not what happens experimentally [Glaser and Glaser, 1982]. For word-naming, no amount of head-start for colour-information is sufficient to create interference. For colour-naming, the appearance of the word at any point up to 300 ms after the appearance of the colour (close to the asymptotic limit for reaction times) causes interference. Additionally, the appearance of the word before the colour always causes interference, however long the subject is given to accommodate to the presence of the word.

The Cohen *et al* model can simulate limited features of the Stroop SOA paradigm. However, if the model is tested beyond the range presented in the original paper, serious flaws are revealed. Trends in the simulation data which can be seen over the original range of SOA values continue at longer SOAs, as the to-be-ignored dimension of the stimulus is presented increasingly before the to-be-responded to dimension (see Figure 2). By convention SOAs which involve the to-be-ignored dimension being presented first are labelled negative. Thus, for colour naming, in the conflict condition, the model response time increases as the SOA gets more negative until eventually the word-aspect is presented early enough to force an incorrect response. In Figure 2 this is represented by the peak in the line showing the colour-naming conflict condition reaction times. RTs start to decrease with increasing negative SOA because the model is more and more quickly selecting the wrong response. If the word is congruent to the colour information then there is comparable interference, but this reveals itself as a speeding of the correct response; this dynamic continues until, ultimately, the model responds before the colour information has even been presented. In Figure 2 this shown by the point at which the line representing RTs for colour-naming congruent condition crosses the dotted line representing zero on the RT axis. For the same fundamental reasons, in the word-naming task the conflict and congruent conditions diverge in the same way (albeit over a longer time span). Thus, the model behaves in accordance with the experimentally disproved horse-race model: presenting colour information ahead of word information creates a reverse Stroop effect — colour information interferes with word-reading.

The reason for these failures may be traced to the evidence accumulation response mechanism. Because the model, like all connectionist models, works on graded signals there is always some signal change due to the to-be-ignored, even if this is very small due to the attentional inhibition. In the case of the colour-naming task, it is integral to the model's function that some influence of the word-aspect of the stimulus survives attentional selection and comes to influence the response stage. Without this feature the basic effect of Stroop interference would not be present. However, in SOA conditions, this influence of the to-be-ignored aspect may accumulate indefinitely. This affects selection time to an extent proportional to the time it is presented multiplied by the strength of evidence conveyed. So arbitrarily small amounts of evidence can provoke erroneous selection if presented for long

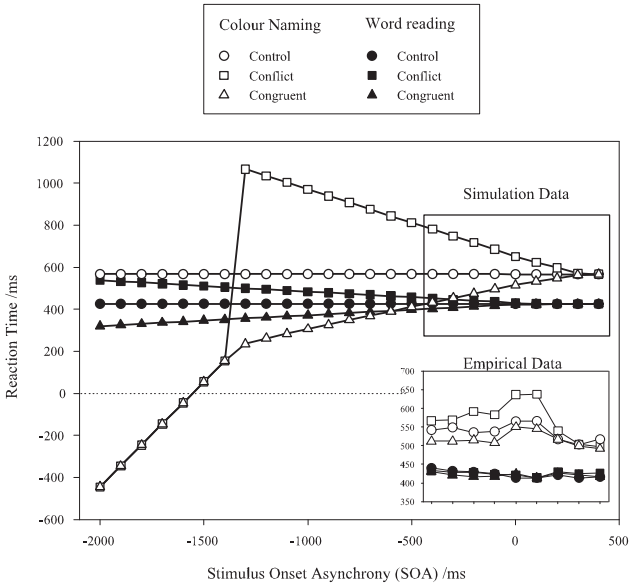

Figure 2: The SOA simulation of the original Cohen model. The empirical data is shown inset. The simulation data corresponds roughly to the empirical data over the range originally reported (-400 to +400 ms) but beyond that diverges.

enough, or they can massively slow correct selection (because accumulated evidence for the opposite response must be overcome).

Adding a more biologically realistic response mechanism — based on the basal ganglia — overcomes these deficiencies and considerably extends the model's explanatory power.

3 The Basal Ganglia model as a response mechanism for a cognitive task

The neural network component of Cohen *et al*'s model performs what is normally thought of as the cognitive elements of the task: stimulus–response translation, attentional control and learning. Only one minor change was required to this 'front-end' to make is compatible with using the basal ganglia model as the response mechanism. The output units of the Cohen model originally had resting values of 0.5. This was changed to 0.1, to make the output signals interpretable by the basal ganglia model as indicative of the salience of the corresponding response [1].

In all other respects the combined model is exactly as published by Cohen *et al* [1990], except with the basal ganglia model [Gurney *et al.*, 2001b; Humphries and Gurney, 2002] replacing evidence accumulation as the method of final response selection.

[1]For consistency this entails changes in the initial weights the networks is given before training, but these are not discussed here as there is no substantive effect on the simulation results; as should be expected from a good model the principle findings are robust under parametric variation, and this aspect of the model is an implementational detail which is irrelevant to overall behaviour of the model. For details see Stafford [2003].

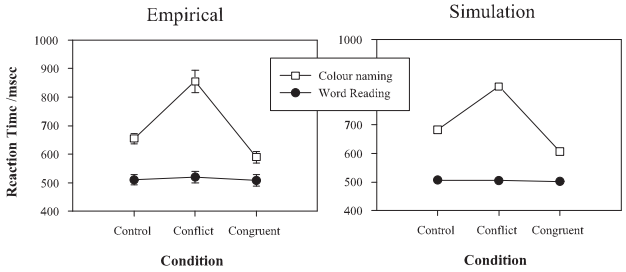

Figure 3: Empirical and simulation reaction times in the basic Stroop conditions for word-reading and colour-naming tasks. Empirical data is from Dunbar & MacLeod [1984], for which standard error bars are shown

4 The Simulation Results

4.1 Matching and Improving on the performance of Cohen's Model

We tested the Cohen connectionist front-end with the basal ganglia model as the response mechanism (hereafter 'The Model') on the first three simulations presented by Cohen *et al* [1990]. The model simulates the basic Stroop task (simulation 1), matching the empirical data as well as the original Cohen model does (Figure 3).

The ability to realistically model learning phenomena is a key benefit of connectionist models. The model mimics the power-law function of learning (Figure 4), just as the original Cohen model does. This demonstrates that the learning dynamic captured by the connectionist front-end is not interfered with by the use of the basal ganglia response mechanism; graded changes in the signals from the front end are converted into appropriately graded changes in reaction times.

The SOA task (Simulation 2) shows up the superiority of the basal ganglia as a response mechanism over the original response mechanism. As discussed, over long negative SOAs the Cohen response mechanism makes wrong selections, due to the small but significant influence of the distracting stimulus dimension. The input units of the basal ganglia model filter out small salience inputs (as discussed section 1). This creates a minimal salience threshold, below which inputs are ignored. Thus, using the basal ganglia response mechanism, the model makes the correct selection at all SOA values. Furthermore, the distracting influence of the to-be-ignored aspect of the stimulus is limited. This is reflected in the stabilisation of reaction times at SOAs below –400 ms (see Figure 5).

5 General Discussion

5.1 Strengths of the Model

This work validates our model against the basic Stroop phenomena. Use of the basal ganglia model as the response mechanism improves the fit that can be made to the empirical data and highlights necessary features response mecha-

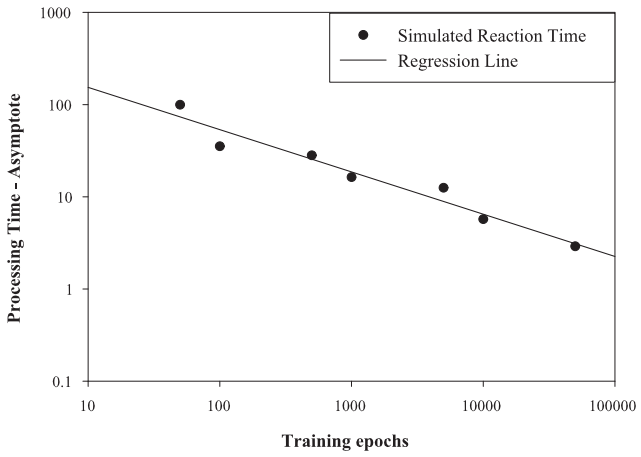

Figure 4: The model conforms to the power law of practice [Logan, 1988]. Both axis use a log scale. Simulation results are shown as dots. The simple regression for the data is shown as a straight line and follows the form log_{10}(Processing Time) $= 2.65 - 0.46 \times log_{10}$(Epochs). $R^2 = 0.948$.

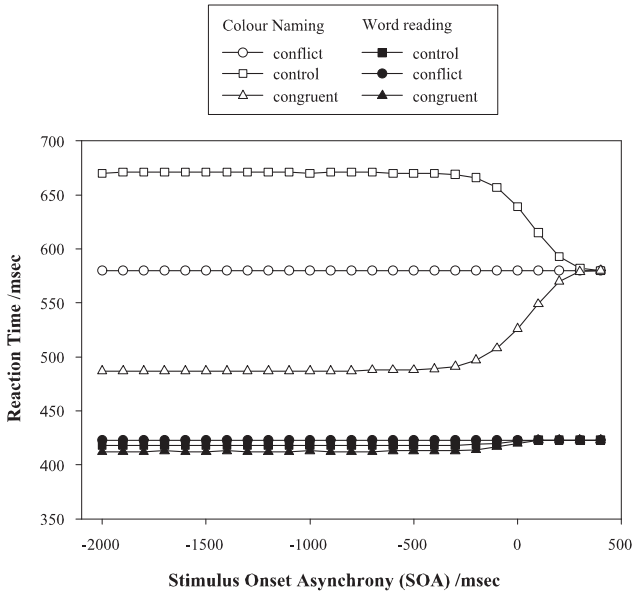

Figure 5: Model SOA data.

nisms should contain, the lack of which was overlooked in the previous account. Use of the basal ganglia model also extends the account of Stroop processing to connect with the neurobiology of selection. The basal ganglia model includes anatomical specific pathways and an account of the dopamine system. This allows future tests of the model against various pathologies, such as schizophrenia.

A better account of the data is one benefit of this model. There is also a theoretical purity to testing models by utilising them in new areas that they were not developed with in mind. It is testament to the basal ganglia model's value as a general model of selection that it deals appropriately with signals provided by a connections model of a cognitive task.

5.2 Why Does The Model Work?

The model captures the basic Stroop (Figure 3) and learning (Figure 4) phenomena because, for moderately sized saliences, selection time is based on the relative difference between the to-be-selected salience and the competing salience (if any). It is with small saliences, and when dealing with successive rather than simultaneous inputs, that the basal ganglia model shows its superiority as a selection mechanism. Both of these cases are revealed by comparison of the SOA simulations (Figures 2 and 5).

The failure of the Cohen model on the SOA simulations is because of a model feature which is neither trivial nor irrelevant. The existence empirically of the basic Stroop interference effect demonstrates that response activation from the to-be-ignored word aspect of the stimulus must break through any initial attentional inhibition. Arriving at the response mechanism before the response activation of the colour aspect, this activity is enough, in Cohen's model, to cause selection. The erroneous selection produced at long SOAs shows that a response mechanism must not make selections based on inconsequentially low inputs. Our basal ganglia model avoids this by having a minimum salience threshold, below which no action is selected.

This minimal threshold was included in the basal ganglia model because of the neurobiology of medium spiny neurons in the striatum – the main input nucleus of the basal ganglia. These neurons possess upstate / downstate functionality, which means that they only start to release action potentials if their input is above a certain threshold. This feature has the effect of filtering out noise in the inputs which is below threshold. The Cohen model evidence accumulation mechanism has no such minimal threshold, and no decay of accumulated evidence, and because of this it always makes a selection if left for long enough. By extension, the diffusion model, the general form of the evidence accumulation mechanism used, contains no capacity for not making a selection. This is a serious flaw. It means that evidence accumulation and the diffusion model alone cannot provide a full account of action selection.

In the basal ganglia model, competition on other channels, even if below selection level, can affect selection time. This priming, whether positive or negative, occurs because activity on other channels alters the resting level of output signals in GPi, and thereby affects the time it takes for outputs to drop to the point whereby selection occurs. The amount of this priming is limited because the model uses units with an output range restricted between 0 and 1. Compare this with the Cohen response mechanism, and by extension the diffusion model, which, with no constraints on where the selection threshold is set, contains the capacity to retain infinitely large values and thus can generate arbitrarily large amount of interference (as seen in the SOA simulations, Figure 2). This benefit of the basal ganglia model demonstrates the value of considering the mechanisms of action selection within a (neural) signal processing context.

5.3 What properties must a selection mechanism possess?

At a minimum these issues indicate that the context within which the diffusion model of selection is used cannot be ignored or assumed. The simulation of the SOA paradigm highlights two properties which the basal ganglia as a selection mechanism brings to the combined model to improve the account of the data. Together both of these features mean that not only is the wrong response not selected, but also the right response is selected efficiently. This is an example of the 'clean switching' property which has been identified as a necessary feature of any selection mechanism [Redgrave *et al.*, 1999].

The first feature is that the basal ganglia model limits the maximum possible influence on selection of concurrently or consecutively active competing inputs. So, in the SOA paradigm with negative SOAs the interference on reaction time does not get progressively longer with increasing SOA, but instead levels off – there is a maximum amount of interference that a distracting stimulus can produce on reaction times. This benefit is due to the wider context of adaptive control that the basal ganglia model arose from. A response mechanism needs to work in real-time, continuously, dealing with the successive selection of actions and interruption of old actions by new. Because the BG model is designed to operate continuously it has equilibrium final states, in which no action is selected. All patterns of input, if unchanging, eventually produce unchanging output states (although such a situation is unlikely to arise). For some patterns of input, the final output state indicates that no action is selected. The evidence accumulation response mechanism, on the other hand, has only one type of final state – that of selecting an action – and it continuously moves towards this state. The existence of equilibrium final states allows the successive switching between actions, without those actions interfering more with the selection of new actions the longer they have been selected.

The second necessary feature is that the basal ganglia will not make selections based on arbitrarily low inputs. This is because, due to the physiological properties of the medium spiny neurons in the striatum, the input nucleus of the basal ganglia, the model has built into it a minimum input threshold below which signals are ignored. Without such an input threshold, any level of input will cause the evidence accumulation counter to inextricably increase towards the selection threshold. A minimum input threshold is not the only way of preventing this kind of erroneous selection. Usher & McClelland [2001], in their model of perceptual choice, present an alternative strategy to a minimal input threshold, but one which has the same functional role. They argue that models of perceptual choice – they discuss the same kind of choice algorithms that are the basis for the Cohen *et al* [1990] response mechanism [Luce, 1986] – require the addition of activation decay on the choice representations. A decay mechanism can fulfill the same role as a minimal input threshold, since for situations where input is less than the decay that input is effectively filtered out. Another way of solving this erroneous selection problem might be to send a no-go signal which prevents selection until appropriate. This would not be feasible with a evidence accumulation model of selection, but it would be feasible with the basal ganglia response mechanism because it does not allow previous signal values to carry potentially unlimited weight when selecting new actions (i.e. clean switching, as discussed above). In the SOA task a no-go signal could be provided by the front-end to the basal ganglia on a third channel. This no-go signal, by being itself selected, can prevent selection until the relevant stimulus dimension has appeared. Although possible, this type of solution is perhaps not theoretically desirable because it relegates the problem of selection to another part of the system and hence begs the question of how correct selection is achieved.

A third possible way of accounting for the basic Stroop effect but avoiding erroneous selection in the SOA conditions is to include in the model a kind of reactive attentional inhibition, which suppresses activity based on the to-be-ignored dimension but only after it has occurred. Just such a stimulus-evoked inhibition mechanism is the focus of the cognitive control hypothesis of Botvinick *et al* [2001]. Initial investigations suggest that this mechanism, because it reduces interference from the to-be-ignored dimension of the Stroop stimulus but not until that interference has first arisen, would allow the accurate modelling of the course of interference in the SOA paradigm [Stafford, 2003]. Future modelling work may suggest ways in which these ways of limiting interference and preventing selection based on arbitrarily small values can be experimentally distinguished.

References

[Besner and Stolz, 1999] D. Besner and J. Stolz. Context dependency in stroop's paradigm: When are words treated as nonlinguistic objects? *Canadian Journal of Experimental Psychology*, 53(4):374–380, 1999.

[Besner *et al.*, 1997] D. Besner, J. A. Stolz, and C. Boutilier. The stroop effect and the myth of automaticity. *Psychonomic Bulletin & Review*, 4(2):221–225, 1997.

[Bogacz *et al.*, submitted] R. Bogacz, E. Brown, J. Moehlis, P. Hu, P. Holmes, and J.D. Cohen. The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced choice tasks. submitted.

[Botvinick *et al.*, 2001] M. M. Botvinick, T. S. Braver, D. M. Barch, C. S. Carter, and J. D. Cohen. Conflict monitoring and cognitive control. *Psychological Review*, 108(3):624–652, 2001.

[Chevalier and Deniau, 1990] G. Chevalier and J.M. Deniau. Disinhibition as a basic process in the expression of striatal functions. *Trends in Neurosciences.*, 13:277–281, 1990.

[Cohen *et al.*, 1990] J. D. Cohen, K. Dunbar, and J. L. McClelland. On the control of automatic processes - a parallel distributed-processing account of the stroop effect. *Psychological Review*, 97(3):332–361, 1990.

[Dishon-Berkovits and Algom, 2000] M. Dishon-Berkovits and D. Algom. The stroop effect: It is not the robust phenomenon that you have thought it to be. *Memory & Cognition*, 28(8):1437–1449, 2000.

[Dunbar and MacLeod, 1984] K. Dunbar and C. M. MacLeod. A horse race of a different color - stroop interference patterns with transformed words. *Journal of Experimental Psychology - Human Perception and Performance*, 10(5):623–639, 1984.

[Durgin, 2000] F. H. Durgin. The reverse stroop effect. *Psychonomic Bulletin & Review*, 7(1):121–125, 2000.

[Ellis and Humphreys, 1999] R. Ellis and G. Humphreys. *Connectionist Psychology : a text with readings*. Psychology Press Ltd, Hove, UK, 1999.

[Girard *et al.*, 2003] B. Girard, V. Cuzin, A. Guillot, K. N. Gurney, and T. J. Prescott. A basal ganglia inspired model of action selection evaluated in a robotic survival task. *Journal of Integrative Neuroscience*, 2(2):179–200, 2003. Journal version of the SAB paper.

[Glaser and Glaser, 1982] M. O. Glaser and W. R. Glaser. Time course analysis of the stroop phenomenon. *Journal of Experimental Psychology-Human Perception and Performance*, 8(6):875–894, 1982.

[Gold and Shadlen, 2000] J. I. Gold and M. N. Shadlen. Representation of a perceptual decision in developing oculomotor commands. *Nature*, 404(6776):390–394, 2000.

[Gurney *et al.*, 2001a] K. Gurney, T.J. Prescott, and P. Redgrave. A computational model of action selection in the basal ganglia i: A new functional anatomy. *Biological cybernetics*, 85(6):401–410, 2001.

[Gurney *et al.*, 2001b] K. Gurney, T.J. Prescott, and P. Redgrave. A computational model of action selection in the basal ganglia: Ii: Analysis and simulation of behaviour. *Biological cybernetics*, 85(6):411–423, 2001.

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

[Logan, 1988] G.D. Logan. Toward an instance theory of automatization. *Psychological Review*, 95(4):492–527, 1988.

[Luce, 1986] R.D. Luce. *Response Times: Their Role in Inferring Elementary Mental Organisation*. Oxford Psychology Series. Clarendon Press, New York, 1986.

[MacLeod and MacDonald, 2000] C.M. MacLeod and P.A. MacDonald. Interdimensional interference in the stroop effect: uncovering the cognitive and neural anatomy of attention. *Trends in Cognitive Sciences*, 4(10):383–391, 2000.

[MacLeod, 1991] C. M. MacLeod. Half a century of research on the stroop effect - an integrative review. *Psychological Bulletin*, 109(2):163–203, 1991.

[Masterman and Cummings, 1997] D. L. Masterman and J. L. Cummings. Frontal-subcortical circuits: The anatomic basis of executive, social and motivated behaviors. *Journal of Psychopharmacology*, 11(2):107–114, 1997.

[Montes-Gonzalez *et al.*, 2000] F. M. Montes-Gonzalez, T. J. Prescott, K. Gurney, and P. Redgrave. The robot basal ganglia: Control of robot action selection by an embodied model of the mammalian basal ganglia. *European Journal of Neuroscience*, 12:134–134, 2000.

[Ratcliff *et al.*, 1999] R. Ratcliff, T. Van Zandt, and G. McKoon. Connectionist and diffusion models of reaction time. *Psychological Review*, 106(2):261–300, 1999.

[Ratcliff, 1978] R. Ratcliff. A theory of memory retrieval. *Psychological Review*, 85:59–108, 1978.

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

[Sharkey and Sharkey, 1995] A.J.C. Sharkey and N.E. Sharkey. Cognitive modeling: psychology and connectionism. In M.A. Arbib, editor, *The Handbook of Brain Theory and Neural Networks*, pages 200–203. The MIT Press, Cambridge, MA., 1995.

[Stafford and Gurney, 2004] T. Stafford and K.N. Gurney. The role of response mechanisms in determining reaction time performance: Pieron's law revisited. *Psychonomic Bulletin & Review*, 11(6):975–987, 2004.

[Stafford, 2003] T. Stafford. *Integrating Psychological and Neuroscientific Constraints in Models of Stroop Processing and Action Selection*. Phd thesis, University of Sheffield. Available at http://www.shef.ac.uk/ ∼abrg, 2003.

[Stroop, 1935] J.R. Stroop. Studies of interference in serial verbal reactions. *Journal of Experimental Psychology*, 18:643–662, 1935.

[Usher and McClelland, 2001] M. Usher and J. L. McClelland. The time course of perceptual choice: The leaky, competing accumulator model. *Psychological Review*, 108(3):550–592, 2001.

[Wilson, 1995] C. Wilson. The contribution of cortical neurons to the firing pattern of striatal spiny neurons. In J. Houk, J. Davis, and D. Beiser, editors, *Models of information processing in the basal ganglia*, pages 29–50. MIT Press, Cambridge, MA., 1995.