

Dynamic Behaviour of a Spiking Model of Action Selection in the Basal Ganglia

Terrence C. Stewart (tcstewar@uwaterloo.ca)

Xuan Choo (fchoo@uwaterloo.ca)

Chris Eliasmith (celiasmith@uwaterloo.ca)

Centre for Theoretical Neuroscience, University of Waterloo
Waterloo, ON, N2L 3G1

Abstract

A fundamental process for cognition is action selection: choosing a particular action out of the many possible actions available. This process is widely believed to involve the basal ganglia, and we present here a model of action selection that uses spiking neurons and is in accordance with the connectivity and neuron types found in this area. Since the parameters of the model are set by neurological data, we can produce timing predictions for different action selection situations without requiring parameter tweaking. Our results show that, while an action can be selected in 14 milliseconds (or longer for actions with similar utilities), it requires 34-44 milliseconds to go from one simple action to the next. For complex actions (whose effect involves routing information between cortical areas), 59-73 milliseconds are needed. This suggests a change to the standard cognitive modelling approach of requiring 50 milliseconds for all types of actions.

Keywords: action selection; basal ganglia; spiking neurons; Neural Engineering Framework; cognitive cycle time

Action Selection

The basal ganglia are generally believed by both neuroscientists (e.g. Redgrave et al., 1999) and cognitive scientists (e.g. Anderson et al., 2004) to be responsible for action selection. Action selection consists of choosing one action to perform out of the many actions in an organism's repertoire. Selection is done on the basis of some sort of context-dependent utility signal for each possible action. Actions that are inappropriate for the current context may have low utility, and a task of the basal ganglia is to select the action that currently has the highest utility value.

Since such a mechanism forms the core of many cognitive models, including all of those based on production systems (where a single production much be chosen to fire), it is useful to develop a computational model of this process. Here, we develop a detailed spiking neuron model that takes into account a broad range of neurological details about the basal ganglia. Other spiking models of action selection exist, but tend to be organized unlike the basal ganglia (Belavkin & Huyck, 2009) and unconstrained by neural properties (Shouno et al., 2009; see Humphries et al., 2006 for an exception and alternate approach).

By directly connecting our model to neuroscientific results, we constrain our parameter values. Every parameter in the model reflects neurological data from the relevant brain areas, resulting in a model that has no free parameters (that affect the results shown here). Furthermore, having a biologically realistic model allows us to make predictions about a wide range of measures, including spike patterns, timing, variability, lesion effects, neural degeneration, the

influence of various drugs, and so on. Importantly, all of these predictions can come from the same model, with no additional parameters.

Neural Structure

The basal ganglia are a group of subcortical structures that are ideally suited for an action selection operation, as they receive input from extremely broad areas of cortex and the limbic system, and send output back to these areas via the thalamus. The basic components are the striatum, the subthalamic nucleus (STN), the globus pallidus internal (GPi), the globus pallidus external (GPe), and the substantia nigra pars reticulata (SNr).

The classic way of thinking about the organization of the basal ganglia is shown in Figure 1A. It consist of a direct pathway, where excitatory inputs from cortex to the D1 cells in the striatum inhibit corresponding areas in GPi and SNr, which then in turn inhibit areas in the thalamus, and an indirect pathway from the D2 cells in the striatum to GPe, STN, and then GPi/SNr (Albin et al., 1989). However, more recent evidence shows other major connections, including a hyperdirect excitatory pathway straight from cortex to STN (Nambu et al., 2002), and other feedback connections, as shown in Figure 1B.




Figure 1: Two schematic diagrams of the basal ganglia. **A** shows the standard direct/indirect pathway. **B** includes the other major connections that have been discovered.

There is also a great deal of topological structure in the inhibitory connections in basal ganglia. Neurons in the striatum project to a relatively localized area in the GPi, GPe, and SNr, while the excitatory connections from STN are very broad (Mink, 1996). This is an important constraint for the model we discuss below.

Simple Action Selection Models

Two simple approaches to neurally modelling action selection are shown in Figure 2. The inputs give the utilities of three possible actions (0.3, 0.8, and 0.5), and the model's task is to choose one of them. Importantly, since the output from the basal ganglia is inhibitory, selecting an action consists of having that particular inhibitory output be zero. In other words, it will no longer inhibit the neurons to which it is connected, allowing the action to occur. Thus, in Figure 2, the selected action is the middle one, whose output value is zero in both cases.

The model in Figure 2A is the most straight-forward. Each input neuron inhibits its corresponding output neuron and excites all others. For the first action, this results in an output of $-0.5 \cdot 0.3 + 0.5 \cdot 0.8 + 0.5 \cdot 0.5 = 0.5$. The action with the largest input will have the smallest output, and if the weights are in suitable ranges, only one output neuron will be turned off. One problem with this approach is determining suitable weights, although this can be helped by introducing recurrent connections, as in our earlier model (Stewart & Eliasmith, 2009). However, a more fundamental problem is that real neurons are typically either excitatory or inhibitory, and seldom both, as they are in this model.

An alternate approach is shown in Figure 2B. Here, instead of each neuron being both excitatory and inhibitory, a separate inhibitory interneuron is introduced. These are found throughout the brain, and can be used here to divide up the excitatory and inhibitory parts of the task. This approach is commonly used in neural models of action selection (e.g. Hazy et al., 2007; Stocco et al., 2010).

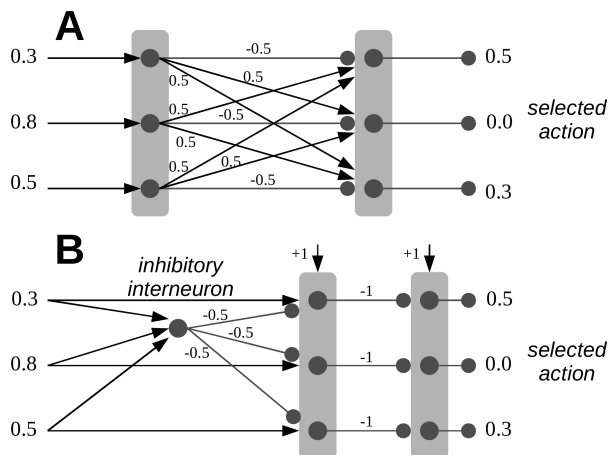


Figure 2: Two simple models of action selection. Inputs are the utilities of three possible actions, and an output of zero indicates the selection of a particular action. Each neuron (circle) outputs the sum of its weighted inputs.

A Realistic Rate Neuron Model

Gurney, Prescott, and Redgrave (2001) have developed a computational model of the basal ganglia that is well-suited to reimplementation using more realistic spiking neurons. While their model uses rate neurons, they have carefully followed the known biological constraints on the connectivity and types of neurons in the basal ganglia.

One of the main differences between their model and other computational models (e.g. Hazy et al., 2007; Stocco et al., in press) is that it does not make use of inhibitory interneurons in the striatum to perform action selection (as in Figure 2B). This is important for two reasons. First, while the striatum does include inhibitory interneurons, the actual behaviour and biological characteristics of these neurons is unclear, making them difficult to model. Second, there seems to be little evidence of the sort of broad, diffuse connectivity required by figure 2B (Gurney, et al., 2001). Tepper and Bolam (2004) identify three different types of striatal interneurons, and demonstrate their ability to affect spike timing in the rest of the striatum. These interneurons are highly influenced by dopamine (Bracci et al., 2002), acetylcholine (Koos & Tepper, 2002), and serotonin (Blomeley & Bracci, 2009), indicating that their role may be more to do with learning and other large-scale cognitive processes than with action selection.

Instead, Gurney, Prescott, and Redgrave (2001) present a model where the inhibitory output from the striatum and the excitatory output from the subthalamic nucleus (STN) combine to produce the desired output. That is, instead of treating the striatum as the primary input to the basal ganglia, neurological evidence shows that the STN receives excitatory connections directly from the cortex, and then produces diffuse excitation in the output nuclei. Figure 3 shows how this leads to an action selection mechanism that separates the inhibitory and excitatory connections.




Figure 3: Action selection via the striatum D1 cells and the subthalamic nucleus (STN). Connections from the STN are all excitatory and set at a weight of 0.5. The input with the highest utility (0.8) causes the corresponding output in the globus pallidus internal (GPi) or substantia nigra (SNr) to drop to zero, stopping the inhibition of that action.

While the model shown in Figure 3 is sufficient for action selection in some circumstances, it turns out not to be fully general. In particular, it has difficulty adjusting to situations where there are many actions with large utilities or where all actions have low utilities. For this reason, a control system is needed to modulate the behaviour of these neural groups. Gurney et al. (2001) argue that the globus pallidus external (GPe) is ideally suited for this, as its only outputs are back to the other areas of the basal ganglia, and it receives similar inputs from the striatum and the STN as does the globus pallidus internal (GPi). In their model, the GPe forms a circuit identical to that in Figure 3, but its outputs project back to the STN and the GPi. This regulates the action selection system, allowing it to function across a full range of utility values. The final network is shown in Figure 4.




Figure 4: The model of action selection in the basal ganglia presented by Gurney, Prescott, and Redgrave (2001). The striatum D1 cells and the subthalamic nucleus (STN) are as in Figure 3, while the striatum D2 cells and globus pallidus external form a modulatory control structure.

Converting Rates to Spikes

The model discussed so far is capable of performing action selection and reproducing a variety of single-cell recording results from electrostimulation and lesion studies (Gurney et al., 2001). However, it does so with rate neurons; that is, the neurons do not spike and instead continually output a numerical value based on their recent input. This makes it difficult to make precise numerical timing predictions or to make use of more accurate neural models. Furthermore, the model has no redundancy, since exactly one neuron is used per area of the basal ganglia to represent each action. The model shown in Figure 4 uses a total of 15 neurons (dark circles) to represent 3 possible actions, and if any one of those neurons is removed the model will fail.

To make timing predictions and to constrain our model with a broader range of neurological details, we needed to adapt the rate model of the basal ganglia into one that uses spiking neurons. For the results shown here, we use the standard leaky integrate-and-fire (LIF) model of spiking neuron behaviour, although our initial results with a more detailed implementation of the medium spiny neurons in the striatum (Gruber et al., 2002) are similar.

For LIF neurons, current is constantly leaking out of the neuron as per the membrane resistance R . If enough input current is gathered to cause the voltage to be above a certain threshold, then the neuron will fire. After firing, the voltage is set to 0 for a fixed refractory period (~ 2 milliseconds) before starting to gather current again. Given a constant current input J and membrane resistance R , the voltage level of the LIF neuron changes over time as given in Equation 1 and shown in Figure 1. The timing of this behaviour is controlled by τ_{RC} , the membrane time constant of the neuron.

$$V(t) = JR(1 - e^{-t/\tau_{RC}}) \quad (1)$$




Figure 5: LIF neuron with constant input current.

For a constant input, we can measure the average firing rate of a given LIF neuron, and this will be dependent solely on the neurophysiological details of the resistance R and the membrane time constant, which tend to be fixed for any particular type of neuron. However, for real *in vivo* neurons, their output will also vary based on any background current flowing into the neurons, and their activity can be scaled by the strength of the incoming synaptic connection. Thus, even among neurons of the same type, their responses will vary, as shown in Figure 6A. The behaviour of a neuron as its input varies is known as its *tuning curve*, and the ones shown in Figure 6A are typical for neurons throughout the brain.

In Figure 6B, we show the tuning curve for the rate neurons used by Gurney et al. (2001). This does not look like the realistic tuning curves of Figure 6A. However, Figure 6C shows that we can implement the effects of such a tuning curve by adding together the realistic tuning curves of 6A. This allows a group of realistic neurons to provide a similar effect to that assumed by the model.

When adding the outputs of the spiking neurons, we scale each one by a factor d_i , producing a weighted sum. We can compute the optimal d_i values using Equation 2, where the integration is over all possible inputs x , a_i is the average firing rate of neuron i given input x , a_j is the same for neuron j , and $f(x)$ is the desired output (Figure 6B). This calculation determines the least-squared-error solution for mapping the neural tuning curves onto the function $f(x)$. The method extends to complex functions and multiple dimensions, making it the basis of the Neural Engineering Framework (Eliasmith & Anderson, 2003).

$$\mathbf{d} = \Gamma^{-1} \mathbf{Y} \quad \Gamma_{ij} = \int a_i a_j dx \quad \mathbf{Y}_j = \int a_j f(x) dx \quad (2)$$

While there clearly must be a developmental or learning-based mechanism to determine these weights, we do not consider this here, just as we do not consider the developmental process for the creation of these separate brain areas in the first place. Instead, we assume that whatever such mechanisms exist converge to weights near the values determined by Equation 2.

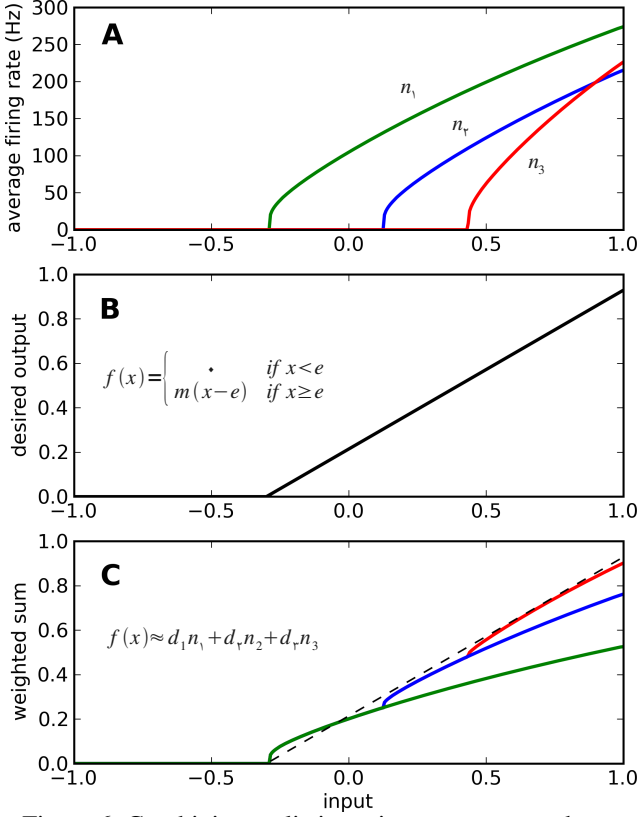


Figure 6: Combining realistic tuning curves to produce a desired function. **A** shows the average firing rate of three different neurons as the amount of input to the neurons increases. **B** shows the neural output function used by the rate neuron model. **C** shows how **B** can be approximated by taking a weighted sum of tuning curves in **A**.

Given these weighting values d_i , we can construct a spiking version of the model shown in Figure 4. Each single neuron in the original model is replaced by a set of 20 spiking neurons (increasing this value does not change our results). These all have the same time constant ($\tau_{RC}=20\text{ms}$; common throughout the brain), but have varying background currents and scaling factors to produce the range of tuning curves seen in Figure 6A. Each connection in the original model from rate neuron A to rate neuron B is replaced by a set of connections from all of the spiking neurons replacing A to all of the spiking neurons replacing B. The actual synaptic connection weight from the i th neuron in A to the j th neuron in B is $w\alpha_j d_i$, where α is the neuron's scaling factor and w is the original rate model's connection weight.

Finally, the timing effects of a neuron firing must be considered. This is vital for producing realistic temporal predictions from a model of spiking neurons. When a

neuron fires, it sends current into all of the neurons to which it is connected. This current $h(t)$ can be characterized by Equation 3, where τ_s captures the effects of neurotransmitter re-uptake and dispersal. As shown in Figure 7, a small τ_s provides a fast, short-lasting effect ($\sim 10\text{ms}$), while others last for hundreds of milliseconds.

$$h(t) = t e^{-t/\tau_s} \quad (3)$$




Figure 7: Post-synaptic currents for common synapses.

Importantly, different neurotransmitters are used by the different types of connections in the basal ganglia. All of the inhibitory connections involve GABA ($\tau_s=6.1\text{ms}$ to 10.5ms ; Gupta et al., 2000), while the excitatory ones of concern for this model involve fast AMPA-type glutamate receptors ($\tau_s=2\text{ms}$; Spruston et al., 1995). This means that the excitation and inhibition in the model act at different times scales, a factor not taken into account in the original model. As we show below, the time constants of these neurotransmitters have a strong impact on the temporal behaviour of our model.

Results

Figure 8 demonstrates that the model is capable of correctly performing action selection. Initially, action B has the highest utility, and the output shows that B is the only action that is not inhibited by the GPi/SNr outputs. In the middle, C is selected and has the highest activation, followed by A.




Figure 8: Spikes produced (bottom) for three possible actions (A, B, and C) as their utility changes (top).

Response Latency

One of the key advantages of using a realistic neural model is that timing predictions emerge from the neural parameters. We start by determining how long it takes the model to select an action when there is a sudden change in the input. Figure 9 shows the output for an action when its utility is suddenly increased at $t=0$. This matches empirical findings that in the rat basal ganglia, output neurons stop spiking 14 to 17 milliseconds after a similar input pulse (Ryan & Clark, 1991).




Figure 9: Spiking produced (bottom) for a sudden change in utility (top). Firing for action A stops 15.1ms after its utility is increased.

We can also examine how long it takes the model to decide between two actions as we adjust the difference between the top two utility values. Figure 10 indicates how the latency changes from very similar utility values (38ms mean latency, standard deviation 8.8ms) to highly differing utility values (14ms mean latency, standard deviation 1.5ms). As far as we are aware, this is a novel prediction.




Figure 10: Mean and standard deviation of basal ganglia response latency as for varying differences between utilities. Error bars are 95% confidence intervals over 200 runs.

Cognitive Cycle Timing

In a full cognitive system, the output of the basal ganglia would be used to affect the firing of other areas of the brain (via the thalamus). This, in turn, will affect the input to the basal ganglia, perhaps causing a different action to be selected. This is the basis of our ongoing development of a full production system using spiking neurons (Stewart, Choo, and Eliasmith, 2010). To investigate how long this whole cycle requires, we need to include the thalamus and a simple cortical area in our model.

For the cortex, we create a group of 5000 spiking neurons representing the current state. These are connected to the inputs to the basal ganglia so that the utility input for each action will be the similarity (measured as the dot product) between the current state and the ideal state for that action. This is done using Equation 2, where $f(x)$ is this similarity measure. For the thalamus, we create neurons representing the actions of switching to each possible state. They are connected to the cortex similarly, such that the firing of one group of neurons in the thalamus will cause the cortical neurons to fire in a pattern representing that state.

To implement the chaining of actions one after the other, we connect the output of the basal ganglia to the thalamic neurons such that if the basal ganglia selects action A, this will stop the inhibition of the thalamic neurons representing state B, thus causing the cortex to go to state B, and the basal ganglia to select action B. The actions are chained so that A leads to B, B leads to C, C leads to D, and so on. This can be thought of as a set of production rules of the form “If A then B; If B then C; If C then D; etc.” The newly added connections are excitatory, using AMPA-type receptors ($\tau_s=2\text{ms}$). All other parameters remain the same.

With this model, we can measure the time taken to change from one action to the next. This provides a measure of the minimum amount of time needed to go from one step to the next in a sequence of cognitive actions. In cognitive models that use production systems, extensive behavioural data has been gathered indicating that this value should be around 50 milliseconds (Anderson et al., 1995).

Figure 11 shows the mean and standard deviation of the cycle times produce by our model. The shaded area shows the timing produced when the correct realistic time constants for the inhibitory GABA neurotransmitter are used. Importantly, there are no parameters in our model that we can vary to affect this performance. In should be noted that our model predicts cycle times between 34 and 44 milliseconds, which is somewhat shorter than the standard 50 milliseconds value. However, this result is only for simple actions: more complex actions are considered next.




Figure 11: Cognitive cycle times produced by our model as the time constant τ_s of the inhibitory neurotransmitter GABA varies. The shaded area indicates parameter settings consistent with neurophysiology (Gupta et al., 2000).

Cognitive models generally use a cycle time of 50ms.

To be cognitively useful, an action selection mechanism needs to be able to trigger more complex actions than those considered so far. In particular, production system rules generally allow actions that can send a value stored in one

brain area to another. To model this we can create connections between cortical areas such that driving a cortical area to a particular value causes a second cortical area to send its value to a third cortical area. This can be implemented using Equation 2 (see Stewart, Choo, and Eliasmith, 2010 for more details). The timing of these types of actions are shown in Figure 12. While simple actions require less than 50 milliseconds, complex actions require more than 50 milliseconds.




Figure 12: Cognitive cycle times produced for complex actions by our model as the time constant of the inhibitory neurotransmitter GABA varies.

Conclusions

We presented a spiking neuron model of action selection that matches the anatomy of the basal ganglia and does not assume the presence of diffuse inhibitory interneurons in the striatum. By constraining the neurons' behaviour to match that of real neurons in the basal ganglia, we produce timing predictions from our model without parameter fitting. Figure 9 shows that these predictions match well for single-cell recordings in rats, and Figure 11 shows a close match for a wide range of cognitive psychology results. Our model thus provides a neural explanation of the commonly used 50 millisecond cognitive cycle time (e.g. Anderson et al., 1995). It also produces novel predictions of increases to this cycle time for situations where two possible actions have similar utilities (Figure 10) and for actions involving information transfer between brain areas (Figure 12).

References

- Albin, R. L., Young, A. B. & Penney, J. B. (1989). The functional anatomy of basal ganglia disorders. *Trends in Neurosciences*, 12(10): 366-375.
- Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., & Qin, Y. (2004). An integrated theory of the mind. *Psychological Review* 111(4), 1036-1060.
- Anderson, J. R., John, B. E., Just, M. A., Carpenter, P. A., Kieras, D. E., & Meyer, D. E. (1995). Production system models of complex cognition. *17th Annual Meeting of the Cognitive Science Society*.
- Belavkin R. & Huyck, C. (2009). A model of probability matching in a two-choice task based on stochastic control of learning in neural cell-assemblies. *9th International Conference on Cognitive Modelling*.
- Blomeley, C.P. & Bracci, E. (2009). Serotonin excites fast-spiking interneurons in the striatum. *The European journal of neuroscience*, 29(8):1604-1614.
- Bracci, E., Centonze, D., Bernardi, G., & Calabresi, P. (2002). Dopamine excites fast-spiking interneurons in the striatum. *Journal of neurophysiology*. 87(4):2190-2194.
- Eliasmith, C. & Anderson, C. (2003). *Neural Engineering: Computation, representation, and dynamics in neurobiological systems*. Cambridge: MIT Press.
- Gupta, A., Wang, Y., & Markram, H. (2000). Organizing Principles for a Diversity of GABAergic Interneurons and Synapses in the Neocortex. *Science* 287(5451), 273-278.
- Gurney, K., Prescott, T., & Redgrave, P. (2001). A computational model of action selection in the basal ganglia. *Biological Cybernetics* 84, 401-423.
- Hazy, T.E., Frank, M.J., & O'Reilly, R.C. (2007). Towards an executive without a homunculus: computational models of the prefrontal cortex/basal ganglia system. *Philosophical Transactions of the Royal Society B*.
- Humphries, M., Stewart, R., & Gurney, K. (2006). A physiologically plausible model of action selection and oscillatory activity in the basal ganglia. *The Journal of Neuroscience*, 26(50), 12921-12942.
- Koós, T. & Tepper, J.M. (2002). Dual cholinergic control of fast-spiking interneurons in the neostriatum. *The Journal of Neuroscience*, 22(2), 529-535.
- Mink, J. W. (1996). The basal ganglia: Focused selection and inhibition of competing motor programs. *Progress in Neurobiology*, 50, 381-425.
- Nambu, A., Tokuno, H. & Takada, M. (2002). Functional significance of cortico-subthalamo-pallidal 'hyperdirect' pathway. *Neuroscience Research*, 43, 111-117.
- Redgrave, P., Prescott, T., & Gurney, K. (1999). The basal ganglia: a vertebrate solution to the selection problem? *Neuroscience* 86, 353-387.
- Ryan, L. & Clark, K. (1991). The role of the subthalamic nucleus in the response of globus pallidus neurons to stimulation of the pre-limbic and agranular frontal cortices in rats. *Exp Brain Res*, 86, 641-651.
- Shouno, O., Takeuchi, J., & Tsujino, H. (2009). A spiking neuron model of the basal ganglia circuitry that can generate behavioral variability. *The Basal Ganglia IX: Advances in Behavioral Biology*. Springer.
- Spruston, N., Jonas, P., & Sakmann, B. (1995). Dendritic glutamate receptor channel in rat hippocampal CA3 and CA1 pyramidal neurons. *J. Physiology* 482, 325-352.
- Stewart, T., Choo, X., & Eliasmith, C. (2010). Symbolic reasoning in spiking neurons: A model of the cortex/basal ganglia/thalamus loop. *32nd Annual Meeting of the Cognitive Science Society*.
- Stewart, T., & Eliasmith, C. (2009). Spiking neurons and central executive control: The origin of the 50-millisecond cognitive cycle. *9th International Conference on Cognitive Modelling*.
- Stocco, A., Lebiere, C., & Anderson, J. R. (in press). Conditional routing of information to the neocortex: A network model of basal ganglia function. *Psych. Review*.
- Tepper, J.M. & Bolam, J.P. (2004). Functional diversity and specificity of neostriatal interneurons. *Current opinion in neurobiology*, 14:685-692.