

# Chapter 5

## Biological cognition – control

- mention somewhere that there are other ways of doing action selection (hard coding in cortex?), that come into play when bg is destroyed...???

### 5.1 The flow of information

To this point in the book, we have looked at two standard issues in cognitive science, semantics and syntax. Both of these are centrally concerned with representations – both the construction and manipulation of representations. A third crucial, and often over-looked, topic for understanding flexible cognitive systems is control. That is, controlling the flow of information through the system in order to determine what to represent, how to manipulate representations given current context, and deciding what an appropriate next course of action is. In short, determining how to guide the system through the world.

I suspect control is somewhat under-examined because it does not become absolutely indispensable until we start to construct large and complex models. Consequently, many models – which focus on specific biological systems, or specific tasks – do not really face the problem of control. Most models in the behavioral sciences are built to explain tasks restricted to a single often tightly constrained domain, be it reading, navigating, eye control, memory performance, pattern recognition, simple decision making, or what have you. Cognitive systems need to be able to perform all of these tasks, switching between them as appropriate. Consequently, any cognitive architecture will need to specify how information can be routed to different areas of the brain to support a given task, how the same information can be interpreted differently depending on the context, how the system

can determine what an appropriate strategy is, and so on.

The process of control can be usefully broken down into two parts: 1. determining what an appropriate control signal (or state) is; and 2. applying that control signal to affect the state of the system. The first of these tasks is a kind of decision making. That is, determining what the next course of action should be based on currently available information. The second task is more of an implementational issue: how can we build a system that can flexibly gate information between different parts of the system. For instance, if we are speaking on the phone and asked by a friend to report only what we are currently seeing, our brain needs to both decide what action to pursue (e.g., answering the friend's question, which necessitates translating visual to verbal information), and then actually pursue that action (e.g., allowing only visual information to drive our language system to generate a report). If the same friend asked for a report on what we were currently smelling, then our brain would configure itself differently, to allow olfactory information to drive the language system. If we could not significantly change the flow of information between our senses and our language system based on slightly different inputs, we would often generate irrelevant responses, or perhaps nonsense. In a sense, this re-routing of information is a kind of cognitive action. The importance of rapidly reconfigurable control is perhaps even more obvious for guiding motor action.

In this chapter, I describe the aspects of the semantic pointer architecture (SPA) most important for control. I begin by drawing heavily on the work of Terry Stewart in my lab on the role of the basal ganglia in action selection, with a focus on cognitive actions. I then describe work pursued by Bruce Bobier, also in my lab, on how actions that demand the routing of information can be implemented in a neurally realistic circuit. That work focusses on routing information through the visual cortex to explain aspects of visual attention. However, I describe how similar processes can be exploited throughout the SPA. These two discussions form the foundation of an example circuit that combines all of the aspects (semantics, syntax, learning, control) of the SPA to do blah???

## 5.2 The basal ganglia

The basal ganglia are a highly interconnected cluster of brain areas found underneath the neocortex and surrounding the thalamus (see figure 5.1). For well over twenty years, they have been centrally implicated in action selection, that is, selecting one from among several alternative actions. Consequently, damage to the

Figure 5.1: From bryan's thesis

Figure 5.2: A schematic diagram of the basal ganglia showing the standard direct/indirect pathway, the hyperdirect pathway and other major connections that have been recently discovered. permission???

basal ganglia (BG) is known to occur in several diseases of motor control, including Parkinson's and Huntington's diseases. Interestingly, these 'motor' diseases have more recently been shown to result in significant cognitive defects as well (ref frank???). Consequently, both neuroscientists (e.g., Redgrave et al., 1999) and cognitive scientists (e.g., Anderson et al., 2004) have come to understand the basal ganglia as being responsible for 'action' selection broadly construed to include both motor and cognitive action.

The basal ganglia circuit shown in figure 5.1 highlights the connectivity accounted for by the classic Albin/Delong model of basal ganglia function (DeLong, 1990; Albin et al., 1989). This model is able to qualitatively account for many of the symptoms of Parkinson's and Huntington's diseases. In this model, there is 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. 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 5.2.

- implicated in pulvemuller and lots of other language work for sequencing...

A recent model of the basal ganglia has suggested that these extra connections can effectively underly action selection (Gurney et al., 2001). This model also takes advantage of the 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). The Gurney et al. model, which forms the basis of our model, makes use of this structure to explain basic action selection.

To see how, consider the striatum, STN and GPi/SNr circuit shown in figure 5.3. This circuit includes the direct and hyperdirect pathways. In this figure, there are three possible actions, which have different 'desirabilities' (or 'utilities') indicated by the input weights to the striatum and STN (0.3, 0.8, 0.5). The equally

Figure 5.3: 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. permission???

Figure 5.4: The model of action selection in the basal ganglia presented by Gurney, et al. (2001). The striatum D1 cells and the subthalamic nucleus (STN) are as in figure 5.3, while the striatum D2 cells and globus pallidus external form a modulatory control structure. permission???

weighted inhibitory connections in the direct pathway cause the most active input to most greatly suppress the action to be selected. However, if not for the additional excitation provided by the STN, all of the actions might be concurrently suppressed (meaning more than one action had been selected). These very broad STN connections thus take the input from the hyperdirect pathway, and combine it to provide a level of background excitation that allows only the most inhibited action to be selected.

It should be evident that the example shown in figure 5.3 has carefully selected utilities. In fact, if there are many actions with large utilities or all actions have low utilities, this circuit will not function appropriately. 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 task, 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 5.3, 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. This network is shown in figure 5.4.

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. 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 timing predictions. Furthermore, the model has no redundancy, since exactly one neuron is used per area of the basal ganglia to represent each action. The original version of the model shown in figure 5.4

Figure 5.5: Spikes produced (bottom) for three possible actions (A, B, and C) as their utility changes (top). The highest utility action is selected, as demonstrated by a suppression of the spikes of the corresponding action. In order, the highest utility action is B, C, A. permission???

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.

However, given the resources of the NEF and the SPA, these short-comings can be rectified. In particular, the SPA suggests that instead of a single neuron representing potential actions, a mapping from a high-dimensional semantic pointer into a redundant group of neurons should be used. As well the NEF provides a means of representing these high-dimensional semantic pointers and their utilities in spiking neurons, and reproducing the transformations suggested by the original model. I return to the role of SPAs shortly. For the time being, I consider a simpler model that represents actions and utilities as scalar values. The ability of this new model to select actions is demonstrated in figure 5.5. There it can be seen that the highest utility action (B then C then A) is always selected (i.e., inhibited).

Crucially, this new implementation of the model allows us to introduce additional neural constraints into the model which could not previously be included. In particular, the types of neurotransmitters employed in the excitatory and inhibitory connections of the model have known effects on the timing of a receiving neuron's response. All of the inhibitory connections involve GABA receptors (with time constants between about 6.1ms to 10.5ms; Gupta et al., 2000), while the excitatory ones involve fast AMPA-type glutamate receptors (with time constants of about 2ms; Spruston et al., 1995). The time constants of these neurotransmitters have a crucial impact on the temporal behavior of the model.

For instance, we can use this model to determine how long it takes the model to select an action when there is a sudden change in the input. Figure 5.6a shows the output for an action when its utility is suddenly increased. This matches empirical findings that in the rat basal ganglia, output neurons stop spiking 14 to 17 milliseconds after a similar input pulse (Ryan and Clark, 1991). However, matching a single run of a model to a single experiment only provides minimal support to the model. Instead, we can run that same experiment on the model while exploring a much broader range of utility inputs. This leads to the novel prediction shown in figure 5.6b. There we have used this same set up to generate a figure that indicates how the latency changes from very similar utility values (38ms mean latency, standard deviation 8.8ms) to highly differing utility values (14ms mean

Figure 5.6: a) Spiking produced (bottom) for a sudden change in utility (top). Firing for action A stops 15.1ms after its utility is increased. b) Prediction of the effect of changing the utility difference between two actions on the response latency of the basal ganglia. Error bars are 95% confidence intervals over 200 runs. permission???

latency, standard deviation 1.5ms). All of the inputs are the same as in figure 5.6, but the difference between the height of the gray and black lines have been altered.

### 5.3 Basal ganglia, cortex, and thalamus

We have seen how a spiking neural model of the basal ganglia can be used to select simple actions. However, the purpose of the SPA is to provide a framework for building large-scale cognitive models. These, of course, require complex actions, driven by sophisticated perceptual input. The previous chapters demonstrated how the SPA can support representations sufficiently rich to capture the syntax and semantics of such representations. Here, I consider how these can be used by the basal ganglia to drive cognitive behavior. Demonstrating this will proceed in steps: first I consider sequences of fixed action, then I consider flexible sequences of action, finally I present a system choosing among and applying cognitive strategies based on perceptual input, and using the results to drive motor control.

To construct such models, we can rely on the well-known and central cortex/basal ganglia/thalamus loop through the brain (see figure 5.7). Roughly speaking, the SPA assumes that cortex provides, stores, and manipulates representations, the basal ganglia map current brain states to future states by selecting appropriate actions, and the thalamus provides for real-time monitoring of the entire system.

While able to perform intricate tasks, cortex in the SPA is built out of combinations of three basic functions: integration (for working memory; section 6.6), multiplication (specifically convolution for syntactic manipulation; section 4.8), and the dot product (for clean-up memory and other linear transformations; sections 4.5 and 3.7). I have discussed each of these functions in detail in past tutorials. Learning clearly plays a large role in establishing and tuning these basic functions, as considered earlier. I leave further consideration of learning until later (sections 6.5). All of the cortical models generated in our lab are combinations of these functions, including many I have not discussed in this book, such as models

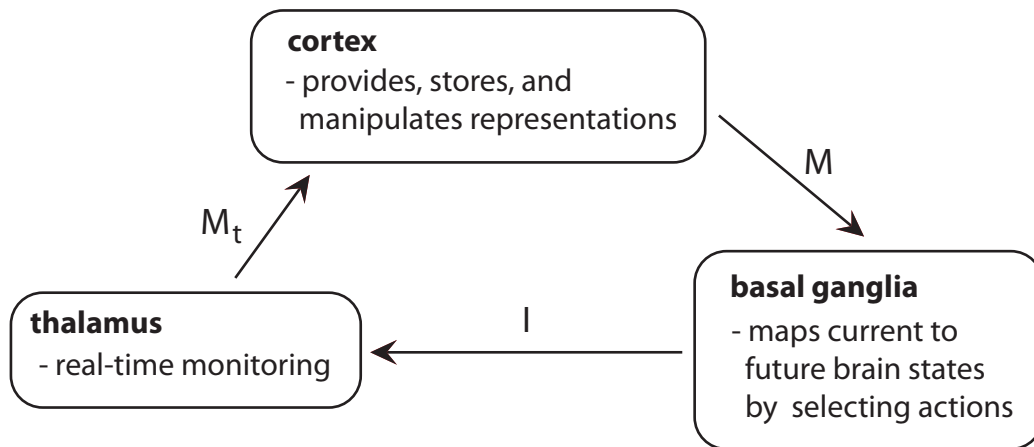


Figure 5.7: simple figure with all three and function labels, and M and M<sub>t</sub> and just I (identity) inhibition from bg to thalamus.

of serial memory (ref xuan???), path integration (rat???), and hemineglect (steve ???).

In order for these cortical operations to be flexibly exploited, however, it is essential to control the flow of information between them. This is the role of the basal ganglia model that I have just described. While there is evidence that cortex can perform ‘default’ control without much basal ganglia influence (ref???), the basal ganglia make the control more flexible, fluid, and rapid (ref???). The central location, and massive projections from (and back to, via thalamus) cortex make the basal ganglia ideal for playing the role of action selection. All areas of neocortex, except the primary visual and auditory cortices, project to the basal ganglia.

Finally, because basal ganglia and cortex are in many ways diffuse, performing many functions over a wide area, and because they are built on top of much older, and more fundamental control systems (e.g. brain stem and the thalamus itself), they need to be integrated with signals coming from other areas of cortex and these older systems. It would not do to be stuck ‘cognizing’ about an auditory input while in imminent danger. The thalamus is unique in its central location and structure in a way ideally suited to playing the role of co-ordinator of these systems. For instance, all output from the basal ganglia goes through the thalamus before returning to cortex. As well, the thalamus receives projections from every sense (except smell), and from all cortical areas, which it also projects to. These projections are organized by cortical region providing a somewhat topographic

map of cortex. Interestingly, however, the reticular nucleus of the thalamus communicates with and regulates the states of the other nuclei in the thalamus. Consequently, the thalamus is ideally structured to allow it to monitor a summary of the massive amounts of information moving through cortex and from basal ganglia to cortex. Not surprisingly, thalamus is known to play a central role in major shifts in system function (e.g. from wakefulness to sleep), and participates in controlling the general level of arousal of the system. In the models presented here, thalamus acts much like a basic relay (as it was long thought to be), because the cognitive coordination of the system can be accounted for by basal ganglia for the considered tasks. However, its contribution to timing effects are important and ultimately projections out of the thalamus control cognitive function so it is included in these examples.

As depicted in figure 5.7, communication between cortex and basal ganglia consists of mapping the contents of current cortical states to the striatum and STN (the basal ganglia input nuclei) through the  $\mathbf{M}$  matrix. This determines the utilities that drive the basal ganglia model in section 5.2. One natural and simple interpretation of the rows of this matrix is that they specify the antecedent of a rule. Consider the rule ‘if there is an A in working memory then set working memory to B’. The  $\mathbf{M}$  matrix can examine the contents of working memory, and output a list of similarities between its rows and working memory with a simple linear transformation (i.e.  $\mathbf{s} = \mathbf{M}\mathbf{w}$ , where  $\mathbf{s}$  are the similarities between the each of the rows of  $\mathbf{M}$  and the vector  $\mathbf{w}$ , the input from working memory). That vector of similarities then acts as input to basal ganglia, which selects the highest similarity (utility) from that input.

The output from basal ganglia results in a release from inhibition of the connected thalamic neurons, which are then mapped back to cortex through  $\mathbf{M}_t$ . This matrix can be thought of as specifying the consequent of a rule, resulting, for example, in setting working memory to a new state B. More generally, the  $\mathbf{M}_t$  matrix can be used to specify any consequent control state given a current cortical state. This loop from cortex through basal ganglia and thalamus and back to cortex forms the basic control structure of the SPA. Notably, all of the representations mentioned in this loop (i.e., of cortical states, control states, etc.) are semantic pointers.



## 5.4 Example: Fixed sequences of action

A simple but familiar example of a fixed sequence of action is rehearsal of the Roman alphabet. This is an arbitrary sequence, with no systematic rule connecting one state to its successor (unlike counting). Consequently, we need to specify 25 rules to be able to traverse the entire sequence. All such rules would be of the form:

IF working memory contains **letter** + **A**  
 THEN set working memory to **letter** + **B**

where bold indicates that the item is a 250 dimensional semantic pointer. For present purposes, these pointers are randomly generated. So, the inclusion of **letter** in each of the pointers provides minimal semantic structure that can be exploited to trigger actions appropriate for any letter (see section 5.6).

To implement the IF portion of the rule, an **M** matrix with rows consisting of all the letter representations is constructed and embedded in the connection weights between working memory and the basal ganglia input (using the methods detailed in section 3.7). This allows the basal ganglia to determine what rule is most applicable given the current state of working memory.

The THEN portion of the rule is then implemented by the **M<sub>r</sub>** matrix in a similar manner, where only one row is activated by disinhibition (as determined by the basal ganglia), sending a given letter representation to working memory. As working memory is being constantly monitored by the basal ganglia, this new state will drive subsequent action selection, and the system will progress through the alphabet.

To run the model, it is initialized by setting the working memory neurons to represent **letter** + **A** semantic pointer. After this, all subsequent activity is due to the interconnections between neurons. Figure 5.8 shows the model correctly following the alphabet sequence. From the spiking pattern we see that the correct action for each condition is successfully chosen by turning off the appropriate inhibitory neurons in the GPi. The top plot is generated by comparing the semantic pointer of each of the 26 possible letters to the current semantic pointer in working memory (decoded from spiking activity) by using a dot product, and plotting the top eight results.

- This model can be downloaded from... URL???

This model demonstrates that a well learned set of actions with a specific representation as an outcome can be implemented in the SPA. However, this model is




Figure 5.8: Contents of working memory (top) generated by taking the dot product of all possible semantic pointers with the decoded contents of working memory (top eight values are shown). The spiking output from GPi indicating the action to perform (bottom) demonstrates that the population encoding the currently relevant IF statement stops firing, disinhibiting thalamus and allowing the THEN statement to be loaded into working memory.




Figure 5.9: The (erroneous) result of permanently connecting the visual system to working memory. The perceptual system’s data (in this case **letter + B**) continually drives working memory (top) and prevents it from properly moving to **letter + C** and subsequent states. The spiking out from GPi (bottom) indicates that it is no longer able to select actions.

not particularly flexible. For instance, we might assume that our working memory is being driven by perceptual input, say from vision. If so, we would have a permanent connection between our visual system and the working memory system that is driving action selection. However, if this is the case, then changing the visual input during the fixed action sequence will cause the sequence to shift suddenly to the new input, as shown in figure 5.9. In fact, just leaving the visual input ‘on’, would prevent the model from proceeding through the sequence, since the working memory would be constantly driven to the visually presented input.

In short, the current model is not sufficiently flexible to allow the determined action to be one which actually changes the control state of the system. That is, we are currently not in a position to gate the flow of information between brain areas using the output of the basal ganglia. However, routing information flexibly through the brain seems to be a fundamental neural process, one which often goes by the name of ‘attention’.

Figure 5.10: Anatomical structures relevant to the dynamic routing model (DRM).  
 ???name all the structures here?

## 5.5 Attention and the routing of information

A attention has been most thoroughly studied in the context of vision. And while visual attention can take on many forms (e.g., attention to color, shape, etc.) the spatial properties of visual attention of the most thoroughly examined. There are two main considerations when it comes to understanding visuospatial attention: selection and routing (notably analogous to the two aspects of any control problem). Selection deals with the problem of identifying what the appropriate target of the attentional system is given current task demands and perceptual features. Routing, in contrast, deals with how, once a target has been selected, a neural system can take the selected information and direct additional resources towards it. Figure 5.10 identifies the brain areas thought to be involved in both selection and routing. In this section, I will only discuss the routing problem, which in general has not received as much consideration as selection, and considers only those brain areas outlined in grey in figure 5.10.

In the last 15 years there have been several proposed models of attentional routing (Olshausen et al., 1993; Salinas and Abbott, 1997; Reynolds et al., 1999; Wolfrum and von der Malsburg, 2007; Womelsdorf et al., 2008). However, none of these uses biophysically plausible spiking neurons, and most are purely mathematical models. Consequently, most, if not all, have been criticized as not being scalable (i.e. there are not enough neurons in the relevant brain structures to support the required computations). Given the importance of both scalability and biological plausibility to the SPA, Bruce Bobier in my lab has developed a novel model of attentional routing called the ‘dynamic routing model’ (DRM).

The DRM draws on these past models in several ways. For instance, it relies on nonlinearities to perform the routing. It also incorporates connectivity constraints, general architectural considerations, and a hierarchical structure from several of these models. It is most directly a descendant of the ‘shifter circuit’ model (Olshausen et al., 1993). Nevertheless, it is uniquely biologically plausible and scalable, and so it provides an especially good account of routing that can be generalized to other parts of the SPA, as I will suggest in the next section. So, let us consider the DRM in more detail.

As shown in figure 5.10, the DRM consists of a hierarchy of visual areas, V1, V2, V4, and PIT, which receive a control signal from a part of the thalamus called

Figure 5.11: Architecture of the dynamic routing model (DRM) of visual attention. Each level has a columnar and retinotopic organization, where columns are composed of visually responsive pyramidal cells and control neurons (black dots). Grey circles indicate columns representing an example focus of attention. Neurons in each column receive feedforward visual signals (grey lines) and a local attentional control signal from control neurons (black lines), and these signals interact nonlinearly in the terminal dendrites of pyramidal cells (small open circles). Coarse attentional signals from Pdm are relayed through each level of the hierarchy downward to control neurons in lower levels (dashed lines). Connectivity is highlighted for the rightmost columns only, although other columns in each level have similar connectivity. ???this needs to be black & white

the pulvinar, specifically Pdm (dorsal medial pulvinar). A variety of anatomical and physiological evidence suggests that Pdm is responsible for providing a top-down control signal to the highest level of the visual hierarchy (Petersen et al., 1985, 1987; Stepniewska, 2004).

A more detailed picture of the connectivity between the visual areas is provided by figure 5.11. In that figure, an example focus of attention that picks out approximately the middle third of the network is shown. To realize this effective routing (i.e. of the middle third of V1 up to PIT), Pdm provides a control signal to PIT indicating the position and size of the current focus of attention in V1. The control neurons in PIT use this signal to determine what connections to ‘open’ between V4 and PIT, and then send their signal to control neurons in V4.

To ‘open’ a connection essentially means to multiply it by a non-zero factor. In essence, the control neurons act as a gate to determine which lower-level (e.g. V4) neurons are allowed to project their information to the higher level (e.g., PIT). In the DRM, this gating is realized by the non-linear dendritic neuron model mentioned in section 4.2.2. This kind of neuron essentially multiplies parts of its input, and sums the result to drive spiking. Consequently, they are ideal for acting as a gate, while using fewer neurons than a two-layer network performing the same function.

The signal that is sent to the next lower level of the hierarchy is similarly interpreted by that level to determine what gating is appropriate, and then passed on. Again, the gating allows the flow of information only from those parts of the next lower level that fall within the focus of attention. And so this process of computing and applying the appropriate routing signal continues to V1.

The mapping of these computational steps to cells in specific cortical layers

Figure 5.12: Proposed laminar arrangement of the DRM neurons.

Figure 5.13: An example of routing in the DRM.

of these visual areas is shown in figure 5.12. As shown here, two sets of neurons in layer V receive the top-down control signal specifying the size and position of the desired routing. These then project to layer VI neurons that determine an appropriate sampling and shift consistent with the desired routing. The results of this computation are sent to layer IV neurons, which act to gate the feedforward signals into layer II/III neurons, which project to higher levels of the visual hierarchy. A more detailed discussion of the anatomy and physiology underlying this mapping can be found in (cite bruce ref???)

This organization, and the computations it under-writes is repeated throughout the model (consistent with most contemporary accounts of cortical organization). Consequently, the model, while detailed, is reasonably straightforward as its central features are simply repeated over and over throughout the width and depth of the hierarchy shown in figure 5.11.

Figure 5.13 shows an example of the shifting and resizing of information presented to V1 made possible by this mechanism. As can be seen in that example ???more on figure???. It should be evident that the circuit is thus effectively taking the information on the lowest level of the hierarchy and ‘moving’ or ‘routing’ it to always fit within the available resources at the highest level. There is strong psychophysical evidence that precisely this kind of ‘normalization’ occurs in the visual system (refs???). As well, such a model is able to account for several detailed observations about attention.

For example, it accounts for the known increase in size of the receptive fields of neurons farther up the hierarchy, is consistent with the patchy connectivity in the visual hierarchy (Felleman et al., 1997), and captures the topographic organization of receptive fields in these areas (Tanaka, 1993). The DRM also accounts for the observation that attentional modulation of neural activity starts at the top of the hierarchy and proceeds down it (Mehta, 2000; Buffalo et al., 2010). Several more detailed changes in receptive field size, response strength, and timing are also captured by the DRM (ref bruce paper/thesis???)

However, to demonstrate the close tie the model has to biological detail, let me consider just one set of experiments in some detail. In this task, recordings were taken from monkeys performing a spatial attention task (see figure 5.14).

Figure 5.14: The Womelsdorf task. blah???Based on womelsdorf figure 1A...

Specifically, the animals foveated a fixation point, after which a cue stimulus (S1) was presented for 440ms at one of three target locations, indicating where the animal should covertly attend. Following a delay, three stimuli were presented, one at the target location, and two distractor stimuli (S2 and S3) one inside and one outside of the recorded cell's receptive field. The animal had to indicate a brief change in the stimulus at the cued location, some random interval after the three stimuli were presented. During the interval, the animal was taken to have sustained spatial attention to S1, and the receptive field of the cell was mapped. This experimental design allowed Womelsdorf et al. to map the receptive field during sustained states of selective attention to the stimuli inside the receptive field or to the stimulus outside the receptive field.

A similar experiment was run on the model, though mapped to a 1D input space. The same methods of fitting the spiking data to determine neuron receptive fields were used in the model as in the experiment. The main difference between the model and experimental runs is that in the model, all spikes could be collected from all neurons, though the number of neurons is smaller. As a result, we ran the experiment on 100 different versions of the model. Each 'version' has the same DRM architecture, but the neurons themselves are randomly chosen from a distribution of parameters that matches the known properties of cells in the relevant part of cortex. Consequently, rather than having just over a hundred neurons from two animals as in the experiment, we have thousands of neurons from hundreds of model-animals. This means that we have a much better characterization of the overall distribution of neuron responses in the model than is available from the experimental data for the animals.

To compare the model and data, we considered the three main effects described in the experimental work. These effects could be seen by comparing attention being directed at a stimulus inside the receptive field to attention being directed at a stimulus outside the receptive field. The effects were 1) a change in the peak firing rate, 2) a shift of the receptive field center, and 3) a change in the receptive field size. In each case, several statistics were calculated to compare model and experimental data, though the details of the analysis are beyond the scope of the book (ref bruce???)

In sum, the model and data are consistent on each of these effects. More specifically, the data suggested an increase in the peak firing rate over all neurons of  $4.7\% \pm 8\%$  (2 standard deviations), while the model had a 8.13% mean (95%

confidence interval (CI) = [7.4%, 8.9%]). The shift of the receptive field centers was calculated from the data to be  $31.4\% \pm 5.6\%$ , while for the the model it was 31.41% (CI=[31.18%, 31.65%]). And finally, for the change in receptive field size, the mean calculated from the data was  $-12.1\% \pm 3.8\%$ , while for the model it was 8.68% on average (CI=[8.10%, 9.22%]).

The last results may not seem consistent, but in fact are. This is because the distribution, which we can characterize very well with the model shows that 95% of the time, the mean of a single animal experiment will fall between -18% and 25%. Consequently, choosing two animals with negative means of the size reported in the experiment is not especially unlikely (it will happen about 30% of the time). However, it does suggest that if the model and experiments are likely to disagree anywhere, it will be in the change of receptive field size. Consequently, it would be very helpful to run additional animals in this condition to determine if the model distribution is accurate in this regard.

One final note about this model is that there are only two free parameters. Both of these parameters were set to match the known receptive field sizes in visual cortex. That is, they were in no way tuned to ensure the ability to predict this specific experiment. The remaining parameters were randomly chosen between different model-animals from distributions known to match general physiological characteristics of cells in visual cortex. Nevertheless, the model is able to provide a very promising characterization of changes in several subtle aspects of individual cell activity. This makes us reasonably confident that the central computational principles related to routing information in the biological brain are accurately captured by this model.

## 5.6 Example: Flexible sequences of action

With this characterization of routing in hand, we can return to the model of section 5.4, and repair its shortcomings. As you will recall, the model as presented was unable to appropriately ignore visual input (see figure 5.9). However, if we include the ability of the thalamus to generate gating signals that control the flow of information between visual input and working memory, we can prevent such errors.

The model of visual attention presented in the preceding section provides a biologically realistic mechanism for exactly this function. While it is clearly geared to a different part of cortex, it demonstrates how thalamic gating signals (in that case, from pulvinar) can be used to control the flow of information through a cor-



tical circuit. As well, there is nothing computationally unique about visual cortex, so it would be unsurprising to see similar functions computed in similar ways in other parts of cortex. In short, this means that computing a nonlinearity between a thalamic input and cortical signals can be considered a plausible and effective means of routing information in many areas of cortex. So, we can introduce a similar circuit into our action-selection model, and provide for much more flexible control.

The addition of this mechanism will allow us to effectively gate the visual input to the selection model, meaning it can selectively ignore and be driven by that input. For instance, we can set the action to take after working memory is equal to ‘Z’, to be one that routes information from visual input to working memory (i.e. ‘look’). All of the other actions would not allow this information flow. As shown in figure 5.15, this model is now able to flexibly ignore and consider the visual input as appropriate.

Notably, in this network a much more general rule is employed than was used before. Specifically, there is now a rule of the form

IF visual cortex contains **letter**+?  
THEN copy visual cortex to working memory

Such a rule applies to every letter, not just the one that happens to be in working memory at the moment. This allows for rules to be defined at a more general category level than in the previous example. This demonstrates an improvement in the flexibility of the system, in so far as it can employ instance specific, or category specific rules.

In addition, the ‘copy visual cortex’ command is the specification of a control state that consists in gating the information between visual inputs and working memory. This demonstrates a qualitatively new kind of flexibility that is available once we allow actions to set control states. In particular, it shows that not only the content of cortical areas, but also the communication between such areas, can be controlled by our cognitive actions.

However, routing can also provide more flexibility than simply gating the information flow between different cortical areas. In the simple alphabet model above, the routing action was essentially ‘on’ or ‘off’. However, we can use the same structure to actually *process* the signals flowing between areas. Recall that the method of binding semantic pointers is to use circular convolution, which is a linear transformation followed by multiplication. As described in the section on attention, gating can also be accomplished with a linear transformation followed by multiplication. Hence, we can use the same gating circuits to bind and




Figure 5.15: Routing information. Contents of working memory are shown on top; the lower half of the graph shows spiking output from GPI indicating the action to perform. The *look* action takes information from visual cortex (in this case, **letter + F**) and routes it to working memory.

unbind semantic pointers, not only routing, but concurrently processing signals. Essentially, we can give the gating signals useful content.

Consequently, same network structure as above can be used to perform syntactic processing, such as question answering. In essence, we can implement a dynamic, controlled version of the question answering network described in section 4.3. In this network, we define semantic pointers that allow us to present simple language-like statements and then subsequently ask questions about those statements. So, for example, we might present the statement

**statement + blue  $\otimes$  circle + red  $\otimes$  square**

to indicate that a blue circle and red square are in the visual field. We might then ask a question in the form

**question + red**

which would be asking ‘What is red?’. To process this input, we can define the following rules

IF the visual cortex contains **statement**+?  
THEN copy visual cortex to working memory

which simply gates the visual information to working memory as before. We can also define a rule that performs syntactic processing while gating

IF visual cortex contains **question**+?  
THEN apply visual cortex to the contents of working memory

Here, ‘apply’ essentially indicates that the contents of visual cortex are to be convolved with the contents of working memory, and the result is stored in the network’s output. More precisely, the contents of visual cortex are moved to visual working memory store (to allow changes in the stimulus during question answering, as above), and the approximate inverse (a linear operation) of visual working memory is convolved with working memory to determine what is bound to the question. This result is then stored in an output working memory to allow it to drive a response. The results of this model answering two different questions from the same remembered statement are given in figure 5.16. These two generic rules can answer any question provided in this format.

To be clear, this example is not intended to suggest that connection weights between cortex and basal ganglia must be changed depending on how an alphabet processing task is specified. As cognitive modellers have explored in the past,




Figure 5.16: Answering two different questions starting from the same statement. Grey areas indicate the period during which the stimuli were presented. The similarity between the contents of network's output and the top 7 possible answers is shown. The correct answer is chosen in both cases.

the interpretation of task commands is likely to depend strongly on a ‘declarative memory’, which I have not considered here. But notice that routing circuits could allow an arbitrary rule from memory to be rapidly and temporarily ‘loaded’ into the  $\mathbf{M}$  and  $\mathbf{M}_t$  matrices by gating a memory store. More likely, there are cognitive actions which control the routing of information through cortex itself to flexibly interpret currently provided rules. Only in cases where actions are often performed would the connections to basal ganglia be updated through reinforcement learning. I return to some of these issues in section 6.5.

Notably, this exact model can reproduce all of the control examples presented to this point. This means that the introduction of these more flexible control structures does not adversely impact any aspects of the simpler models’ performance as described above. This is crucial to claims of flexibility. The flexibility of the SPA needs to be independent of its particular use: flexibility needs to reside in the overall design of the system, not the task specific changes introduced by a modeller. (???URL of model that lets people run all of these examples and more???)

In addition, we can be confident that this control circuit will not adversely affect the scaling of the SPA. Only about 100 neurons need to be added for each additional rule in the basal ganglia. Of those, about 50 need to be added to striatum, which contains about 55 million neurons (Beckmann and Lauer, 1997), 95% of which are input (medium spiny) neurons. This suggests that we can estimate that about one million rules can be encoded into a scaled-up version of this model. In combination with the reasonable scaling of the representational aspects of the SPA (see section 4.5), this suggest that the SPA as a whole will scale appropriately.

- pointer here (and lots of other places?) to movies of a network like this in action... gives a much better sense of dynamics, etc.

## 5.7 Timing and control

Because the model of the basal ganglia that we have presented is constructed using a variety of biological constraints, we are able to ask questions of this model that have not been addressed adequately in the past. Specifically, because we know the kinds of neurons, their spiking properties, and the temporal properties of the neurotransmitters in basal ganglia, we can make specific predictions about the timing action selection (refs terry’s stuff)???

For example, Ryan and Clark (1991) showed that in the rat basal ganglia the output neurons stop firing 14 to 17ms after a rapid increase in the utility of one of

the possible actions. It is a simple matter to run such an experiment on the basal ganglia model. To keep things simple, we can include two possible actions, A and B, and rapidly change the utility by directly changing the input to striatum (i.e., by controlling the best matching action to current cortical activity). An example run of such a model is shown in figure 5.17a, which results in a cessation of firing in approximately 15ms.

More interestingly, we can explore the effects of the difference in utility between the two actions on the length of time such firing persists. These more informative results are presented in figure 5.17b, where it can be seen that the latency can increase to about 38ms for actions that are only slightly different in utility after the change. For the largest utility differences, the latency drops to about 14ms, the same lower bound as the Ryan and Clark experiment. As far as I am aware, this latency profile remains an untested, but strong prediction of the model.

We can also look at the effects on timing of the basal ganglia in the context of the entire cortex-basal ganglia-thalamus loop. In fact, this timing is implicitly demonstrated by figure 5.8. There it can be seen that in the fixed action selection case, it takes about 40ms for the system to switch from one action to another. This is more fully characterized in figure 5.18a, where the mean and standard deviation of timing is shown over a range of time constants of the neurotransmitter GABA. GABA is the main neurotransmitter of the basal ganglia, and has been reported to have a decay time constant of between 6 and 11ms (refs??? terry). We have identified this range on the graphs with a grey bar. On this same graph, we have drawn a horizontal line at 50ms because this is the standard value assumed in most cognitive models for the length of time it takes to perform a single cognitive action (Anderson et al., 1995) (other refs terry???).

Interestingly, the cycle time of this loop depends on the complexity of the action being performed. Specifically, figure 5.18a shows the simplest fixed action cycle time, where figure 5.18b shows the more complex flexible action cycle time, like those discussed in section 5.6. The main computational difference between these two types of action is that the latter has a control step, which can re-route information through cortex. The cycle time is affected by increasing from about 30-45ms in the simple case, to about 60-75 ms in the more complex case. These two instances clearly bracket the standard 50ms value. Notably, this original value was arrived at through fitting of behavioural data. If the tasks used to infer this value include a mix of simple and complex decisions, arriving at a mean between the two values we describe here makes perfect sense.

- (???Terry mentioned possible evidence for needing quicker times for easy

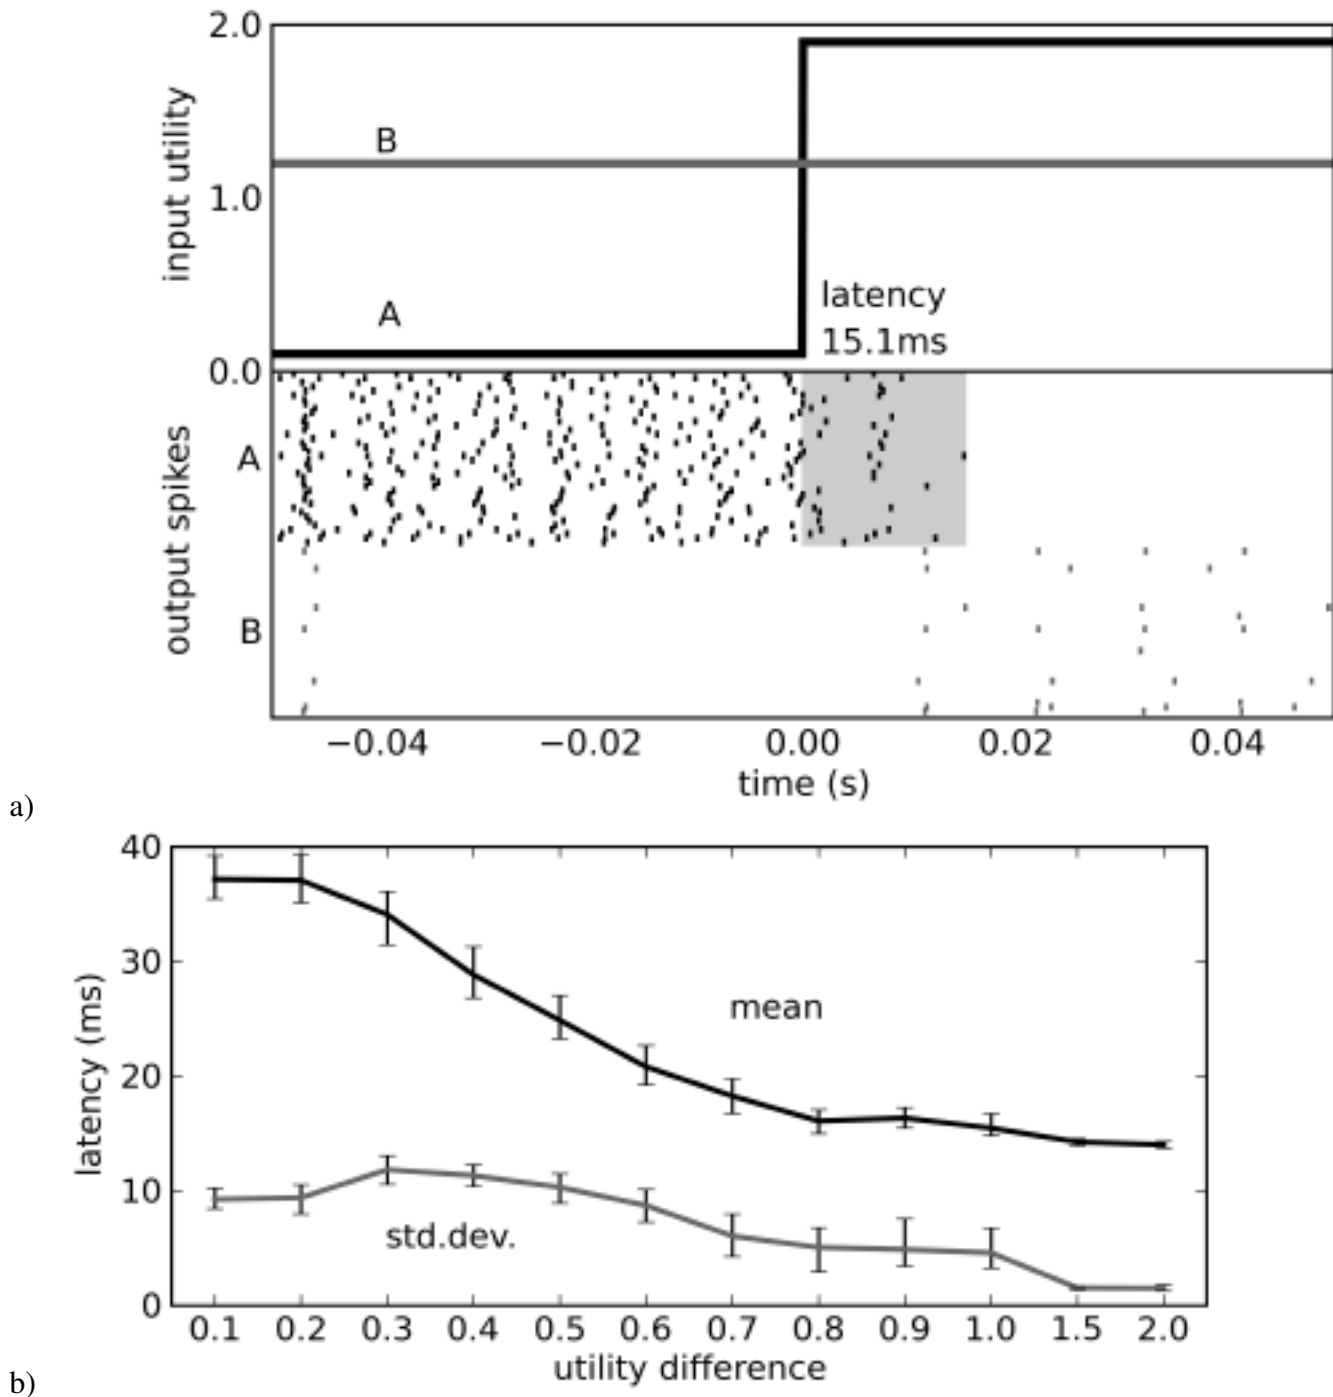
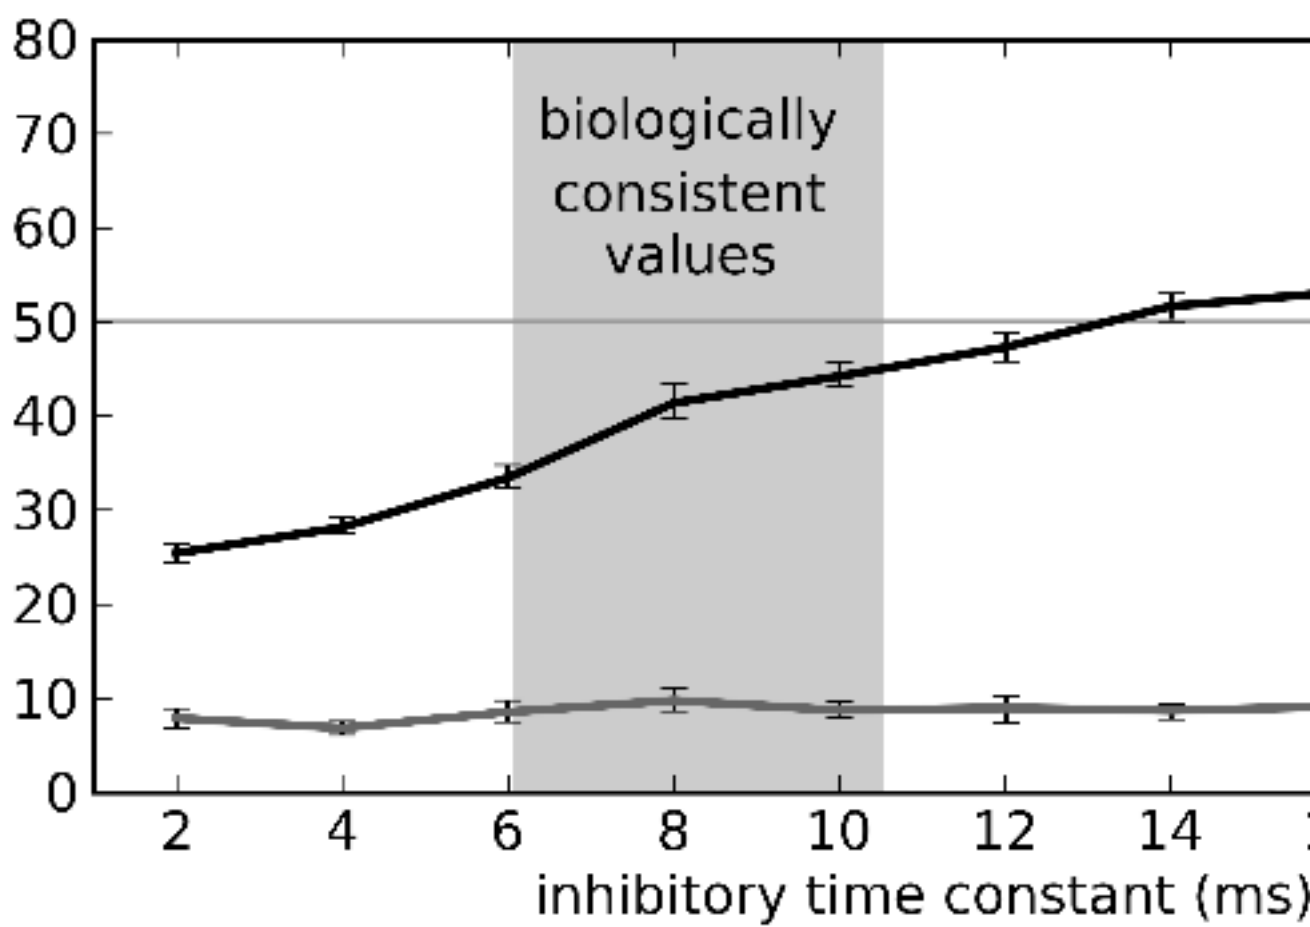
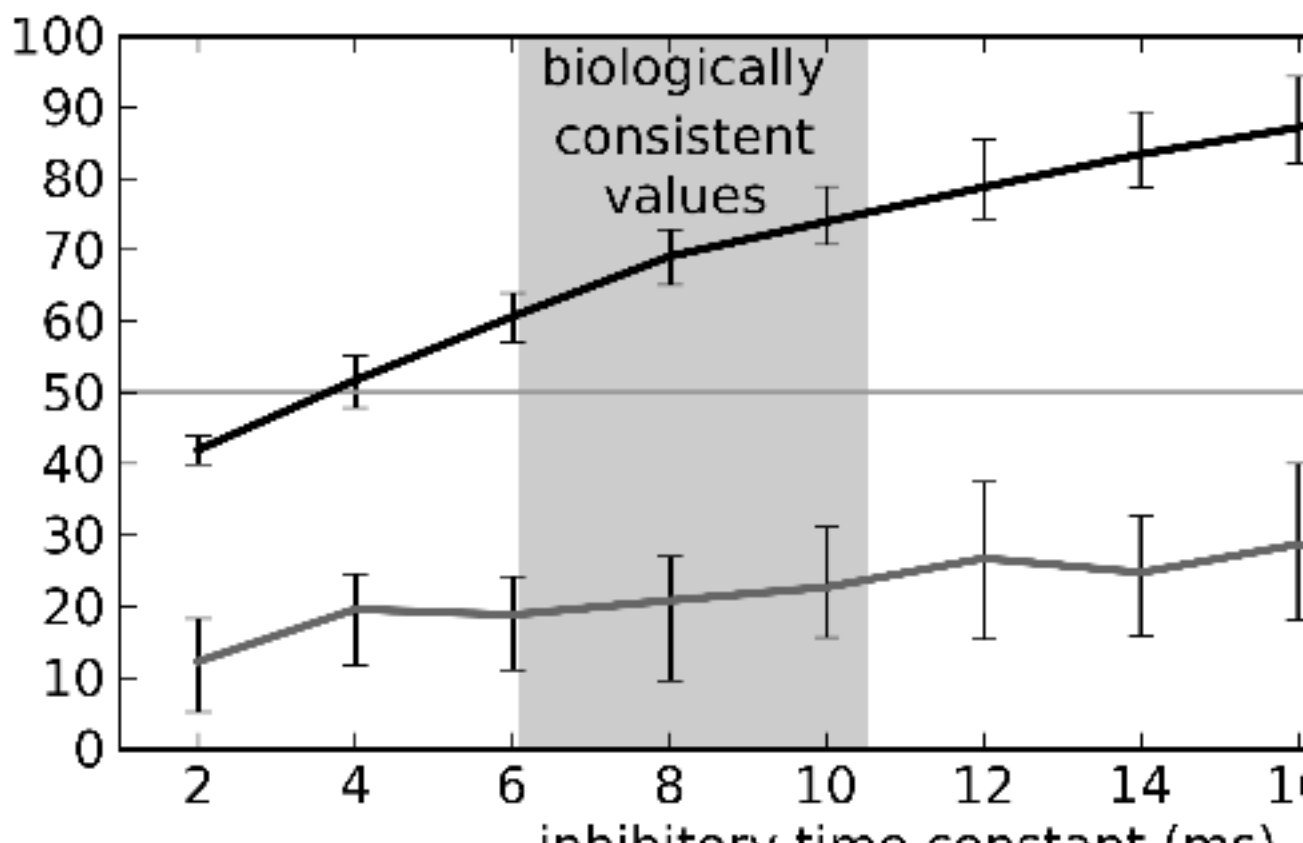


Figure 5.17: Timing predictions for changes in action utility. a) An example of the delay time between a rapid change in utility (top) and the change in spiking activity choosing a new action (bottom). Here, firing for action A stops 15.1ms after the change in utility (hence it is chosen). b) Averages and standard deviations of such changes over many utility differences. Error bars are 95% confidence intervals over 200 runs. ???these are place holders, need better quality from terry, and make one figure???



a)





tasks... citation???)

Together, these timing results help to highlight some of the unique properties of the SPA. For example, it can help provide an explanation of the genesis of certain ‘cognitive constants’ that would not otherwise be available, and it can address available neural data about the impact of utility on selection speed. These explanations cannot be provided by any basal ganglia models implemented in rate neurons, because they do not include the relevant timing information (i.e. neurotransmitter time constants). They are also not available from more traditional cognitive modelling approaches, such as ACT-R, that have determined such constants by fits to behavioural data. As well, these explanations are not available from other neurally-inspired architectures that include basal ganglia as an action selector, such as LEABRA, because the winner-take-all mechanisms in such models do not have temporal constraints (it is a direct winner-take-all calculation).

Furthermore, the dynamical properties of the SPA do not merely help us match and explain more data, but they also suggest specific behavioural and neurobiological experiments to run to test the architecture. In the former case, distinguishing the kinds of action that might take noticeably longer, or shorter, depending on their complexity. In the latter case, suggesting means of increasing and decreasing the length of time certain neurons in basal ganglia fire after actions are switched. Results from such experiments should provide a more detailed understanding of the neural underpinnings of cognitive control.

In general, the central theme of this chapter, ‘control’, is tightly tied to neural dynamics. This is unsurprising, since ‘control’ in most technical disciplines is a dynamical concept. Although I have discussed control more in terms of flexibility of routing information because of my interest in presenting an architecture that can manipulate complex representations, the dynamical properties underlying such control are ever-present and unavoidable. Measuring such dynamical properties is simpler than measuring informational properties, because *when* an event occurs (such as a spike or a decision) is more amenable to quantification (using a clock) than determining *what* such an event represents – but the ‘when’ and the ‘what’ *together* determine behavior. As a result, the dynamics of such processes provide a kind of independent measure of the underlying information processing being performed. Since we are interested in architectures that realize such processes, those that can be constrained by both dynamical and informational properties are subject to stricter constraints, and hence should be more convincing if they can meet both. This is why the ability of the SPA to contact these temporal constraints, while performing interesting information processing, is an important strength of

the approach.

There are, of course, many more kinds of temporal constraints than those I have discussed in this chapter. In the next, I will consider those related to learning and memory. However, though the behavior under consideration is different, the theme will be the same: the SPA seems to be in a position to provide a uniquely biologically realistic account of the detailed dynamical and information processing processes underlying cognition.

## 5.8 Nengo: Question answering

- Theoretical point: Decoding/transforming structured representation
- also large scale model construction and networks of networks
- do tutorial following terry's model for bbs/cog sci