# A UNIFIED MODEL OF RULE-SET LEARNING AND SELECTION

by

**Pierson Fleischer**

**A Dissertation**

*Submitted to the Faculty of Purdue University*

*In Partial Fulfillment of the Requirements for the degree of*

**Doctor of Philosophy**


Department of Psychological Sciences

West Lafayette, Indiana

December 2018

# THE PURDUE UNIVERSITY GRADUATE SCHOOL
# STATEMENT OF COMMITTEE APPROVAL

Dr. Sébastien Hélie, Chair

Department of Psychological Sciences

Dr. Gregory S. Francis

Department of Psychological Sciences

Dr. Richard Schweickert

Department of Psychological Sciences

Dr. Shawn W. Ell

Department of Psychology, University of Maine

**Approved by:**

Dr. David Rollock

Head of the Graduate Program

# TABLE OF CONTENTS

# LIST OF TABLES

# LIST OF FIGURES

# ABSTRACT

Author: Fleischer, Pierson. PhD
Institution: Purdue University
Degree Received: December 2018
Title: A Unified Model of Rule-Set Learning and Selection
Committee Chair: Sébastien Hélie

The ability to focus on relevant information and ignore irrelevant information is a fundamental part of intelligent behavior. It not only allows faster acquisition of new tasks by reducing the size of the problem space but also allows for generalizations to novel stimuli. Task-switching, task-sets, and rule-set learning are all intertwined with this ability. Naturally there are many models that attempt to individually describe these cognitive abilities. However, there are few models that try to capture the breadth of these topics in a unified model and fewer still that do it while adhering to the biological constraints imposed by the findings from the field of neuroscience. Presented here is a comprehensive model of rule-set learning and selection that can capture the learning curve results, error-type data, and transfer effects found in rule-learning studies while also replicating the reaction-time data and various related effects of task-set and task-switching experiments. The model also factors in many disparate neurological findings, several of which are often disregarded by similar models

# INTRODUCTION

The world is not random. The way something looks or feels or smells can provide cues about its possible uses. An example of this would be the knowledge that the hard inedible things inside of fruits are seeds. Objects that share the same features can often be used in the same ways. To keep up with the previous example, that knowledge about seeds could help a person identify the seeds of fruits that they had never encountered before so they could grow more. An intelligent entity can learn these consistencies and use them to build associations between an object's observed properties and its behavior. These associations then allow the agent to predict the behavior of new stimuli that share relevant features with familiar stimuli. Relevant factors can also include previous actions taken and other broader contexts such as goals. These associations are called rules. They cover a broad spectrum of conditional relations that can be as specific as "Don't eat red mushrooms with white spots" to as broad and as multifaceted as "Be polite as possible (and all the specific behaviors that entails) when interacting with customers". The first example relates a specific stimulus feature to a specific action while the second specifies a whole range of different and otherwise unrelated actions to be associated with some internal classification of customer (versus, for instance, coworker). Within the actions that make up polite behavior additional rules specify when each action should be taken. A collection of rules that are mutually exclusive and thus only one rule in the collection can apply to any given stimulus is a rule-set.

## Rules

Rules can cover a wide variety of situations and circumstances. In the simplest case they directly indicate a response. This kind of rule will be referred to as concrete

rules and they directly select a response/category, usually through discrimination of one or more features of a stimulus. Concrete rules can discriminate based on a single feature or a subset of features that is shared across a group of stimuli, e.g. “red stimuli are part of category B” which will apply to multiple stimuli of different shapes as long as all of them are red. A concrete rule-set can also discriminate using multiple features of a stimulus, an example of this being “When I encounter a red traffic light while driving I should press the brake pedal”. In this case, there are two relevant stimulus features and one state-based condition: 1. traffic signal, 2. red light, 3. while driving. In both cases these rules apply to multiple stimuli which include stimuli that have never been encountered before, such as traffic lights in a new city. This means that when a novel stimulus is encountered learning can progress very rapidly or even be skipped altogether, so long as the rule-set is still valid (Helie et al. 2015).

In contrast to concrete rules, the term “abstract rule” is used to define rules that do not select specific responses to stimuli and instead are used to select a set of concrete rules when there are multiple competing sets of concrete rules that could be applied to the stimulus. A good example of this in action is when stimuli can be categorized using either one of two dimensions, for instance color or shape, and there is a set of rules that have been learned for the stimulus for the specific values that have been encountered within those dimensions. A cue, such as background color indicates which rule-set and, by extension, which dimension should be used to categorize a given stimulus. An abstract rule-set associates each of the possible background colors with the correct set of concrete rules. Subjects learn response associations more quickly when the problem space is organized such that it can be represented with a combination of concrete rule-sets that are

selected by abstract rules compared to situations where abstract rules are not useful (Badre, Kayser, & D'Esposito, 2010). This shows that humans will spontaneously take advantage of these kinds of patterns in the environment, even when they do not have any prior knowledge about the usefulness of rules in the current environment.

Not every physical or cognitive response to stimuli is a rule. There is a difference between changes that can be described by rules and changes that occur because of rules. An automatic response, like shifting attention to a surprising stimulus, is not a rule even though a rule can describe it. In this thesis, we define rules as being volitional and accessible to conscious awareness.

## Physiological Basis for Rules

The prefrontal cortex (PFC) is commonly understood as the location of high-order processes and cognitive control (Miller, Freedman, & Wallis, 2002). Rules, especially abstract rules, can be categorized as such and thus investigations into the location of rules should point towards the PFC. An fMRI study by Helie, Roeder, and Ashby (2010) found that subjects that successfully learned the rule-sets of a rule-based task had greater activity in the PFC than subjects that did not. In another fMRI experiment, Badre and Wagner (2004) used a task that, on each trial, primed a particular response but did not always end up asking for that response. In trials with a mismatch between the primed response and the requested response subjects had to use rules to resolve the conflict between responses. In these cases the fMRI results showed increased activity in the dorsolateral PFC (DLPFC).

Non-fMRI studies also show that the PFC has a strong association with rules. Damage to the PFC is known to cause deficits in performance in the Wisconsin Card

Sorting Task (WCST - described in Section 1.3) (Bunge, 2004; Fuster, 2008), which relies on rule-sets and rule-set switching. In 2001 Wallis et al. published an article that used the same-different task (described in Section 1.3) where monkey subjects respond to either matching (same) or non-matching (different) stimuli. Single cell recordings taken during the task found that neurons that exhibited rule-selective firing patterns were more prevalent in the DLPFC than neurons with any other kind of task-related firing pattern. Other studies have shown impairments in the ability to switch between rule-sets in subjects with lesions in the PFC (Koechlin, Ody, & Kouneiher, 2003, O'Reilly, Noelle, Braver, & Cohen 2002, Badre, Hoffman, Cooney, & D'esposito 2009). Further supporting evidence of the localization of rules in the PFC can be found in a review article by Duncan (2001).

## Behavioral Investigations of Rule-Set Switching

Rules are such a flexible and potent tool that they have naturally garnered substantial interest amongst cognitive psychologists. A particularly common area of study involving rules is task-switching studies (Stroop, 1935; Grant & Berg, 1948; Spector & Biederman, 1976; Sudevan & Taylor, 1987; Rogers, & Monsell, 1995; Hübner, Futterer, & Steinhauser, 2001; Logan & Bundesen, 2003; Monsell & Mizon, 2006; Koch & Allport, 2006). In this paradigm, subjects switch between well-defined tasks such as even/odd judgements or color naming. Naturally, each task has its own response to a given stimulus. Since a rule-set is a collection of rules and rules are stimulus/response associations each task therefore uses its own rule-set. That way the rule-set for the color-naming task can have the same response assigned to a red 4 and red 7 while the even/odd task set will have different responses assigned to those stimuli.

Subjects rarely make errors in these familiar tasks but the time it takes to make a response varies based on a number of factors. Experiments of this sort go back to 1927 when A.T. Jersild published a paper examining reaction times across a variety of tasks such as adding and multiplying. The data showed the now classic results of faster response times when the subjects are repeating a task and longer response times when the task to be performed changes. Since then numerous studies have been conducted investigating the dynamics of switching rule-sets (Stroop, 1935; Grant & Berg, 1948; Spector & Biederman, 1976; Sudevan & Taylor, 1987; Rogers, & Monsell, 1995; Hübner, Futterer, & Steinhauser, 2001; Logan & Bundesen, 2003; Monsell & Mizon, 2006; Koch & Allport, 2006).

In 1995, Rogers, & Monsell examined various methods of mitigating switch costs. They found that, while some reduction is possible, at a certain point the switch cost cannot be reduced any further.It appears that some reconfiguration can only be done once the stimulus has been presented. In 2001, Hübner, Futterer, & Steinhauser expanded upon this finding and determined that the switch-cost asymptote for tasks with more cognitive steps was larger than the asymptote for simple tasks. In studies like these, switch costs are hypothesized to be caused by some task-reconfiguration process. In 2003 Logan & Bundesen showed that controlling for task repetition eliminates the switch-cost, which instead suggested that task repetition caused a priming effect. However, this evidence was weakened when Monsell & Mizon (2006) showed that switch-costs reappear when the probability of a task-switch is low. A connectionist model by Reynolds, Braver, Brown, & Van der Stigchel (2006) was used to simulate the residual switch-cost after priming. In the model, the rule-set is selected by the cue and stored in the PFC.

Dopamine probabilistically maintains the rule-set until the stimulus is presented. The stimulus and rule-set (if maintained) form a compound representation which then determines a response. Reynolds, Braver, Brown, & Van der Stigchel (2006) compared reproduced reaction times and neural activity gathered from subjects using fMRI.

Whether or not priming can entirely account for switch costs, its effects on task-switching response times are undeniable. This is demonstrated in a 1987 paper by Sudevan & Taylor wherein subjects were given a preview of the task cue before the cue was presented. The preview was not always accurate however. Sometimes the preview would indicate the wrong cue or be ambiguous (neutral). When the preview correctly predicted the upcoming task cue, subjects' reaction times were faster than with the ambiguous preview. Conversely, when the preview predicted the wrong task cue, the upcoming task reaction times increased from neutral. When subjects are aware of what the new task will be, for instance when tasks alternate regularly, longer delays between stimulus presentations also reduce reaction time (Rogers, & Monsell, 1995).

Switching tasks is not the only way to slow reaction times or interfere with performance. In 1935 Stroop published a famous paper showing that when one task is far more well-rehearsed than the other, the common task will interfere and slow reaction times in the uncommon task. This interference can also occur from tasks that subjects expect to perform, as was seen in Sudevan & Taylor (1987) where a misleading forecast slowed reaction time. Alport, Styles, and Hsieh went farther when, in 1995, they presented evidence that switch costs are the result of interference from past tasks and not some task-switching process. The interference interpretation of switch costs was used in

an ACT-R model presented by Altmann & Gray in 2008. This model was able to replicate reaction time and error effects related to task-switching and repeated task runs.

Spector & Biederman (1976) found that switch costs are particularly severe when there is no external indication of which task to perform and subjects have to keep track of previous tasks to discern which task should be performed. In the same experiment they found that when the stimulus has only one task associated with it switch costs are minimal. In a situation like this, the stimulus itself acts as the cue for the task. Rogers, & Monsell (1995) expanded upon this and showed that adding distractor features to the stimuli which cued the incorrect task slowed subjects' reaction times. A similar approach was used by Koch & Allport in 2006 where every task was compatible with each stimulus but each stimulus was paired with only one task. When this pattern was broken reaction times increased. In 2015 van 't Wout, Lavric & Monsell used a similar experiment to show that, when task frequency and recency are controlled, subjects that switched between five tasks had switch costs equal to subjects who switched between three. This suggests that the results from Koch & Allport were not due to a larger number of tasks that could be paired with a given stimulus but from a stimulus-task pairing that was unfamiliar and unexpected.

All of these experiments use tasks and cues that are already familiar to the subjects. Of course this familiarity had to be developed somehow. How do subjects behave in situations where they do not yet know what the 'correct' response should be? Consider the Wisconsin Card Sorting Task (WSCT) (Grant & Berg, 1948). In this task, subjects attempt to match stimulus cards in one of the three dimensions: object shape, object color, or number of objects. While the tasks are familiar there is no cue informing

subjects which task is correct. Task switching is rare and, through trial and error, the subject learns which dimensional rule-set they should follow.

Performance in a similar task, called the intradimensional/extradimensional task was modeled in O'Reilly et al. (2002). In this task, subjects are presented pairs of cards that vary in two dimensions. Subjects chose a card based on the features in one of the dimensions and the cards can never both have the same feature at the same time. When the target feature changed it could switch to the other feature in the same dimension (interdimensional shift) or could switch to a feature from the other dimension (extradimensional shift). Lesions in specific brain areas causes varying levels of perseveration when switches occur. The model used a connectionist architecture where the rules were hard-coded in the model and the rule-sets were selected using dopamine as the reward signal.

Another model using un-cued task switching was presented in Helie, Ell, Filoteo, & Maddox (2015). The task used in this article asked subjects to categorize lines that varied continuously in two dimensions, length and orientation. Only length defines a stimulus's category, orientation is irrelevant. Halfway through the experiment any combination of these three things occur (including none of them or all of them): 1. A shift in length, increasing the length of the stimuli and categories by a large amount. 2. A shift in orientation, increasing the orientation of the stimuli by a large amount. 3. Relevant dimension shift, henceforth the categories are determined by orientation, not length. The model used the computational cognitive neuroscience approach (Ashby & Helie, 2011) with rule-sets encoded in the PFC and rule-set switching governed by random-walk. Of note is the method used to integrate the stimulus features and the rule-sets. Rather than a

composite layer where each unique pair of rule-set and feature had a neuronal representation, the rule-sets adaptively gate the connections from the features to the response. The reward signal is provided by feedback sensitive cortical cells. The notable result that this model reproduced is that when a relevant dimension shift was accompanied by a length-shift, the new categories were acquired faster than when a relevant dimension shift was not paired with a length-shift.

Another useful task used in the investigation of rule-set learning is the "Same/Different" task (Bamber, 1969). In this task a subject is presented with an image and cue (such as a tone) informs the subject whether to respond when two sequentially presented images are the same (identical match) or different. This task is common because it can be performed by animal subjects and therefore allows the experimenters to use more invasive methods of investigation. Most abstract rule-based tasks can only be performed well by humans and thus the focus will be on results from human studies.

An example of more extensive learning can be found in an experiment by Badre, Kayser and D'Esposito (2010) subjects were asked to learn associations between 3 responses and 18 stimuli which varied across 3 dimensions (with 3x3x2 discrete values). In one condition the associations were "flat" with no pattern to them and no abstract rule-set that would aid learning. In the other condition, the associations were "hierarchical" – the dimension with two values could be treated like a cue, signaling which of the other two dimensions were irrelevant for that trial. Subjects were not informed that any hierarchy existed, they were simply told to learn the correct response for each stimulus. Despite this obfuscation, subjects in the hierarchical condition learned faster and achieved better performance by the end of the experiment compared to subjects in the flat

condition showing that subjects try to apply rule-sets to situations even when there is no prior indication that rule-sets would be useful.

In 2013 Collins and Frank conducted a two phase experiment where subjects had to learn associations between two-dimensional stimuli and four responses. In the first phase there were two possible values in each of the two dimensions (2x2) and each feature combination was associated with a different response. In phase two new values were added in one dimension and the new combinations were assigned responses so that the subjects' performance would differ based on the way the subject represented the environment. For instance, using color as a cue for an abstract rule-set (color as context) would allow subjects to transfer knowledge while using shape as a context would not, resulting in different learning curves between these two groups of subjects. Collins and Frank simulated these two representations using a connectionist model in which rule-sets were associated with cues using a cortico-basal ganglia loop. Another cortico-basal ganglia loop associated responses with stimulus feature/rule-set pairs. The simulation results showed that some subjects built color-as-context representations while others built shape-as-context representations. Task representations other than these two where uncommon to the point of being unidentifiable, suggesting that humans naturally organize information hierarchically.

## Summary and Outline

Concrete rules are any association between a stimulus or stimulus feature and a response. A rule-set is a collection of rules that are not mutually exclusive. Rules in a rule-set typically operate on a single stimulus dimension. Abstract rules are associations between a cue and a concrete rule-set. Two common experimental paradigms are

shift-cost investigations using task-switching and rule-set acquisition and learning. The next chapter reviews some existing biological constraints related to rule-set learning and point out common ways that models violate these constraints. The physiological basis for a new model is then explored and the overarching theory of the model is presented. In chapter 3 this theory will be implemented using a quantitative model that will then be used in chapters 4 through 7 to reproduce a variety of learning and reaction-time results. Chapter 8 will present an overall evaluation of the model and discuss ways the model might be extended and future work that could be done. The next chapter will describe some important aspects of the brain and hypothesize about the implications these features have for rule-set switching and rule-set learning.

# A NEW THEORY OF RULE-SET REPRESENTATION AND LEARNING

A good model needs to do a great many things. A model needs to replicate the empirical data, but a good model should also make interesting predictions about future experiments. If a model makes absurd or impossible predictions outside of the data it was built on it is a very limited model. One way to create interesting predictions is to link the features of the model to the characteristics of the real situation. Ideally, the model will continue to produce accurate results even when outside of the environment it was designed for. Furthermore, if a plausible model fails to make accurate predictions, being able to identify the source of the errors in the model can indicate aspects of the real problem that may be incorrectly understood. Realistic model also produce predictions about the real world and allow the model to be more easily expanded to account for other categories of data. For these reasons, it is a good idea for the model to account for as much biological and neuroscientific data as possible (Ashby & Helie, 2011; Newell, 1992). However, there is a plethora of biological and neurological information about the brain and no model can incorporate all of it. The most important features to include are those that have the greatest functional impact on the model and those that most firmly localize the components of the model.

The goal of this research is to create a model that can make use of cues and multidimensional stimuli to learn task structures. The model should use feedback to learn concrete and abstract rules, populate rule-sets, and select rule-sets using cues. Creating new rule-sets is outside the scope of the model. Using multiple abstract rule-sets is

inferred by the architecture of the model but currently remains unexplored. There are no mechanisms designed specifically to reproduce response-time effects such as priming and switch-costs but the model should be able to selectively maintain recently used rule-sets in working memory and thus reproduce the basic effects of these manipulations.

## Feedback and Learning in the PFC

Most rules are not hard-coded into the brain. They are learned through environmental exploration (any rule-set learning task that shows a learning curve, for example see Badre, Kayser & D'Esposito 2010 or Collins & Frank 2013). Any model that tries to model rule-set learning needs to incorporate some form of plasticity. For a biologically plausible model, this plasticity also needs to be linked to some adaptive process in the brain. The obvious first candidate for this process is synaptic plasticity. Signals are transmitted between cells through chemical signaling agents called neurotransmitters. Neurons that recently underwent an action potential release these neurotransmitters, which bind to receptors on other neurons. When these receptors are triggered, they alter the membrane potential of their host cell, which affects the likelihood that it will undergo an action potential of its own. These chemical signals are transmitted at specialized structures called synapses. The magnitude of the neurotransmitter's effect on the postsynaptic cell is not fixed. A single action potential of the presynaptic cell can make only a small difference in the membrane potential of the postsynaptic cell or it singlehandedly could induce an action potential (through the release and capture of neurotransmitters). Some neurotransmitters inhibit the postsynaptic cell instead, hyperpolarizing the cell and making action potentials less likely to occur. The properties of synapses can be permanently adjusted to change the strength of the signal received by

the postsynaptic cell. This modification can occur through a variety of mechanisms such as increasing or decreasing the number of neurotransmitter receptors on the membrane of the receiving neuron (Kida, Josselyn, de Ortiz, Kogan, Chevere, Masushige, & Silva, 2002; Chklovskii, 2004). When a synapse is modified to increase the strength of the signal, it is called long-term potentiation (LTP) and when it is decreased, it is called long-term depression (LTD). While synapses do weaken over extended periods of disuse, this is beyond the scope of a model of trial-level task learning. Therefore, the focus will be on the synaptic plasticity involving two active neurons. In some brain areas, such as the PFC, the determining factor between LTP or LTD occurring is the degree of recent activity in the post synaptic cell (i.e., Hebbian learning, Ashby & Helie, 2011). In these areas if the activity of the post-synaptic cell is weak (but not completely inactive) LTD will occur. If the post-synaptic cell activity is strong enough, LTP will occur instead. Hebbian learning is useful for reinforcing behavioral patterns that are already partially established (e.g. Reynolds, Braver, Brown, & Van der Stigchel, 2006). However, it does not have any inherent relation to reward and thus cannot use positive feedback to establish new behaviors. In other areas (e.g., the striatum), plasticity is mediated by dopamine, which is associated with reward and feedback. In these areas, if cell activity is high enough for modification to occur, high levels of dopamine will cause LTP while low levels will cause LTD.

It is no surprise that synaptic plasticity is the learning mechanism of choice for neurological models of rule-set acquisition. However many of these models, such as O'Reilly et al. (2002), use dopamine as the reward prediction error signal for LTP and LTD in the PFC. The problem with this approach is that the PFC lacks significant

quantities of the protein DAT which facilitates the rapid re-uptake of dopamine (Seamans & Robbins, 2010). Without this protein dopamine lingers in the synapse and thus cannot provide a signal with the temporal resolution that LTP and LTD would require for the synaptic modifications to be meaningful and improve task performance (Cass & Gerhardt, 1995). This is problematic for any model that relies on dopamine signals in the PFC that can be differentiated across trials, such as Reynolds et al. (2006), not just those that use dopamine as a reward signal to direct learning.

To avoid this issue, other models, such as the one use by Collins and Frank (2013), incorporate the cortico-basal ganglia loop and have the entirety of the reward-modulated learning take place in that brain area, where the presence of DAT allows dopamine to function as a temporally precise reward signal. While studies have not ruled out this possibility, adopting this model consigns all trial-level learning to the cortico-BG loop, and for the most part to just the striatum. This despite the PFC being much larger than the BG. A third, less restrictive option is to incorporate feedback cells that become active when either positive or negative feedback is received. These feedback cells are based on an fMRI experiment by O'Doherty, Kringelbach, Rolls, Hornak, & Andrews (2001) which found areas of the PFC that became more active when reward was received and less active when punishment was received. The magnitude of these changes was proportional to the intensity of the reward/punishment received. In other areas of the PFC, they found reward and punishment had the inverse effect on neuronal activity, with punishment increasing activity and reward decreasing it.

In the PFC feedback representation approach, feedback activates particular populations of cells in a model. The outgoing connections of these cells target other cells,

exciting some and inhibiting others. This pattern of additional excitation and inhibition affects the firing pattern of the targeted cells and, due to the way LTP and LTD are triggered, whether the cells' synapses undergo LTP or LTD. Some feedback cells are activated by positive feedback and others by negative feedback and the pattern of connections differs between these two populations. The end result is that positive feedback influences Hebbian learning between cells in such a way that the response chosen for the current stimulus will be more likely to be chosen again should the same stimulus be encountered. Negative feedback would instead make the chosen response less likely to be selected the next time the stimulus is encountered.

### Presynaptic Inhibitory Gating

The ability to combine rule-sets and stimulus features creates a vast combinatorial space for varied dynamic responses. However, this space cannot be represented at the neuron level. Integrating the combination of active rule-set and current stimulus to select a response using only standard cell-to-cell plastic connections would require some large array of intermediate cells. This array would become unfeasibly large for even moderate numbers of rule-sets and stimulus features. Models such as those by O'Reilly et al. (2002), Reynolds et al. (2006), Altmann & Gray (2008), and Colins & Frank (2013) would fall victim to this problem if the number of tasks or the number of possible responses within each task were increased to real-world levels. The Colins & Frank (2013) model is especially flawed in this respect because the combinatorial explosion occurs in the striatum where there is limited room for it. An alternative approach that does not require an inordinate number of neurons would be to use plastic synaptic connections of a special type, which deliver what is known as pre-synaptic inhibition

(Shepherd, 2004). Pre-synaptic inhibition involves an inhibitory synapse in which the post-synaptic component is not the dendritic terminals of a post-synaptic cell as is common. Instead, the inhibition is received at the axon terminals of the post-synaptic cell and affects only the specific synapse that the terminal is part of. Inhibition received this way does not affect the spiking of the post-synaptic cell at large but instead blocks neurotransmitter release at the targeted synapse. Pre-synaptic inhibition acts like a gate, allowing cells that represent stimulus features to selectively activate only a subset of the cells that they connect to depending on the active rule-set.

Because of the unusual effect of pre-synaptic inhibition, a different synaptic model is needed to implement these connections and they will feature prominently in the proposed model. However the term “pre-synaptic” is used extensively within the topic of neural connections and so, in an effort to avoid ambiguity, the term “non-somatic” shall be used from here on to refer to connections such as these that do not directly affect the spiking of the post-synaptic cell. Non-somatic inhibition was used in the Heterosynaptic Inhibitory Criterion Learning (HICL) model presented in Helie et al. (2015) and the architecture of the HICL model inspired a large part of this theory and the design of this new model. The biggest drawback with this model is that it does not capture cued rule-set switching and its mechanism for un-cued switching is not integrated into the connectionist network that makes up the rest of the model. A new model is proposed that will address these shortcomings.

## Rule-Set Maintenance in Working Memory

The fact that recently used rule-sets can interfere with the rule-set currently in use is evidence that some memory of previous rule-sets tends to be maintained. Ashby, Ell,

Valentin, & Casale presented FROST, a model of working memory maintenance in 2005. FROST is designed to reproduce the delayed-response task, which does not require learning. In this task, the subject (typically a monkey) watches as a reward (such as food) is hidden in one of two containers. The monkey has to remember which container the food is in until it is permitted to open the container. FROST represents the location of the reward using simulated PFC cells, which are activated using stimulus cells representing the visual input of the food being placed in a container. When the food is hidden these stimulus deactivate but the location is maintained in the PFC cells through a set of parallel bidirectional connections with the thalamus (Figure 1). Each PFC cell excites a thalamic cell and that thalamic cell excites the PFC cell in return. The maintenance is, by default, disabled but can be enabled through a connection from the PFC to the striatum. The striatum inhibits the interior segment of the globus palidus (GPi) which inhibits the thalamus. When the thalamus is inhibited the maintenance loop cannot be established. In this regard, the FROST model behaves counter to what the task-switching experiments would suggest. In order to model the unintentional interference from previously activated rule-sets these rule-sets must maintain some degree of activity past the removal of the cue and into the next trial. Since this interference is detrimental, it makes sense to assume that the maintenance mechanism is active by default. Otherwise why purposefully activate it in a situation where it is not helpful? Hence, maintenance should be enabled by default and require time to dissipate or suppress.

One way to make the thalamic maintenance loop enabled by default is to incorporate the exterior segment of the globus palidus (GPe) which also inhibits the GPi and, unlike the striatum, is naturally active and does not require external excitation to

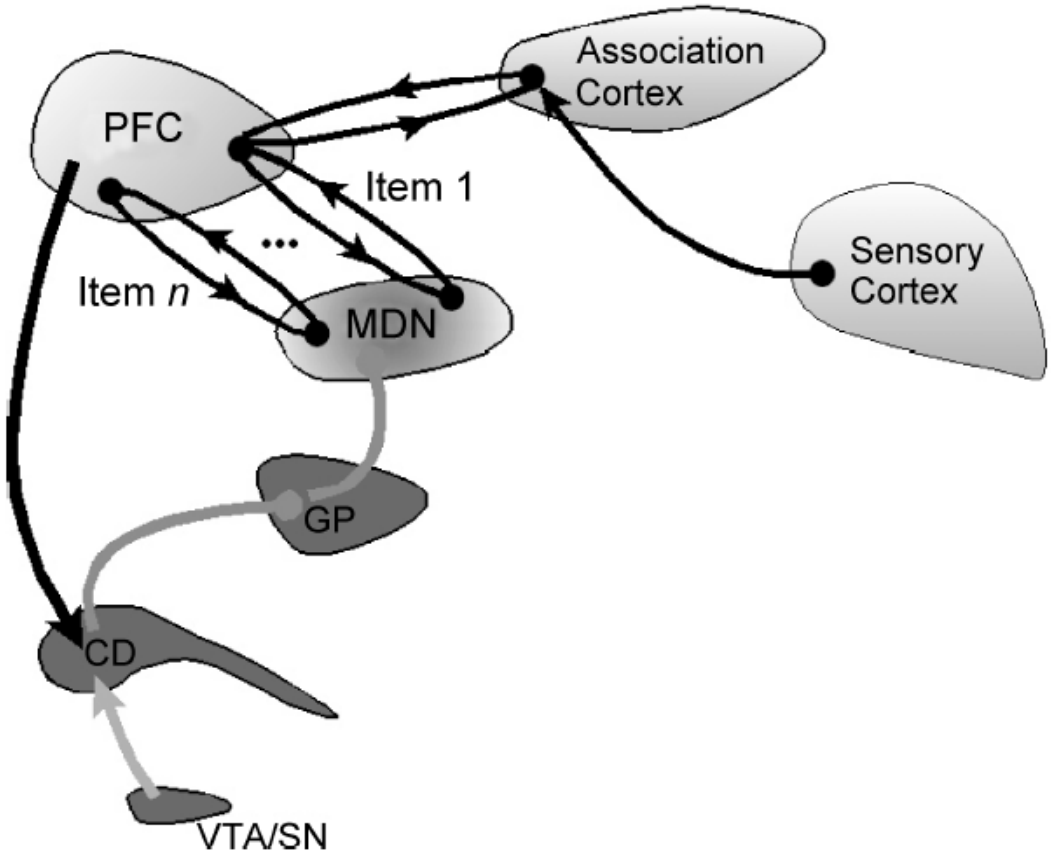


*Figure 1*. The FROST model architecture showing the pathway through the basal ganglia and the parallel maintenance loops between the PFC and the thalamus (from Ashby et al., 2005).

generate an inhibitory signal. With this, the thalamus is freed and the maintenance loop is enabled by default. However the old rule-sets are not maintained indefinitely so there must also be a way to halt the maintenance loop. Results published by Rushworth, Hadland, Paus, & Sipila (2002) and by Dove, Pollmann, Schubert, Wiggins, & von Cramon (2000) indicate that the pre-supplementary motor area (pre-SMA) is active when subjects engage in cue-mediated task switching. Rushworth et al. (2002) also shows that temporary lesioning of the pre-SMA negatively affects performance only in trials involving a task-switch. The pre-SMA excites the subthalamic nucleus (Aron, Behrens, Smith, Frank, & Poldrack, 2007), which excites the GPi. The GPi overcomes the inhibition from the GPe and inhibits the thalamus, disrupting the cycle of excitatory input between the thalamus and the PFC cells that represent rule-sets. Without the excitatory input from the thalamus activity in the PFC cells decay and the obsolete rule-set leaves working memory.

It is unusual for a model that includes the GPe and GPi to not include the striatum as well. The striatum receives input from most cortical areas and selectively inhibits both the GPi (direct pathway) and GPe (indirect pathway) (Alexander, DeLong, and Strick 1986). These connections can give some insight into the role of the striatum regarding working memory maintenance. The direct pathway enables maintenance while the indirect pathway prevents it. In tandem these pathways could dynamically constrain the items, such as task-sets, that can enter working memory, preventing distractions. This kind of functionality does not factor into the current battery of experiments and thus the striatum is omitted from this implementation of the model.

## Hierarchical Organization of the Prefrontal Cortex

Localizing various cognitive processes to specific brain areas is a very common topic of study for cognitive neuroscience. Achieving success does not put an end to research either. If a cognitive process is found to occur only in a specific part of the brain, psychologists and neurobiologists almost immediately start trying to narrow it down further, looking for a specific sub-area responsible for the process. The PFC is no exception.

A large amount of research suggests that the lateral prefrontal cortex (LPFC) is associated with tasks that require rule-sets and rule-set switching. A few examples of such research include an fMRI study by Koechlin, Ody, & Kouneiher (2003) which found that anterior areas of the PFC were associated with greater levels of abstraction (Figure 2). Crone, Wendelken, Donohue, & Bunge (2005) found more activity in the LPFC when abstract rules were used, when a new rule-set is selected, and when the next rule-set that would be used was not predictable. Similarly, Zanolie, Van Leijenhorst, Rombouts, & Crone (2008) found increased activity in the dorsolateral PFC relative to medial areas during various rule-set related processes. The greatest increase in activity was found when the subject made an error due to interference from irrelevant rule-sets - meaning the relevant rule-set needs to be reinforced. A lesser increase in LPFC activation occurred when a new rule-set was selected through trial and error from a finite set of possible rule-sets. The smallest increase in activity occurred when the environment changed and the currently active rule-set was found to no longer be applicable.

We can also look at cases where the anterior PFC is not fully functional. In a study by Bunge and Zelazo (2006) structural MRI scans were used to track PFC

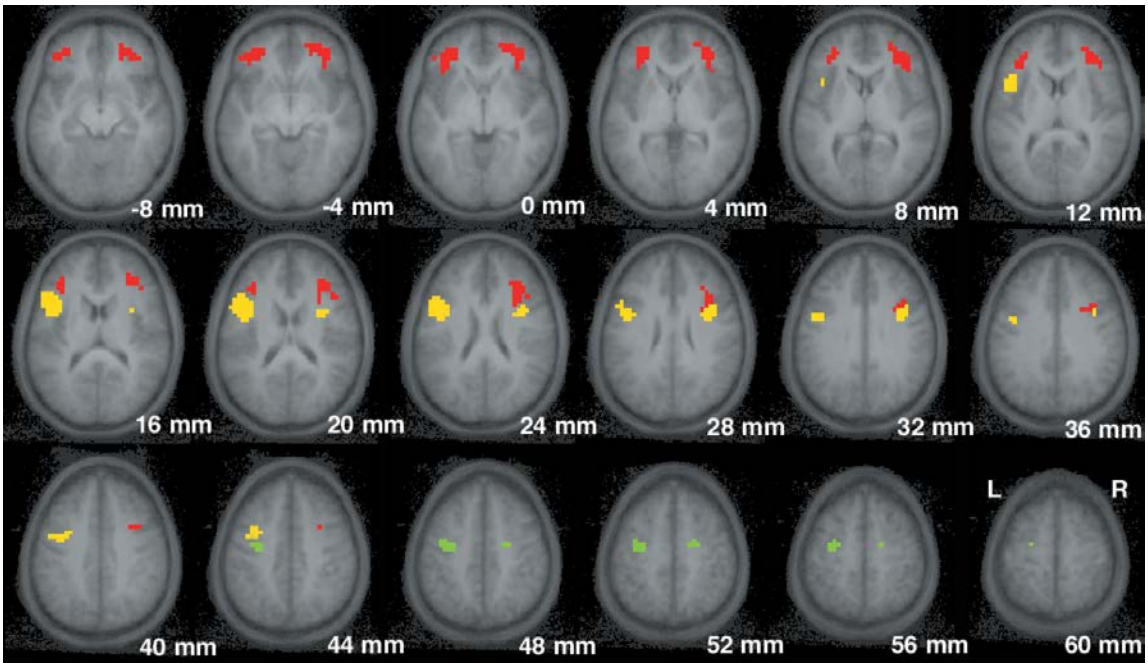


*Figure 2*. From Koechlin, Ody, & Kouneiher 2003. fMRI results showing the hierarchical organization of the PFC. Areas in green were more active during tasks with more possible responses (large concrete rule-sets). Areas in yellow were more active during tasks with more possible cues – each corresponding to a different rule-set (many concrete rule-sets). Areas in red were more active in tasks were the stimulus-response and cue-rule-set associations conflicted with a larger percentage of other task blocks (uncommon abstract rule-set).

development in young children. Over the same period of time their ability to use rules is evaluated through a set of behavioral tasks. Anterior areas of the PFC develop later than posterior areas and as they develop children become more capable of using rules in complex ways. Children go from being able to use only a single rule-set with no ability to switch between rules to some ability to switch so long as the rules are complementary. At this point in development, the children are still unable to switch between conflicting rules. Badre et al. (2009) found that patients with lesions in the PFC that perform categorization tasks become impaired as contingencies are added to the task. The more anterior the lesion the more abstract the contingencies have to be before impairments present. Other supporting evidence can be found in review papers by Badre (2008) and Buckner (2003).

Together these findings suggest a hierarchical design where concrete rule-sets are represented in posterior PFC regions and abstract rules are more anterior. This creates a tiered system where physical responses are coded in motor and premotor area. Anterior to these motor-related areas is the posterior PFC where the concrete rule-sets that govern the selection of motor responses are represented. Anterior to that is the anterior PFC, location of the abstract rule-sets, which selects from among the concrete rule-sets. Combined, these findings form a new theory of rule-set representation and learning.

## The Theory

The literature reviewed above suggests that rule-sets might be maintained as items in WM. Brain activity related to rule-set application is also co-located with brain activity related to WM. Hence, this theory proposes that rule-sets are represented as regular WM items that can be selected, switched between and need to be maintained. Rule-sets that

are more abstract are represented in the anterior areas of the PFC. The integration of rule-sets and stimulus features relies on adaptive gating and presynaptic inhibition. Lastly, learning should not take place entirely in the BG and learning taking place in the cortex cannot use dopamine as the reward signal.

More specifically, highly processed stimulus features are represented in the PFC as reported in papers such as Freedman, Riesenhuber, Poggio, & Miller (2003) and Wallis, & Miller (2003). These stimulus cells excite cells in the premotor cortex (PMC) that encode various responses (Figure 3[1]). These excitatory connections are gated by concrete rule-set cells in the posterior PFC via non-somatic inhibition (Figure 3[2]). The stimulus features encoded in the PFC include those features that cue rule-set selection and thus the stimulus features also excite the concrete rule-set cells (Figure 3[3]). The abstract rule-set cells gate the input to the concrete rule-set cells in the same way the concrete rule-set cells gate the input to the PMC cells (Figure 3[4]). The concrete rule-set cells are paired with thalamic cells that mutually excite each other to maintain an active rule-set (Figure 3[5]). This maintenance loop can be broken by activating the GPi, which inhibits the thalamus (Figure 3[6]). Without outside excitation the GPi is prevented from disrupting the maintenance loop by inhibition from the GPe (Figure 3[7]). Reward-selective feedback cells in the PFC provide excitation or inhibition based on the outcome of the selected response (Figure 3[8]). This reward signal influences the effect of Hebbian Learning at the non-somatic inhibitory connections resulting in gradually improving performance. The next chapter describes a complete model implementation of this theory.

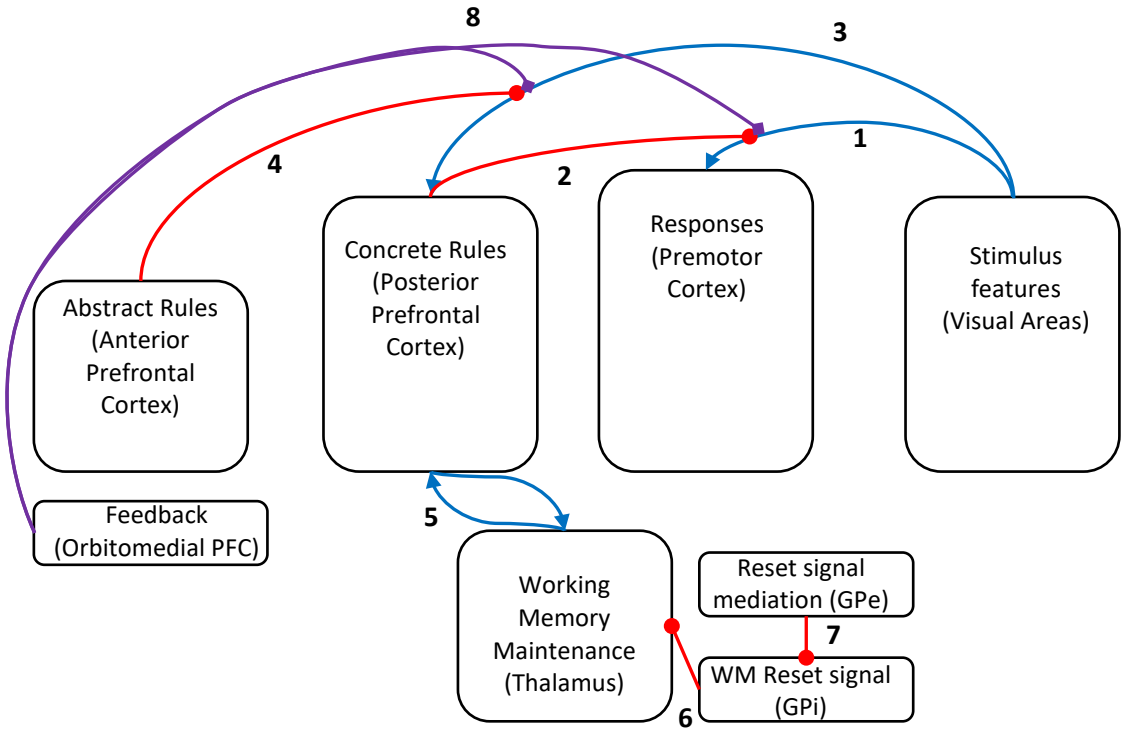


*Figure 3*. The broad connectivity diagram of the proposed theory. Blue connections represent excitatory connections and red connections represent inhibitory connections. Purple connections are used only with the feedback area (orbitomedial PFC) and represent both excitatory and inhibitory connections. Numbered areas are described in the text.

# THE MODEL

## Cortical Pyramidal Cells

The cells' membrane potential is modeled using the following equations, parameterized to model a cortical pyramidal cell (Izhikevich 2007):

$$V_a(t+1) =$$

$$V_a(t) + [0.7[V_a(t)+60][V_a(t)+40] - U_a(t) + E_a(t) - I_a(t) + \varepsilon(t)]/100$$

$$U_a(t+1) = U_a(t) + 0.03\{-2[V_a(t)+60] - U_a(t)\}$$

where $V_a(t)$ is the membrane potential (in mV) of cell *a* at time *t*, $U_a(t)$ is the value of the recovery variable (a dimensionless representation of the slow ionic current) of cell *a* at time *t*, $E_a(t)$ is the excitatory (glutamate) input to cell *a* at time *t*, and $I_a(t)$ is the inhibitory (GABA) input to cell *a* at time *t*. $\varepsilon_a(t)$ is the noise in cell *a* at time *t* and is drawn at each time-step from a Gaussian distribution with a mean of 0 and a standard deviation of 200. The noise and both inputs are represented as mA of injected current. If the above equation would result in a $V_a(t)$ greater than *35* then *a* is said to have spiked and the following adjustments are made: $V_a(t)$ is set to *-50* and $U_a(t)$ is increased by *100*. Additionally, the spike time *t* is concatenated onto the end of $S_a$, the list of the times at which cell *a* spiked. Each cell also generates output at each time step according to the following equations.

$$O^{Glut}{}_a(t) = \sum_{s \in S_a} \left[\frac{t-s}{60} e^{1-\left(\frac{t-s}{60}\right)}\right]^+$$

$$O^{GABA}{}_a(t) = \sum_{s \in S_a} \left[\frac{t-s}{30} e^{1-\left(\frac{t-s}{30}\right)}\right]^+$$

where $O^{Glut}{}_a(t)$ is the glutamate output of cell *a* at time *t* and $O^{GABA}{}_a(t)$ is the GABA output of cell *a* at time *t*. This is an abstraction of how neurons work in the brain. Biological cells release the same mix of neurotransmitters at every one of their synapses. Within the model, some pyramidal cells have both excitatory and inhibitory outgoing connections. These cells are stand-ins for groups of cells that include excitatory pyramidal cells and inhibitory interneurons. $[f(t)]^+$ equals $f(t)$ when $f(t) \geq 0$, and 0 when $f(t) < 0$. Glutamate and GABA release are modeled with separate equations because the release and re-uptake of GABA is faster than that of glutamate (Helie et al. 2015). The outputs of cells are transmitted to other cells through connections. Excitatory connections are denoted with a solid arrow and inhibitory connections with a hollow arrow: $\underset{s}{\rightarrow}$ ; $\underset{s}{\Rightarrow}$. Standard cell-to-cell connections are indicated with the letter "*s*" below the arrow: $\underset{s}{\rightarrow}$ and non-somatic connections with an "*ns*": $\underset{ns}{\rightarrow}$. The effect of connections on the E and I terms of the voltage equation are described with the following equations:

$$E_a(t) = \sum_{b \epsilon C} w_{(b)\underset{s}{\rightarrow}(a)} \left[ O^{Glut}{}_b(t) + P_{(b)\underset{s}{\rightarrow}(a)}(t) \right]^+$$

$$I_a(t) = \sum_{b \epsilon C} w_{(b)\underset{s}{\Rightarrow}(a)} \left[ O^{GABA}{}_b(t) + P_{(b)\underset{s}{\Rightarrow}(a)} \right]^+$$

$$P_{(a)\underset{s}{\rightarrow}(b)}(t) = \sum_{c \in C} O^{Glut} \cdot w_{(c)\underset{ns}{\rightarrow}\left((a)\underset{s}{\rightarrow}(b)\right)} - \sum_{c \in C} O^{GABA} \cdot w_{(c)\underset{ns}{\Rightarrow}\left((a)\underset{s}{\rightarrow}(b)\right)}$$

where $w_{(b)\underset{s}{\rightarrow}(a)}$ is the weight of an excitatory connection from cell a to cell b and $w_{(a)\underset{ns}{\rightarrow}\left((b)\underset{s}{\rightarrow}(c)\right)}$ is the weight of an excitatory non-somatic connection from cell a to the excitatory synapse from cell b to cell c.

Some of these non-somatic connections, like those from the feedback cells, are likely to be located on the axons of the stimulus cells because one weight applies to multiple outgoing synapses of the post-synaptic cell. Other non-somatic connections, like those from the rule-set or premotor cortex (PMC) cells in the model (as detailed later in this thesis), have varying effects across the synapses of the post-synaptic cell. Biologically, the non-somatic connections that cause these varying effects could target either the axon near the terminals or individual synapses themselves. Within the model, those non-somatic connections targeting individual synapses correspond to pre-synaptic inhibition. Whatever the true biological implementation, they can be modeled in the same way because the only effect that the location of the synapse has is different travel times of the action potential. Since the model treats this travel as being instantaneous synapse location is irrelevant.

If there is no connection of a given type between two model components, the corresponding *w* term can be treated as though it were 0. Within the model each connection has its own mean value for its weight term. When the model is initialized the starting value of the term is drawn from the uniform distribution of all numbers within ±1% of the corresponding mean unless the postsynaptic cell is an accumulator (as explained next).

## Non-Cortical Cells

While most of the cells in the model are parameterized as Cortical Pyramidal cells, the model also makes use of thalamic and globus pallidus cells. All cells use the same equations for *O, E, I, P, and S.* However, each cell type has unique patterns of firing and responds differently to excitation. Therefore, in order to properly model these

types of cells, different equations have to be used for V and U and the effect of an action potential must be parameterized differently. Fortunately, Izhikevich (2007) has already calculated these parameters. In cases where a parameter has a value that has no effect on the rest of the equation, such as a factor having a value of 1.0, it is intentionally included in the equation to make the relation to the general form clearer.

The voltage update equation for thalamic cells is:

$$V_a(t+1) = V_a(t) + [1.6[V_a(t) + 60][V_a(t) + 50] - U_a(t) + E_a(t) - I_a(t)]/200$$

$$U_a(t+1) = \begin{cases} U_a(t) + 0.01\{15[V_a(t) + 65] - U_a(t)\} & V(t) \leq -65 \\ U_a(t) + 0.01\{-U_a(t)\} & else \end{cases}$$

where the parameters are the same as in the pyramidal cell (Izhikevich 2007). The two update equations for *U* model the two firing modes of thalamic neurons: bursting (top equation) and tonic (bottom equation). If $V_a(t)$ is greater than $(35 + 0.1U_a(t))$ then the cell has undergone an action potential. $V_a(t)$ is set to $(-60 - 0.1\ U_a(t))$ and $U_a(t)$ is increased by 10.

The voltage update equation for globus pallidus cells is:

$$V_a(t+1) = V_a(t) + [1[V_a(t) + 55][V_a(t) + 40] + 140 - U_a(t) + E_a(t) - I_a(t)]/20$$

$$U_a(t+1) = \{U_a(t) + 0.15\{8[V_a(t) + 55] - U_a(t)\}$$

where the parameters are the same as in the pyramidal cell. If $V_a(t)$ is greater than 25 then the cell has undergone an action potential and $V_a(t)$ is set to -50 and $U_a(t)$ is increased by 200. These cells spontaneously fire unless they are inhibited by some other source. This is modeled by the addition of the constant factor of 140.

## Modular Design

The model is composed of two major modules: a response selection module and a rule-set selection module. The model is designed to use stimulus features called cues that explicitly indicate which rule-set is relevant for the current stimulus. However, there are situations where a change in the relevant rule-set is only apparent through sudden poor task performance (Helie et al. 2015). This possibility, and how the new model could be expanded to handle such a situation, are discussed in Chapter 8.

## Response Selection Module

In the response selection module an array of pyramidal cells represents the stimulus features (Figure 4). If the features are discrete, each feature can be represented by a single cell. If the stimulus has features that vary continuously then each cell will have a receptive field that determines how much excitation a specific feature value within that continuum will generate for that cell. These cells are part of the PFC and thus the features that they represent are highly preprocessed. While they can represent something as simple as the color or orientation of a stimulus, there is no reason that a more complex idea, such as a compound 3D shape or whether or not the current stimulus matches a previous stimulus, could not be symbolized so long as the representation is informative regarding response selection (Freedman et al. 2003). The stimulus cells excite the pyramidal response cells in the premotor cortex (PMC). Within the model these cells represent different responses or categorizations of the stimuli. They are not meant to correspond to specific movements or action plans although in an experimental setting the two are often equivalent. Because responses are mutually exclusive, each PMC cell sends non-somatic inhibition to the incoming excitatory connections of every other PMC cell,

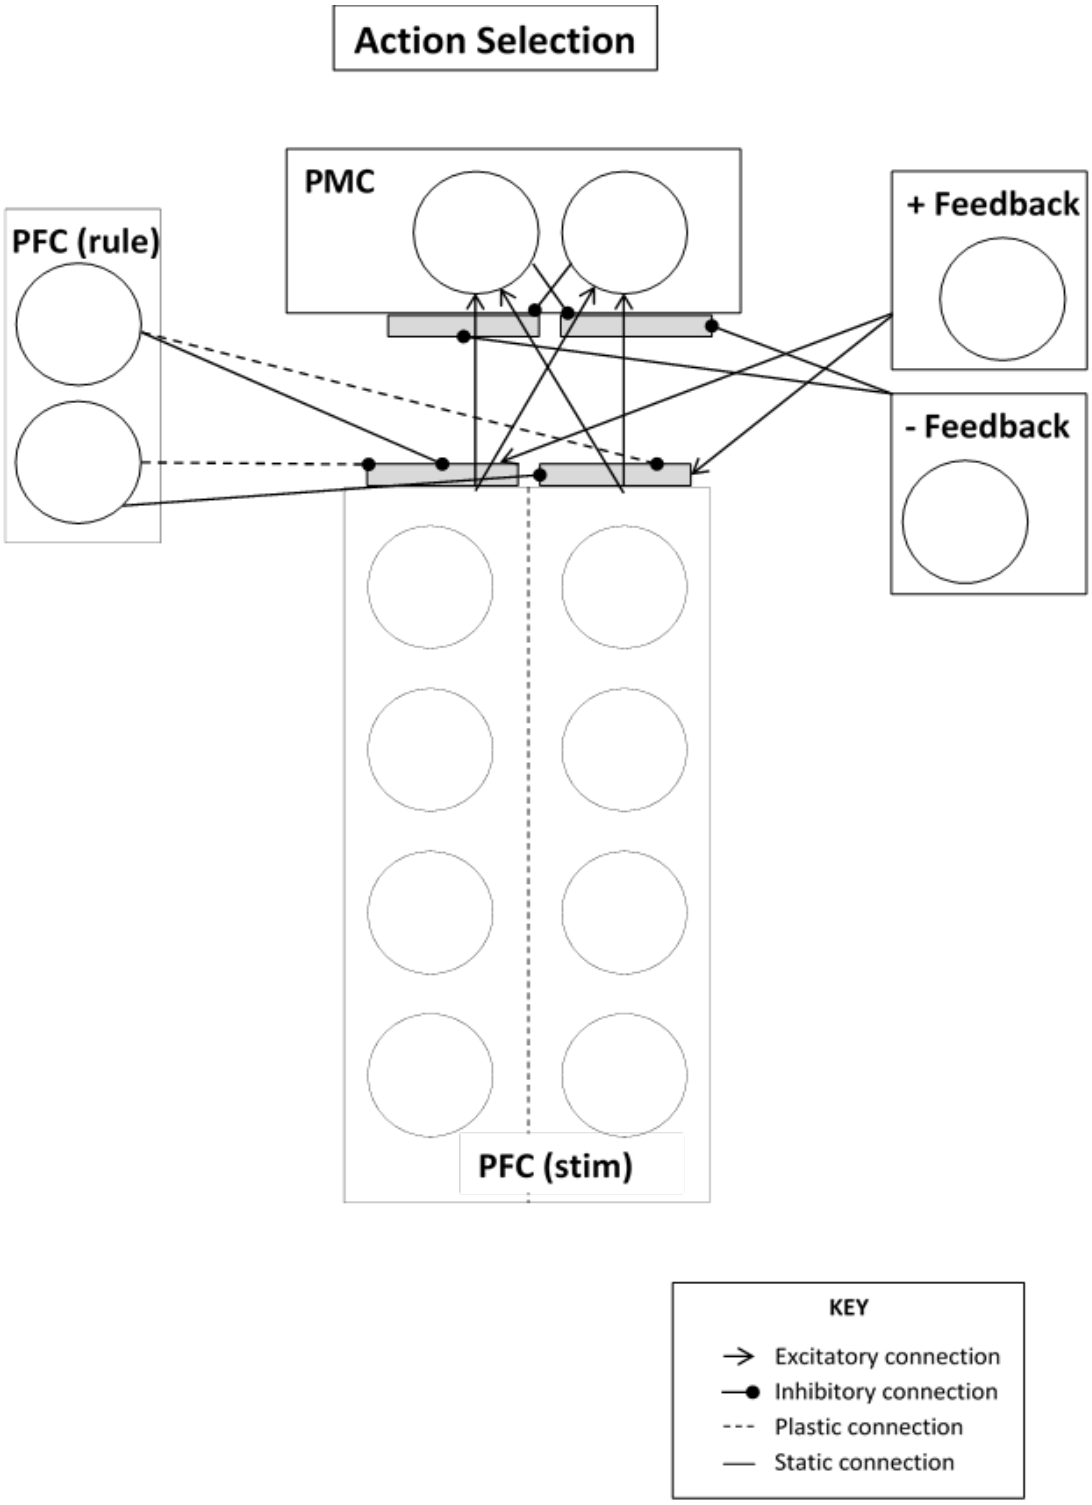


*Figure 4*. The connectivity diagram of the response selection model segment. Non-somatic connections are denoted using a gray field over the receiving somatic connection.

creating winner-take-all behavior (Rumelhart, & Zipser, 1985). Each of the pyramidal rule-set cells also non-somatically inhibits every connection from a stimulus cell to a PMC cell. Each rule-set cell has a dimension of the stimulus associated with it. Connections from stimulus cells that represent features outside of this dimension are inhibited statically and severely. Within the associated dimension the inhibition is weaker and can change based on the level of activity in both the connection receiving the inhibition and the rule-set cell. The feedback cells become active when the model receives feedback. If the feedback the model receives is positive the positive feedback cell becomes active and if it is negative the negative feedback cell becomes active instead. The positive feedback cell non-somatically excites every stimulus cell-to-PMC connection while the negative feedback cell inhibits those connections. At the end of a trial the inhibition between the activated stimulus cell(s) and the active PMC cell will be reduced if the additional excitation from the positive feedback cell was present and increased if the inhibition from the negative feedback cell was present. Thus, if the model receives positive feedback it will be more likely to select the same response cell when it encounters a stimulus with similar features. The equation used to update the strength of every plastic connection in the response selection component is as follows:

$$w_{(a)\underset{ns}{\Rightarrow}\left((b)\underset{s}{\rightarrow}(c)\right)}(n+1) = w_{(a)\underset{ns}{\Rightarrow}\left((b)\underset{s}{\rightarrow}(c)\right)}(n)$$

$$-\eta_1 \cdot \sum_t {O^{GABA}}_a(t) \cdot \left[\sum_t \left[{O^{Glut}}_b(t) + P_{(b)\underset{s}{\rightarrow}(c)}(t)\right]^+ - \theta_1\right]^+$$

$$\cdot \left[w_{max} - w_{(a)\underset{ns}{\Rightarrow}\left((b)\underset{s}{\rightarrow}(c)\right)}(n)\right]$$

$$+\eta_2 \cdot \sum_t {O^{GABA}}_a(t) \cdot \left[\theta_1 - \sum_t \left[{O^{Glut}}_b(t) + P_{(b)\underset{s}{\rightarrow}(c)}(t)\right]^+\right]^+$$

$$\cdot \left[\sum_t \left[{O^{Glut}}_b(t) + P_{(b)\underset{s}{\rightarrow}(c)}(t)\right]^+ - \theta_2\right]^+ \cdot \left[w_{(a)\underset{ns}{\Rightarrow}\left((b)\underset{s}{\rightarrow}(c)\right)}(n) - w_{min}\right]$$

where all the variables are as defined previously, and all of the parameter values for this equation are specific to the task but it is important to note that $\theta_1$ is always greater than $\theta_2$. The second term models the changes in the non-somatic synapse that reduce the inhibition from the rule-set cell. Reduced inhibition means that, under this rule, the range of stimulus values represented by the stimulus cell postsynaptic to the rule-set cell are associated with the response of PMC cell postsynaptic to the stimulus cell. This change only occurs when all three of the following conditions are met: the stimulus cell is active, the PMC cell is active, and the model positive feedback cell is active. Only when all three of these conditions are met is $\sum_t \left[{O^{Glut}}_b(t) + P_{(b)\underset{s}{\rightarrow}(c)}(t)\right]^+$, the input to the PMC cell, greater than $\theta_1$, otherwise this part of the equation will be zero and can be ignored. The last term models what happens when only two of the three conditions are met. In this case, the connection is strengthened and the inhibitory signal becomes stronger. If more than one of the conditions are absent the PMC input does not reach θ2 and so both of the equation segments are zero and the connection strength does not change.

## Cued Rule-Set Selection

The rule-set selection architecture mirrors the response selection architecture using cue features in place of stimulus features and selecting concrete rule-sets instead of responses. As such, the rule-set cells inhibit each other's incoming connections in the same way that the PMC cells do, creating another instance of winner-take-all behavior. The cells representing the cue features activate the rule-set cells and these connections are gated by a cell representing an abstract rule-set (Figure 5). Theoretically, there could be multiple abstract rule-sets, in which case each feature in the cue dimension could be associated with a different rule-set for each possible abstract rule-set. However, none of the tasks used to validate the model incorporate multiple abstract rule-sets. Synaptic plasticity and feedback for these connections operate the same way as in the response selection component, albeit with potentially different parameters in the weight update equation. Once a rule-set is selected a positive feedback loop with the thalamus keeps it active. This activity persists across trials and only ceases when inhibition arrives from the interior segment of the globus pallidus (GPi). As previously described, cells in the globus palidus spontaneously fire at high rates. The GPi is therefore inhibited by the external segment of the globus palidus (GPe). Without the GPe the GPi would be perpetually active and the thalamus would be in a constant state of inhibition. The GPi receives excitation from the pre-supplementary motor area (pre-SMA) via the subthalamic nucleus (Aron et al., 2007). Within the model this excitatory input is the only way for the GPi to overcome the inhibitory influence of the GPe. Studies by Rushworth et al. (2002) and Dove et al. (2000) indicate that the pre-SMA is active when subjects engaged in cue-mediated task switching. Assuming each task has its own set of rules, this excitation of

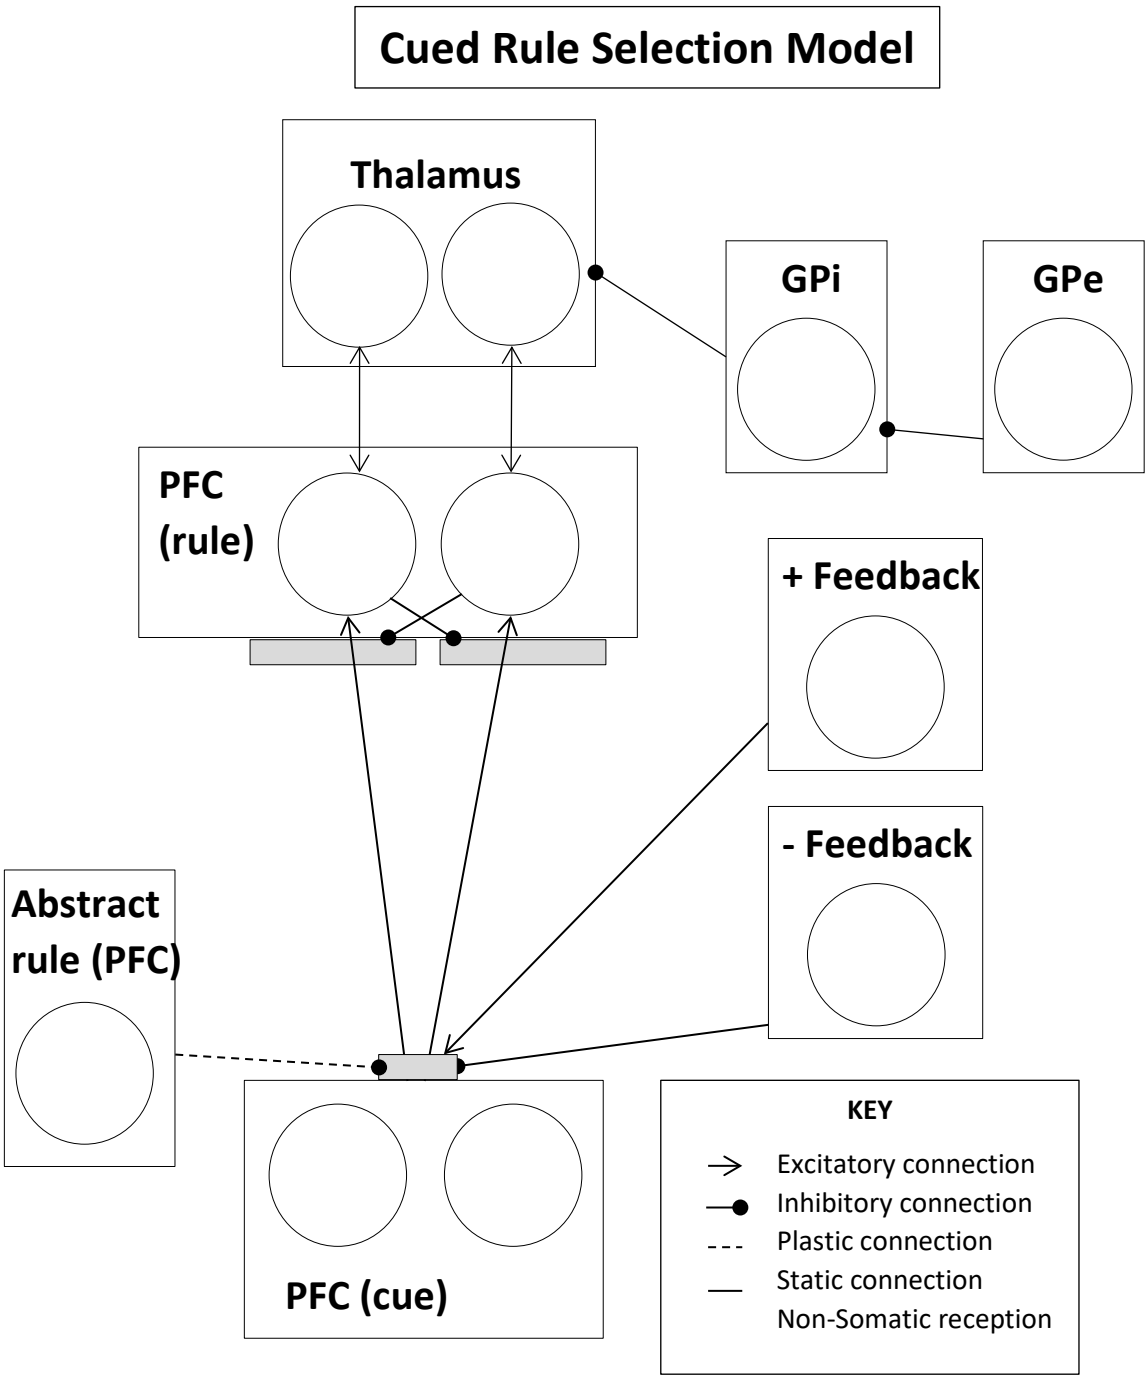


*Figure 5*. The connectivity diagram of the response selection model segment. Non-somatic connections are denoted using a gray field over the receiving somatic connection.

the GPi inhibits the active rule-set set and allows a new cue to select a rule-set without interference from the lateral inhibition. It should also be pointed out that, while this is model learns to associate features of the cue dimension with rule-sets, the distinction between the cue dimension and stimulus dimension(s) are defined beforehand. The model is already assumed to "know" which stimulus dimensions can be used in abstract rules and which can be used in concrete rules. There is no mechanism for building or selecting a representational structure for a task. The representational structure can change between tasks and even between conditions of the same task. When this occurs, the experimenter must adjust the structure of the model to fit the desired representation. Although there is an outside chance that the model could function when every dimension has representation in both the cues cells and the stimulus cells, it was not designed with this implementation in mind, nor has it been tested using such an implementation**.** Any experimental task that involves the subject determining which stimulus dimensions shall be used in abstract vs. concrete rules will require multiple implementations of this model: one for each representational structure. In each of these implementations the cue dimensions will be assumed to be identified beforehand, whether from some outside instruction or an internal decision that is made by cognitive systems outside of the model's scope.

## Default Model Architecture and Input Parameters

There are several experiments to be simulated and each require different model architectures and parameters to properly represent the task. However, there will be many parameters that are common across experiments. A default model architecture and parameter set is presented here to be referred to in cases where an experimental task does not require a unique design. Indeed, one of the experimental tasks will use this default

implementation without a single modification. These parameters as well as those unique to the upcoming experiments were found using grid-search. In the following chapters, any feature not explicitly described as deviating from this description should be assumed to follow it. The default model uses 1 abstract rule-set cell, 1 GPi cell, 1 GPe cell, 2 concrete rule-set cells with 2 corresponding thalamus cells, 2 external cue cells, 1 positive feedback cell, and 1 negative feedback cell. The stimulus cells in the model represent features across 2 stimulus dimensions. The number of cells in each dimension is directly tied to what the cells represent. Each experiment will implement the model with a different number of stimulus cells and therefore this parameter is left undefined in the default model. The number of PMC cells is also dependent upon the number of possible responses in the task so it too will be defined for each task. There should always be one thalamic cell for each concrete rule-set cell. Each stimulus cell receives an input of 500mA when a stimulus with the feature it codes for is presented (see Figure 6 – stimulus presentation occurs at 1550ms rather than 500ms). Simulated trials run for 2800ms including a 500ms period at the start of each trial for which no stimulus or cue is presented (burn-in time) (Figure 6 uses 3500ms and 1300ms respectively). During this 500ms period the GPi cell receives a 200mA input current representing the reset signal from the preSMA. After 500ms the input to the GPi stops. The reset signal is applied on every trial, before it is clear that the task has changed. This is supported by the findings of Monsell & Mizon (2006) who suggest that preemptive task reconfiguration occurs at every trial unless the probability of a task-switch is low. The duration of the reset signal is independent of the length of the burn-in time.

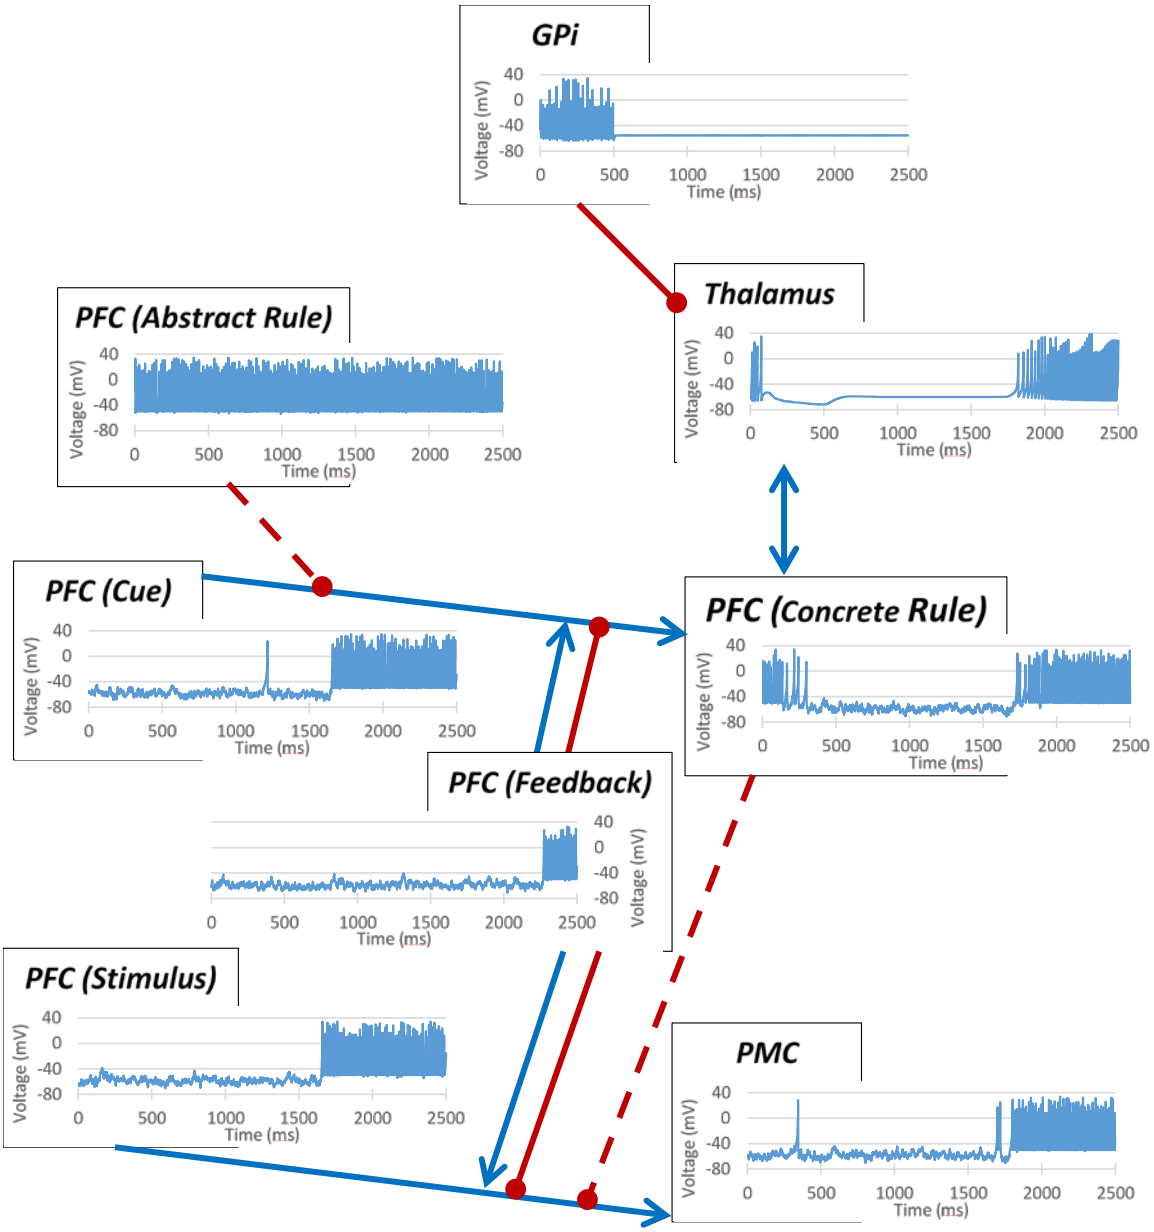


*Figure 6.* A connectivity diagram including spike trains for the major brain areas included in the model. Blue connections are excitatory and red connections are inhibitory. Dotted lines are plastic while solid lines are fixed. Spike train data was taken from a trial from the task from van ‘t Wout, Lavric, & Monsell (2015) (see Chapter 7 for a full description). The spike train shown as from the most active cell in each brain area and the graphs are cut off at 2500ms since there is no change in activity after that point.

If the length of the burn-in time is altered for an experiment the duration of reset signal remains unchanged unless otherwise specified. The stimulus and cue are presented at this point as well, 500ms after the start of the trial. Upon presentation of the stimulus, the stimulus cells that are associated with the features of the presented stimulus receive a constant excitatory signal of 500mA for the remainder of the trial (see Figure 6). Each external cue cell functions exactly the same way, receiving 500mA of excitatory input for the rest of the trial if the cue feature it codes is presented. The abstract rule-set cell receives a constant input of 500mA over the full length of the trial including burn-in (see Figure 6). A running sum of each PMC cell's $O^{Glut}$ term is tracked over the course of the trial. When one of these sums reaches 7000 the model is considered to have made a response. Each PMC cell represents a different response and the first PMC cell to reach the threshold determines the response made. Once a response has been made, it cannot be changed and no other response can be made for that trial. If the response made by the model is correct, the positive feedback cell starts receiving a 500mA input, which continues for the remainder of the trial (see Figure 6 – Input to the feedback cell starts at about 2250ms). If instead the model makes an incorrect response the negative feedback cell receives 500mA input for the remainder of the trial.

Table 1 lists the connection strengths of every non-plastic synapse in the default model. The $(PMC) \underset{ns}{\Rightarrow} \left((stim) \underset{s}{\rightarrow} (PMC)\right)$ and $(concrete\ rule) \underset{ns}{\Rightarrow} \left((external\ cue) \underset{s}{\rightarrow} (concrete\ rule)\right)$ non-somatic connections only target connections where the post-synaptic cell is not the cell sending the non-somatic signal - i.e. PMC cells inhibit the incoming connections to other PMC cells but not their own incoming connections. "Out stim" refers to all stimulus cells that code features that are outside the dimension

Table 1

*Fixed Connection Weights in the Model for the Default Model Implementation*

| Designation | Weight |
|---|---|
| $w_{(stim)\underset{s}{\rightarrow}(PMC)}$ | 55 |
| $w_{(external\ cue)\underset{s}{\rightarrow}(concrete\ rule)}$ | 25 |
| $w_{(concrete\ rule)\underset{s}{\rightarrow}(thalamus)}$ | 40 |
| $w_{(thalamus)\underset{s}{\rightarrow}(concrete\ rule)}$ | 40 |
| $w_{(GPi)\underset{s}{\Rightarrow}(thalamus)}$ | 100 |
| $w_{(GPe)\underset{s}{\Rightarrow}(GPi)}$ | 40 |
| $w_{(PMC)\underset{ns}{\Rightarrow}\left((stim)\underset{s}{\rightarrow}(PMC)\right)}$ | 3 |
| $w_{(pos)\underset{ns}{\rightarrow}\left((stim)\underset{s}{\rightarrow}(PMC)\right)}$ | 0.8 |
| $w_{(neg)\underset{ns}{\Rightarrow}\left((stim)\underset{s}{\rightarrow}(PMC)\right)}$ | 0.4 |
| $w_{(concrete\ rule)\underset{ns}{\Rightarrow}\left((external\ cue)\underset{s}{\rightarrow}(concrete\ rule)\right)}$ | 6 |
| $w_{(pos)\underset{ns}{\rightarrow}\left((external\ cue)\underset{s}{\rightarrow}(concrete\ rule)\right)}$ | 0.8 |
| $w_{(neg)\underset{ns}{\Rightarrow}\left((external\ cue)\underset{s}{\rightarrow}(concrete\ rule)\right)}$ | 0.4 |
| $w_{(concrete\ rule)\underset{ns}{\Rightarrow}\left((out\ stim)\underset{s}{\rightarrow}(PMC)\right)}$ | 6 |

associated with the concrete rule-set cell. In this way, the stimulus cells that represent features in dimensions outside of the rule-set cell's scope are forcefully prevented from influencing the response decision. The pair of connections from the concrete rule-set cells to the thalamic cells are one-to-one, as are the connections from the thalamic cells to the

concrete rule-set cells. Each concrete rule-set cell excites and is excited by a single thalamic cell, creating a positive feedback loop. All other connections in Table 1 exhibit all-to-all connectivity. There are two plastic non-somatic connections, $(concrete\ rule) \underset{ns}{\Rightarrow} \left( (in\ stim) \underset{s}{\rightarrow} (PMC) \right)$ and $(abstract\ rule) \underset{ns}{\Rightarrow} \left( (external\ cue) \underset{s}{\rightarrow} (concrete\ rule) \right)$. In the plastic connection denoted by $(concrete\ rule) \underset{ns}{\Rightarrow} \left( (in\ stim) \underset{s}{\rightarrow} (PMC) \right)$ "in stim" refers to those stimulus cells that represent features in the dimension of the rule-set associated with the concrete rule-set cell. In this way, it is the complement of "out stim" described above. Both of these connections are all-to-all. The other parameters for these connections vary between every experimental model implementation. Therefore, those parameters will be defined separately for each experiment modeled. See the following pages for a detailed step-by-step account of signal propagation and the non-somatic gating mechanism in both the response selection module and cued rule-set selection module. Each step includes a description of the cellular activity, the transmission of signals, and the effect of those signals on other cells at that point in the process.

## Summary

The model uses plastic inhibitory non-somatic connections to learn associations and select responses. The design is recursive with rule-sets selected in the same way as physical responses reflecting the hierarchical organization of the PFC. Rule-sets are maintained in working memory until a reset signal is sent through the basal ganglia. The model is grounded in biological data. The next question is: how well can the model reproduce experimental results.

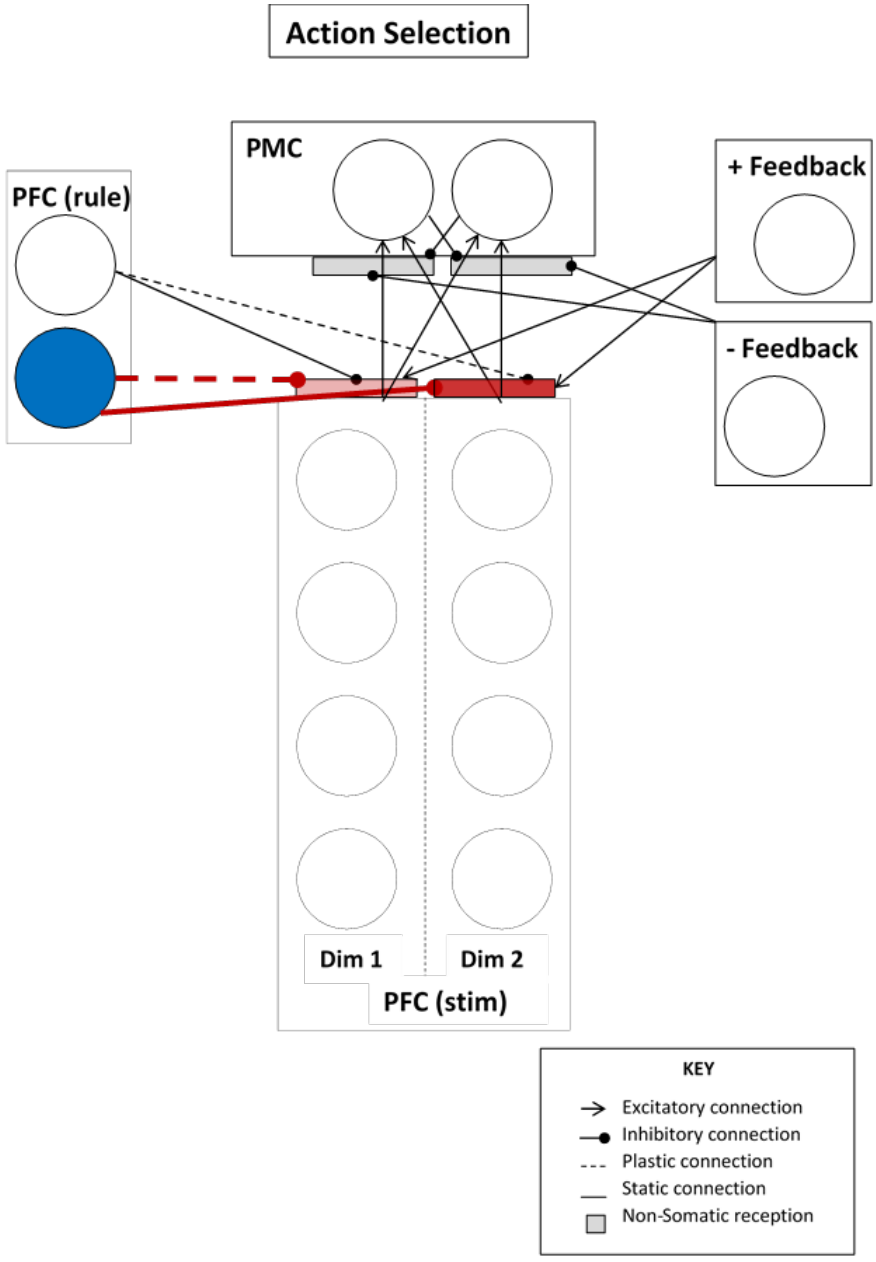


*Figure 7*. We assume a rule-set has already been selected (see other figure for rule-set selection process). The selected rule-set cell sends strong, blanket non-somatic inhibition to the outgoing connections of stimulus cells in the irrelevant dimension (Dim 2) and selective, milder non-somatic inhibition to the outgoing connections of stimulus cells in the relevant dimension (Dim 1).

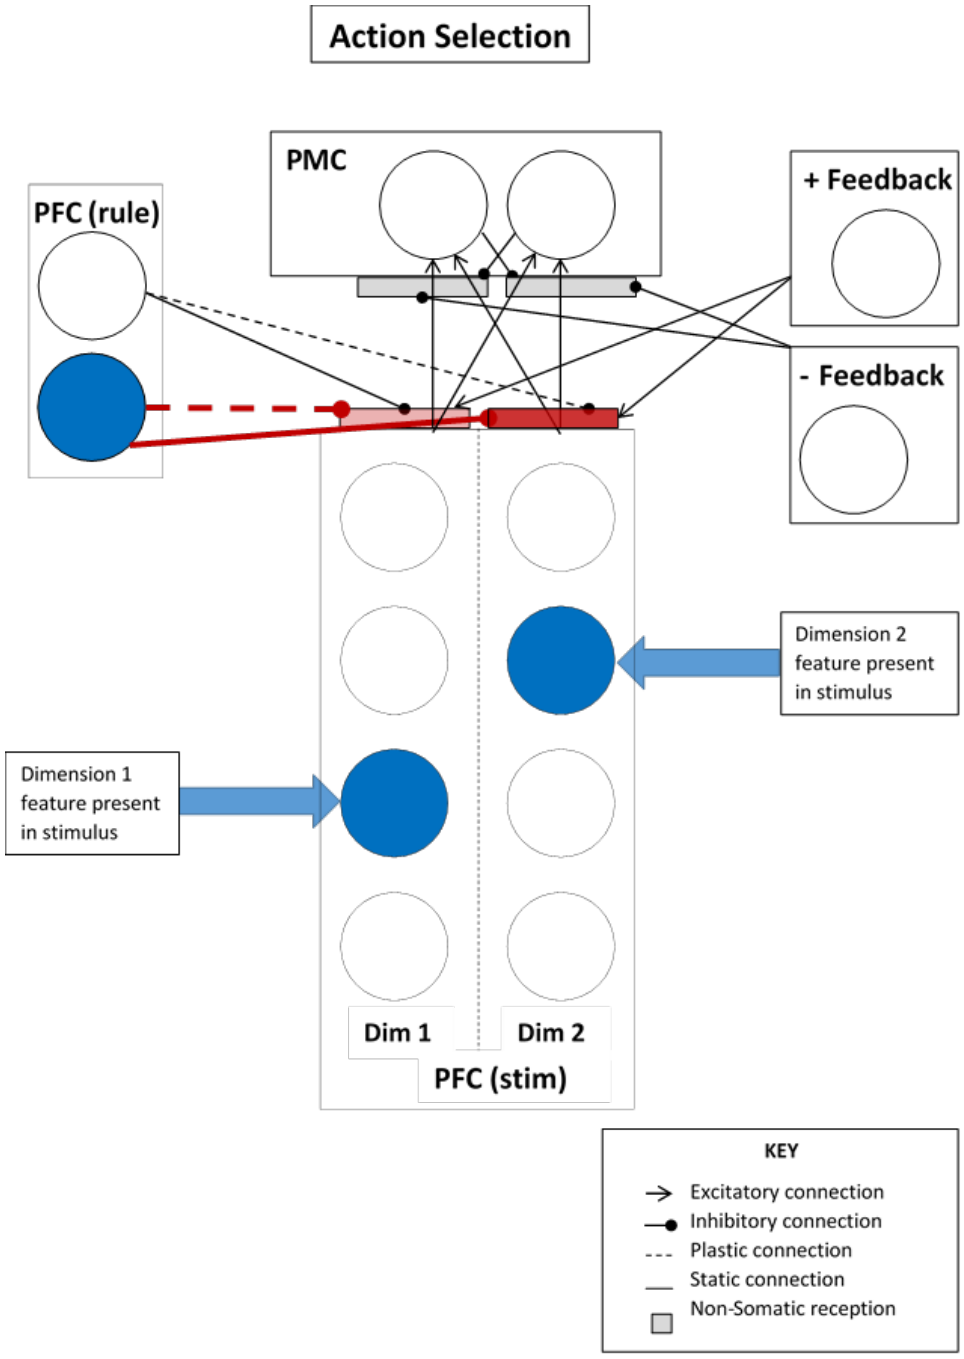


*Figure 8*. The stimulus is presented and the stimulus cells in each dimension that correspond to the stimulus's features receive excitatory input and become active.

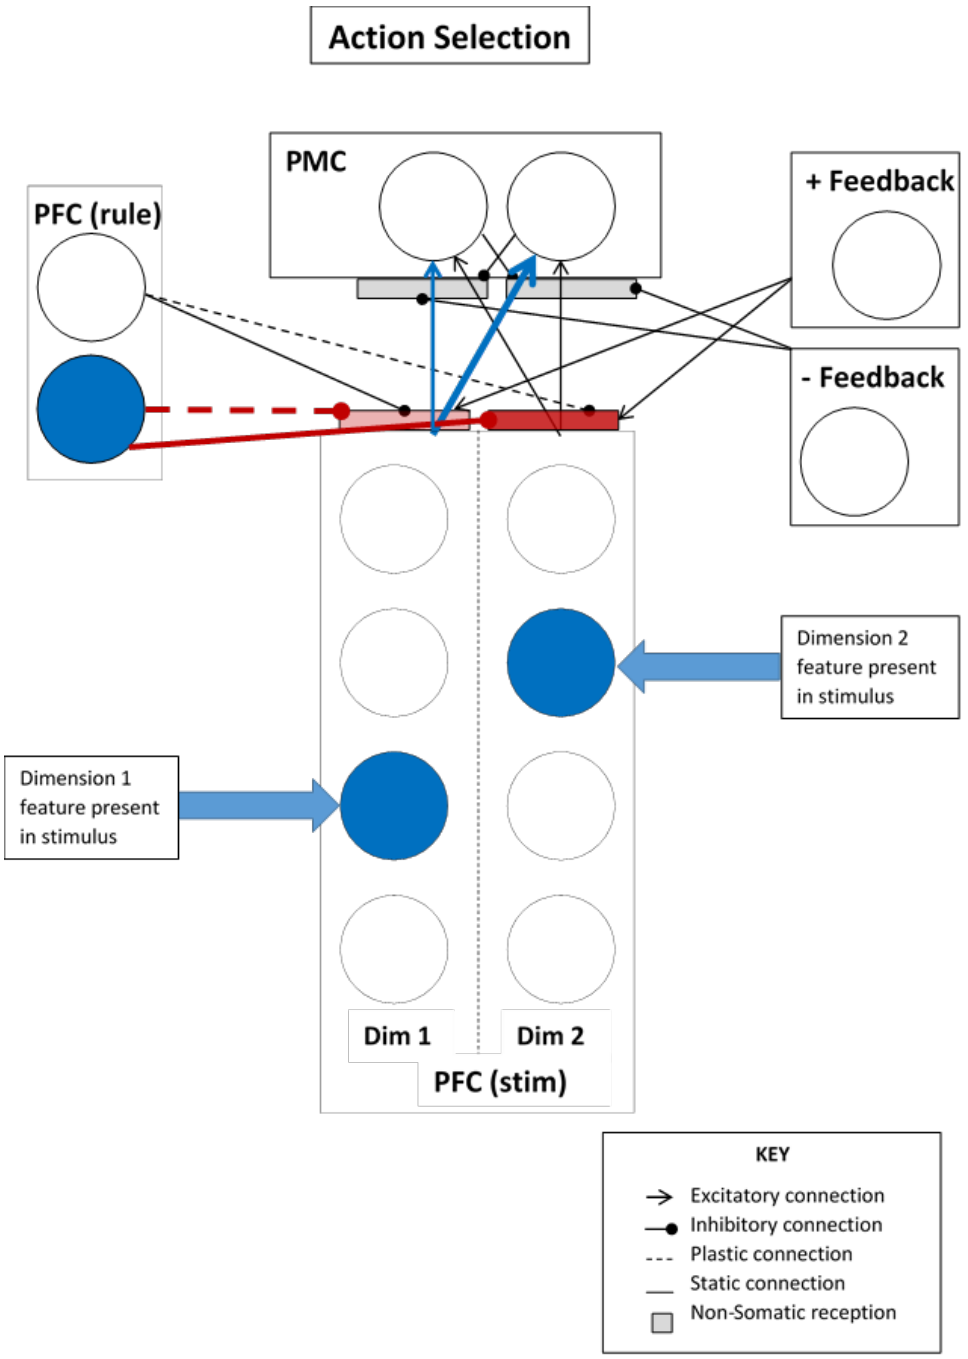


*Figure 9*. The strong inhibition on the outgoing connections from stimulus cells of the irrelevant dimension prevents the active cell in that dimension from sending any signal to the PMC cells. Assume some learning has taken place in the relevant dimension and that the inhibition is stronger on the connection to the left PMC cell. Thus, the signal that the right PMC cell receives is stronger than what the left receives.

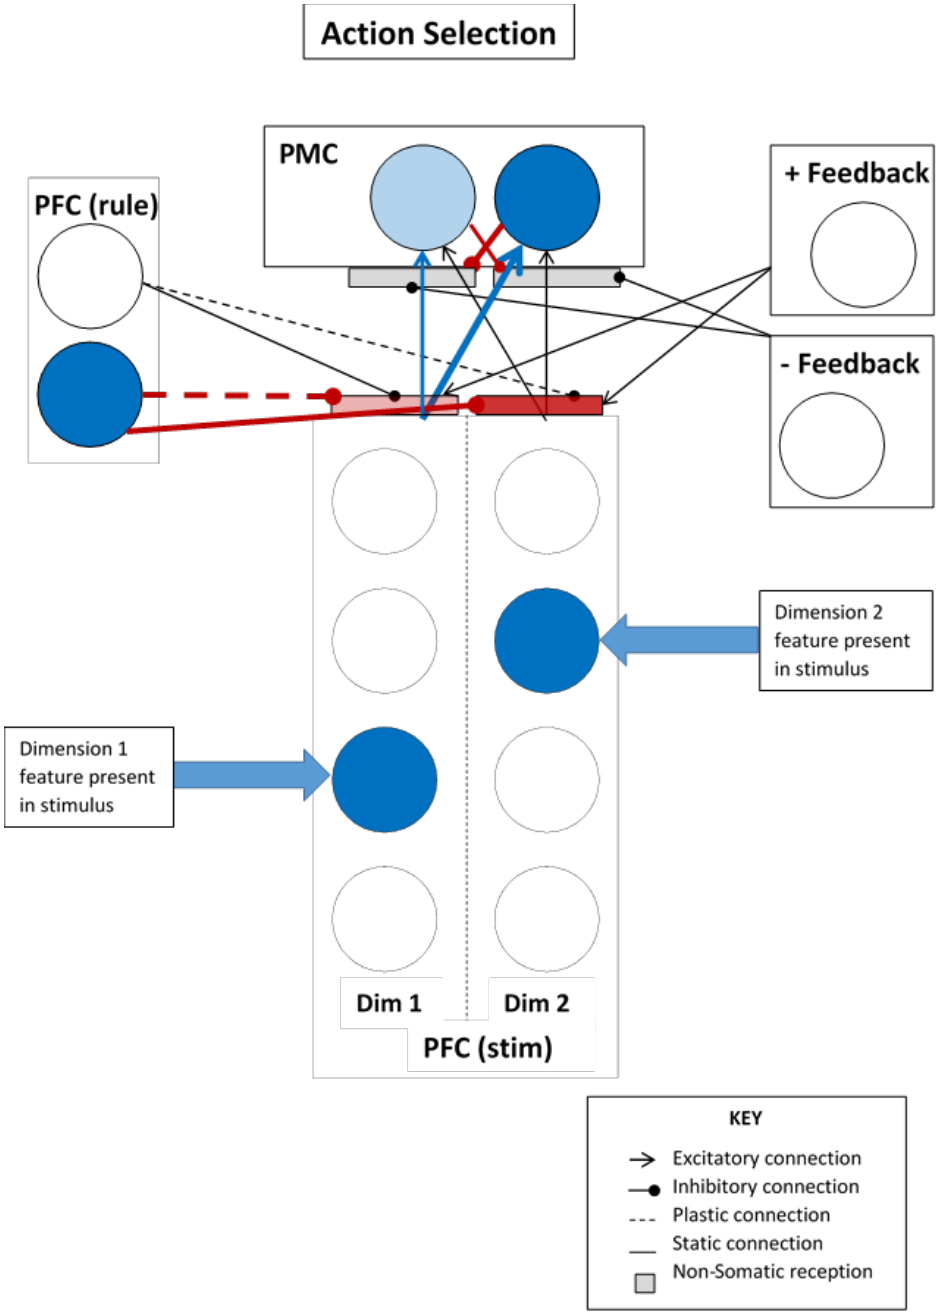


*Figure 10*. Both of the PMC cells become active and begin laterally inhibiting the incoming connections to each other. However, the weaker signal received by the left PMC cell results in less activity and less lateral inhibition compared to the cell on the right. If this were early in the training, the excitation received by the PMC cells would be nearly equal. In this case, the necessary differences in activity would arise primarily through noise.

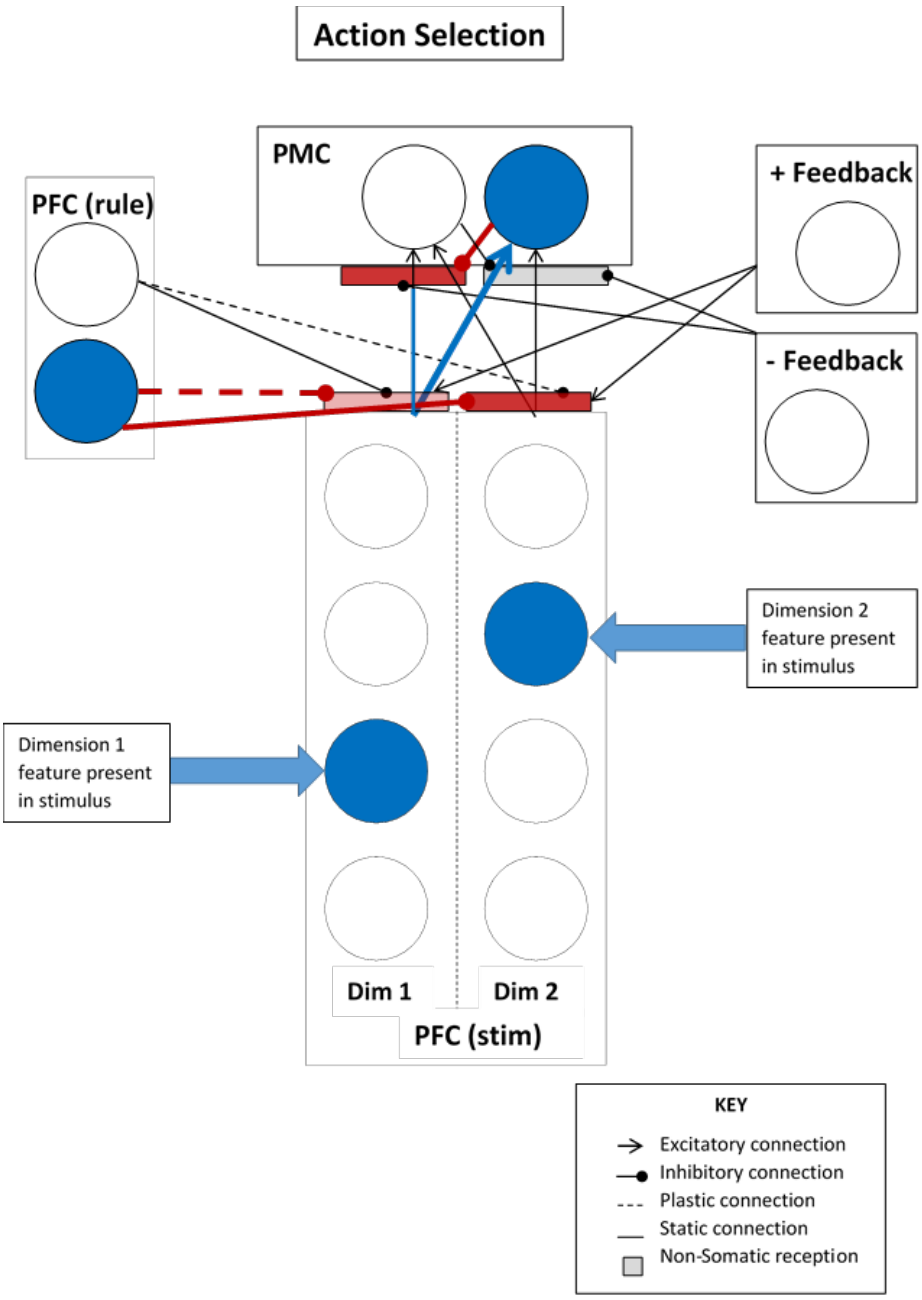


*Figure 11*. The weak inhibition from the left PMC cell has only a small effect on the input and, by extension, the activity of the right PMC cell. Meanwhile, the strong inhibition from the right PMC cell stifles what little excitatory input the left PMC cell received. The right cell is now free of the left's inhibition and there is no possible way for the left or any other PMC cell to become active. The output of the right cell accumulates until it reaches the response threshold and the model responds accordingly.

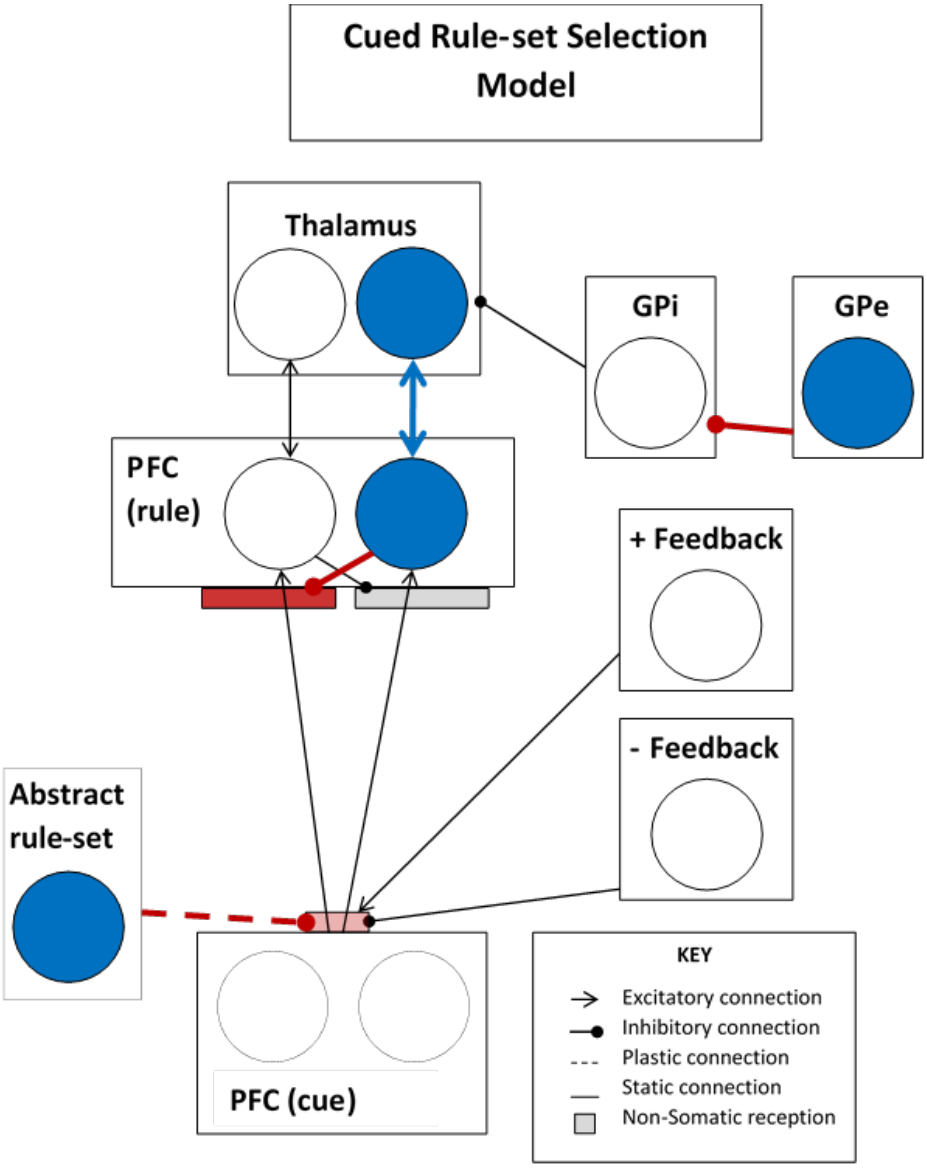


*Figure 12*. In rule-set selection there is a lot of holdover from the previous trial and therefor a lot more going on at the start. The abstract rule-set cell is active and gating the connections from the cue cells to the concrete rule-set cells (although the cue cells themselves are not active). The concrete rule-set cell that was used in the previous trial stays active into the next and its activity is maintained through a pair of excitatory connections with a thalamic cell. The exteranal segment of the globus palidus is naturally active even in the absence of excitatory input. The GPe inhibits the GPi which otherwise would be similarly active.

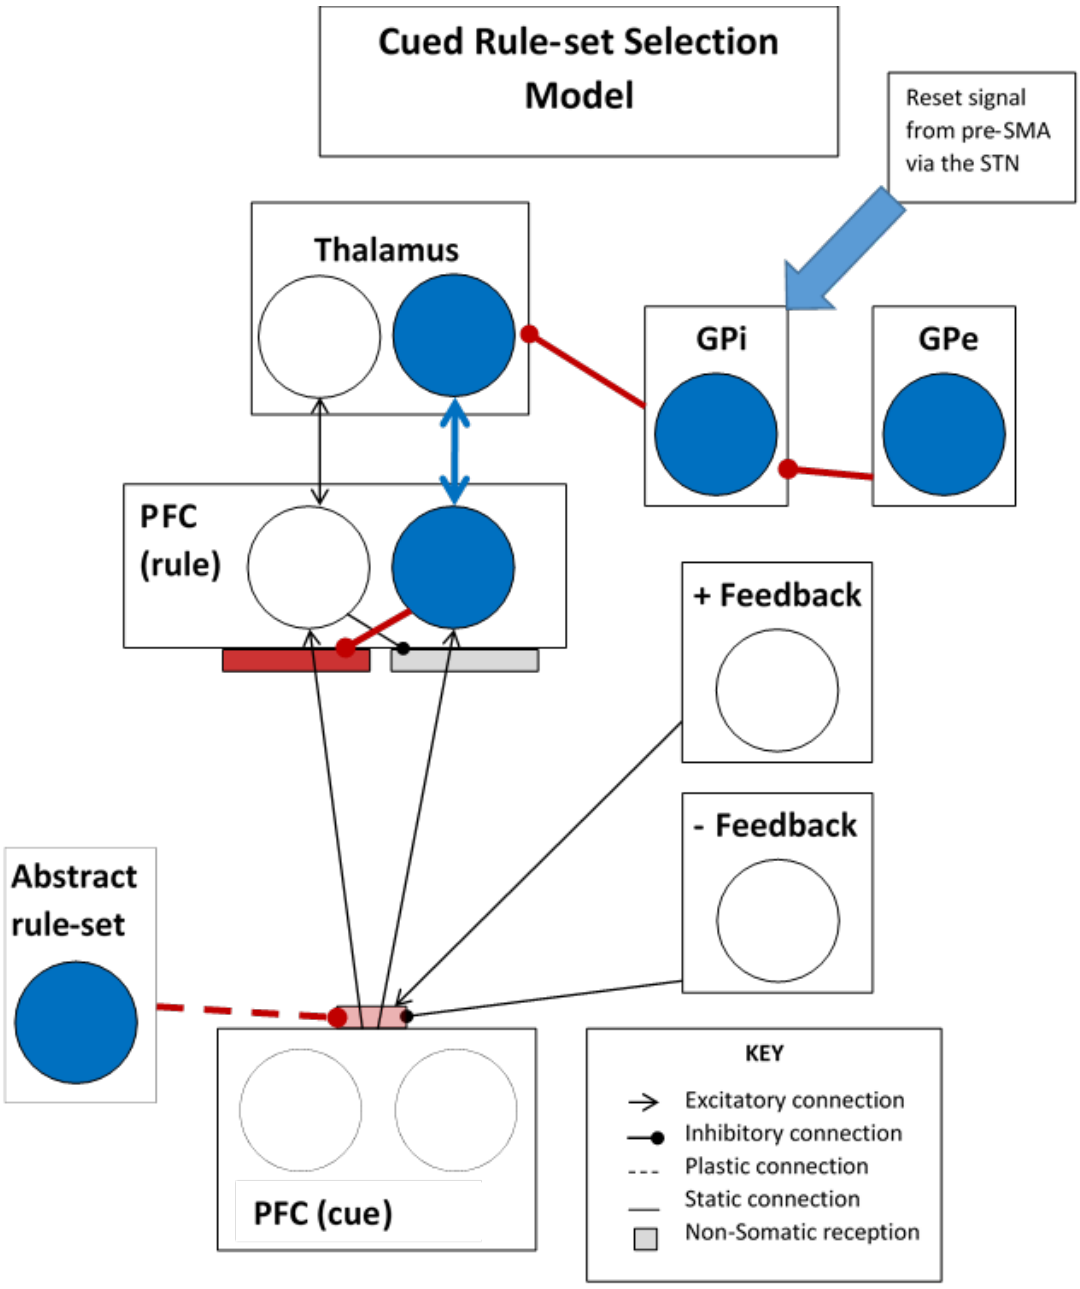


*Figure 13*. Once the trial begins, an excitatory signal is sent to the GPi from the subthalamic nucleus in order to clear the previous rule-set from working memory. This excitation counteracts the inhibition from the GPe and the GPi becomes active. The GPi begins to inhibit the thalamus.

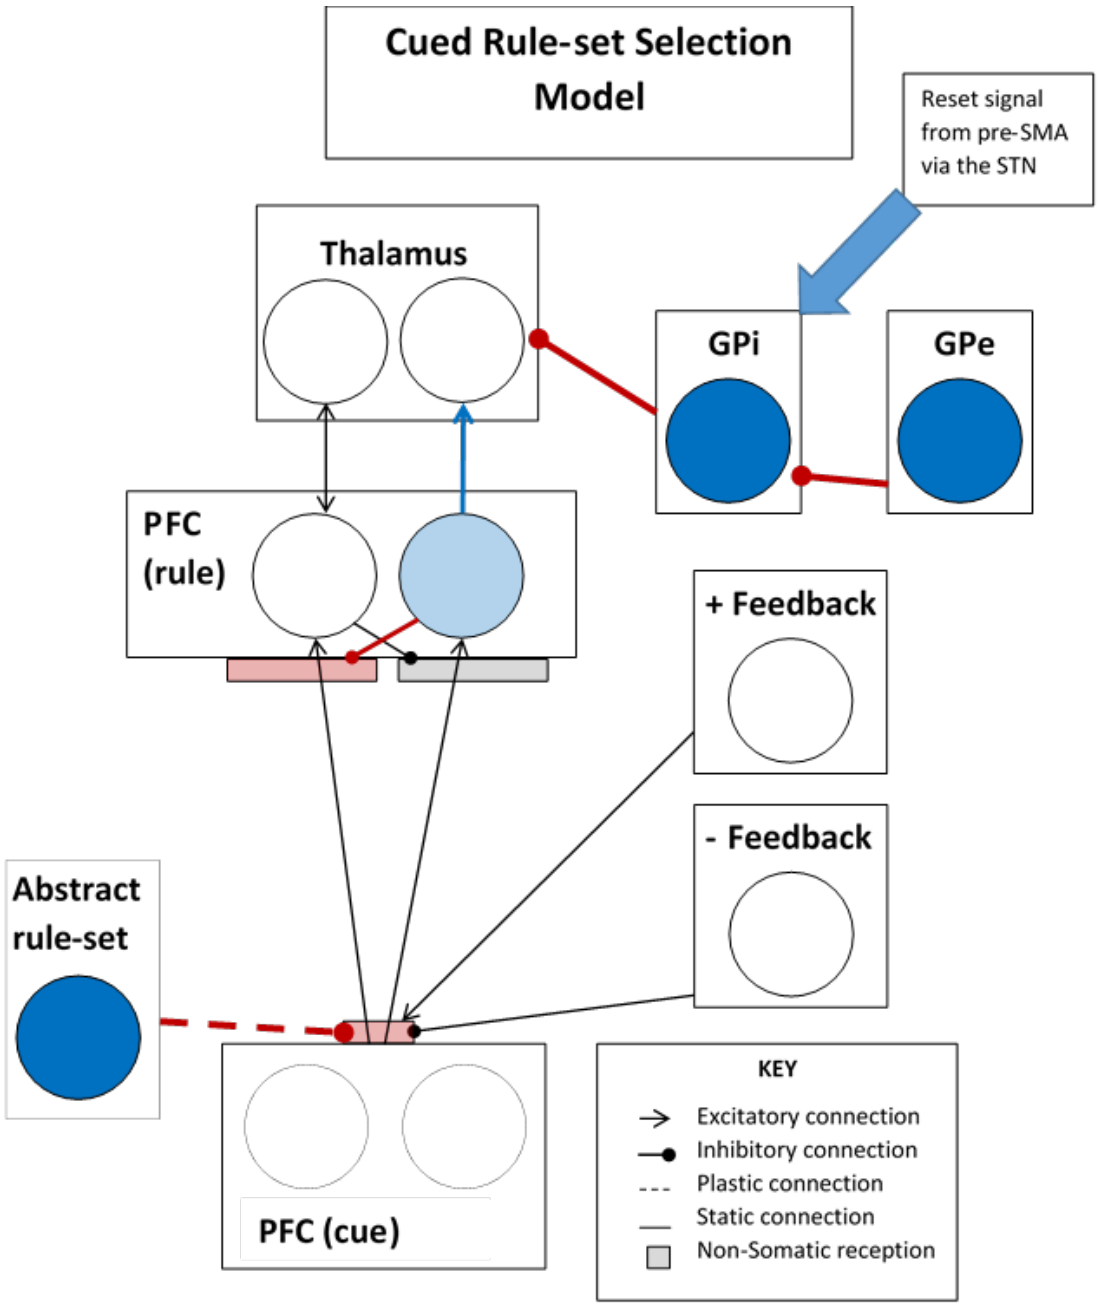


*Figure 14*. The inhibition from the GPi prevents the thalamus from firing. The input to the concrete rule-set cell dwindles and soon both it and the thalamus will be completely inactive.

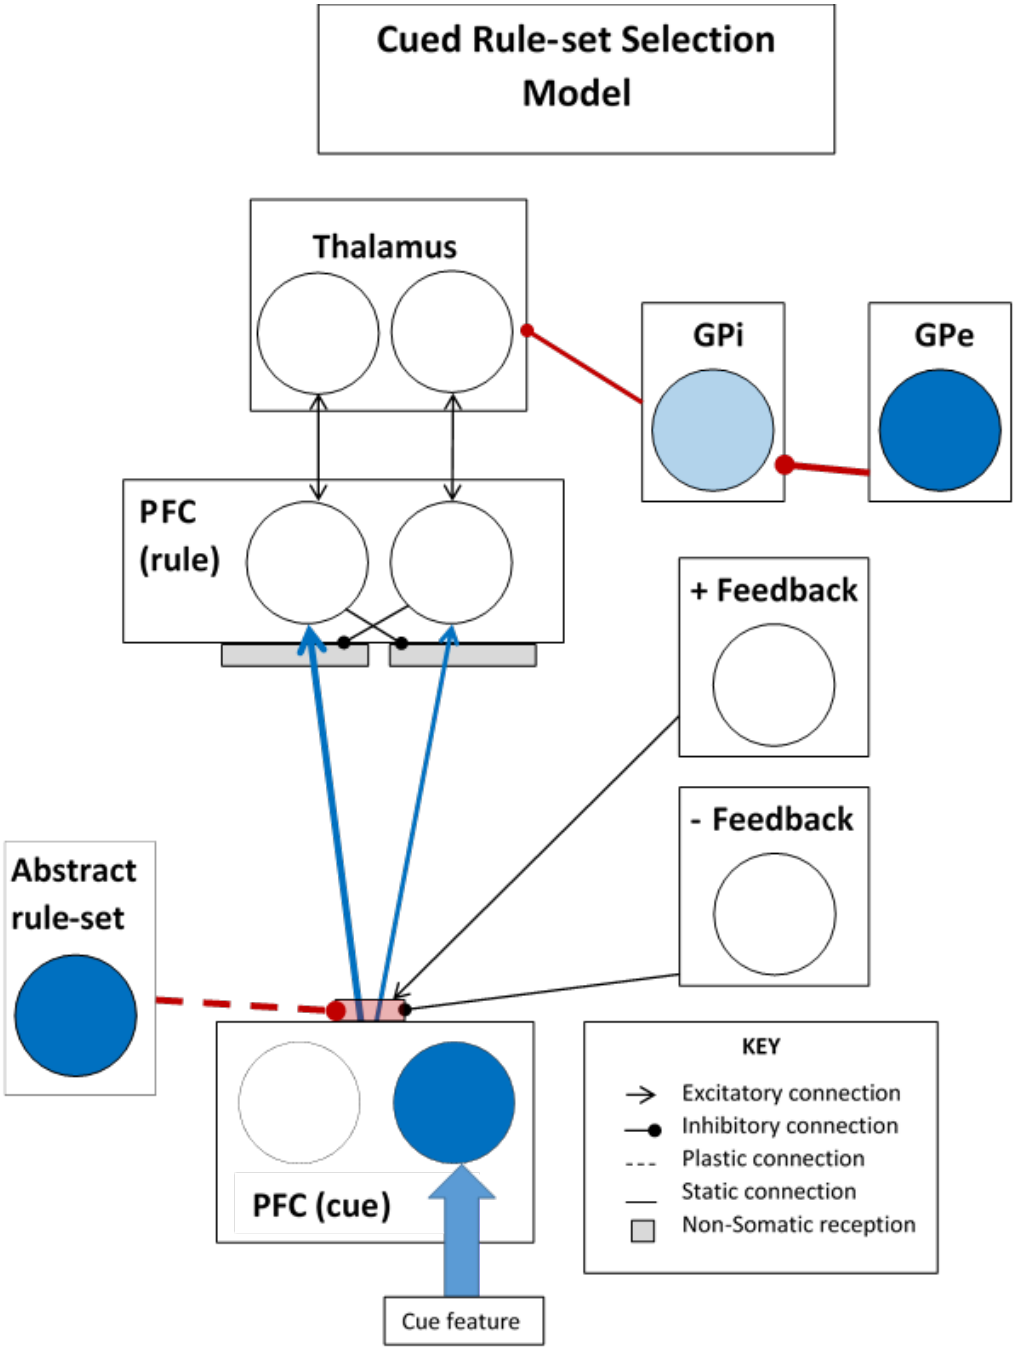


*Figure 15*. For brevity, assume that the cue is presented at about the same time that the reset signal stops. Without the excitation of the reset signal the GPi begins to go quiet. The cue cell representing the cue becomes active and begins sending its signal, gated by the abstract rule-set cell. Once again, assume that some learning has taken place and that the gating inhibition is uneven.

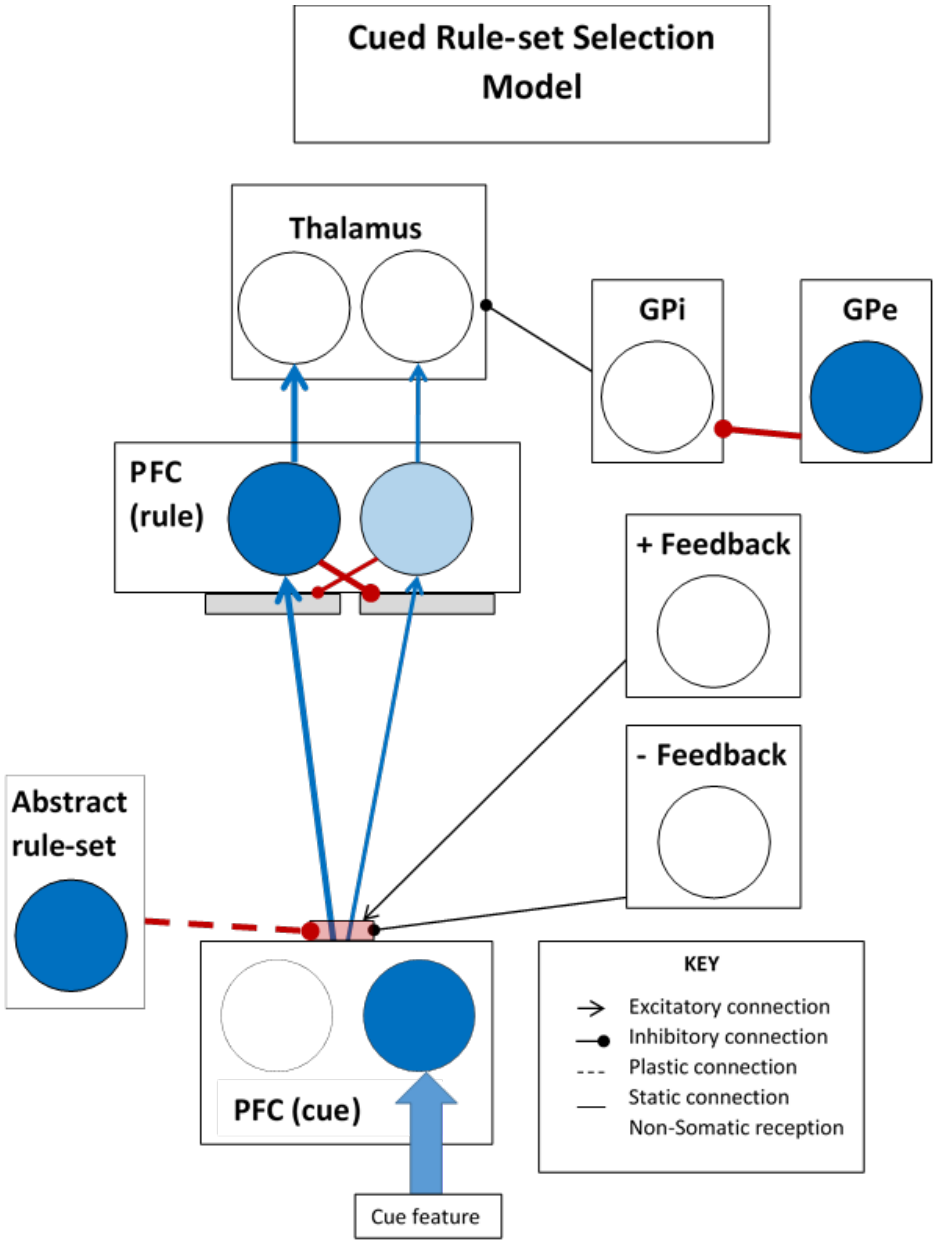


*Figure 16*. Both of the concrete rule-set cells become active and begin exciting their corresponding thalamus cells as well as laterally inhibiting the incoming connections to each other. However, the weaker signal received by the right concrete rule-set cell results in less activity and less lateral inhibition compared to the cell on the left. If this were early in the training, the excitation received by the concrete rule-set cells would be nearly equal. In this case, the necessary differences in activity would arise primarily through noise.

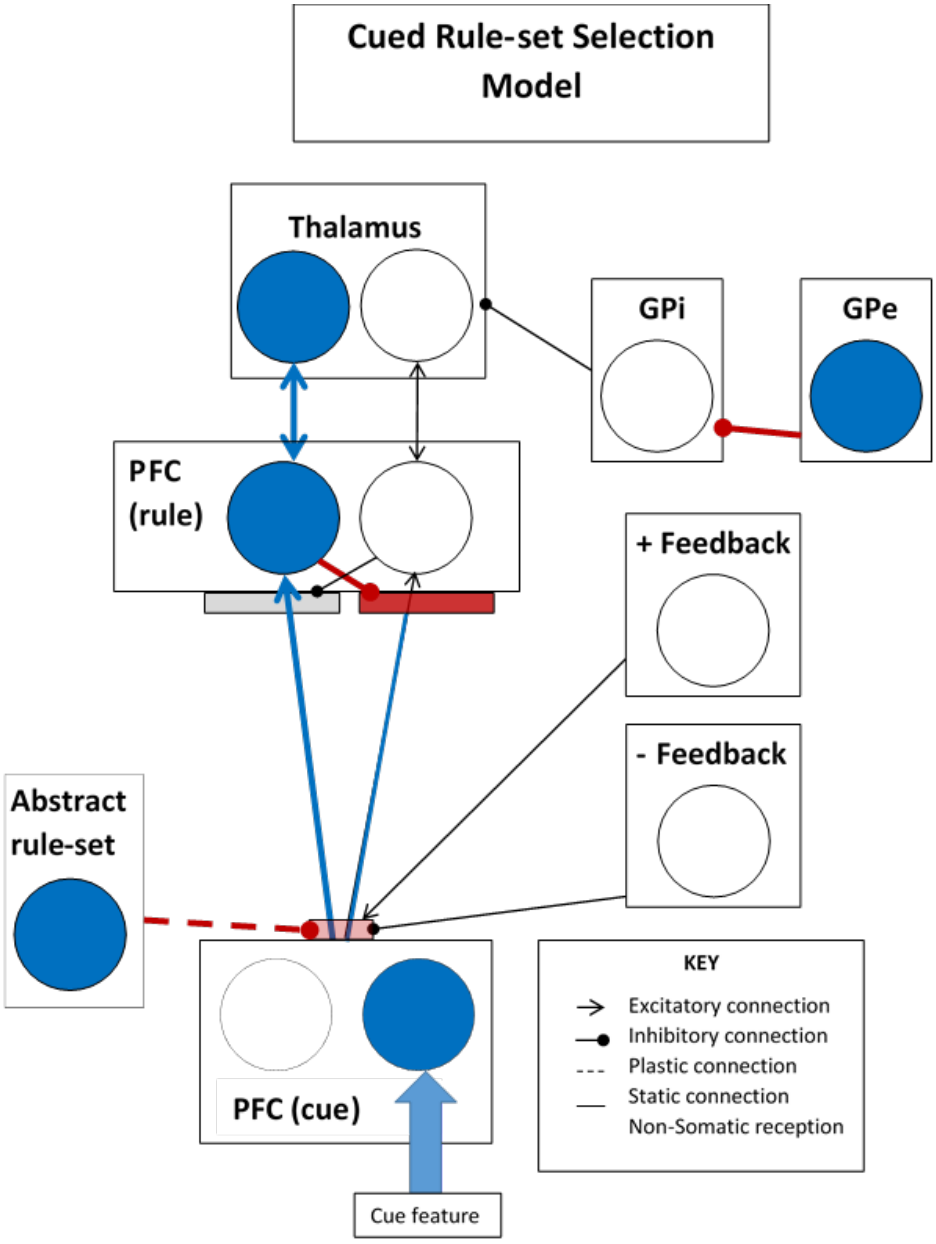


*Figure 17*. The weak inhibition from the right concrete rule-set cell has only a small effect on the input and, by extension, the activity of the left concrete rule-set cell, especially once it begins reverberating with its paired thalamus cell. Meanwhile, the strong inhibition from the left concrete rule-set cell stifles what little excitatory input the right concrete rule-set cell received. What little excitation makes it to the thalamus and back is not enough to start a positive feedback loop. The right cell is now free of the left's inhibition and there is no possible way for the right or any other concrete rule-set cell to become active. The response selection module of the model is now gated solely by the left concrete rule-set c.

# EXPERIMENT 1 – BADRE, KAYSER, AND D'ESPOSITO 2010

The first test for a model of multiple rule-sets is to determine if it performs better when using multiple rule-sets than when using only one rules-set. In Badre, Kayser, and D'Esposito's 2010 paper subjects were tested in two conditions: the hierarchical condition in which using multiple rule-sets would have been beneficial, and the flat condition in which it would not. Subjects in the flat condition learned slower and had a lower final accuracy than subjects in the hierarchical condition. The model was tested on these two tasks to see if it would also perform better in the hierarchical condition.

## Task

In the Badre, Kayser, and D'Esposito task subjects categorized 18 stimuli into 3 categories. The stimuli were pictures of computer generated 3D shapes and varied in three dimensions: the shape itself, the angle at which the shape was viewed, and the color of the border. Subjects saw each stimulus 20 times for a total of 360 trials. There were two conditions that differed in the pattern of associations between the stimuli and categories. In the hierarchical condition, the stimuli with red borders can be correctly categorized solely by using the shape while the stimuli with blue borders can be correctly categorized by using only the orientation (see Figure 18). This means that subjects could represent the stimuli hierarchically with the color of the border acting as a cue to use either a rule-set on shape or a rule-set on orientation. In the flat condition, there is no exploitable structure to the stimulus-response pairings and all three dimensions must be used to correctly categorize each stimulus.

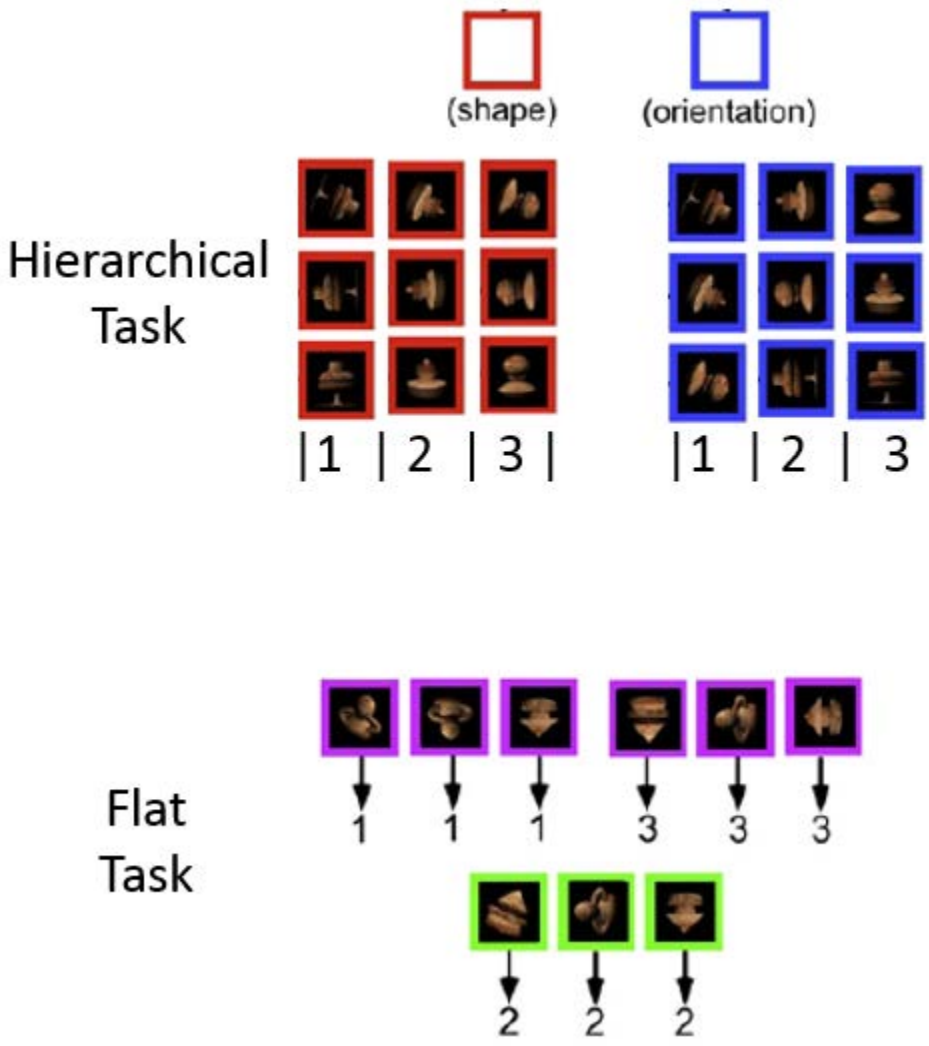


*Figure 18*. A sample of the stimuli from Badre, Kayser, and D'Esposito's (2010). Note how the hierarchical stimuli can be arranged by shape for red-bordered stimuli and by orientation for blue-bordered stimuli while no such pattern exists for the flat stimuli. Figures taken from Badre Kayser, and D'Esposito's (2010).

## Human Data

Subjects were given no information about any organization in the stimulus-response associations yet subjects in the hierarchal condition learned more quickly and achieved a higher final accuracy than subjects in the flat condition (Figure 19. There are two reasons for this improvement in performance. The first is that there are fewer associations to learn when using a hierarchical representation: 2 border colors to 2 rule-sets, 3 shapes to 3 responses, and 3 orientations to 3 responses for a total of 8 associations. In the flat representation each of the 18 combinations of border color, shape, and orientation needs its own association to a response. Secondly, in the flat condition there is overlap in the relevant dimensions. Stimuli with different responses share features with each other, which causes interference. While the same is technically true in the hierarchical condition, this overlap only occurs in the dimension designated irrelevant by the cue. Once the subject has learned to use the cue to ignore one of the dimensions, they are able to disregard misleading interference from the irrelevant dimension. In the flat condition, there are no irrelevant dimensions and therefore subjects are subject to disruptive interference.

It should be noted that Figure 19 is not a representation of mean accuracy across all subjects. Instead, it presents learning curve estimates of the most typical subject for each condition. The model results (presented later) uses mean.

## Task-Specific Model Parameters

As stated in the model description, problem representation is defined by the architecture of the model. This means that different representations will use different architectures. This being the case, the task was run using two versions of the model: one

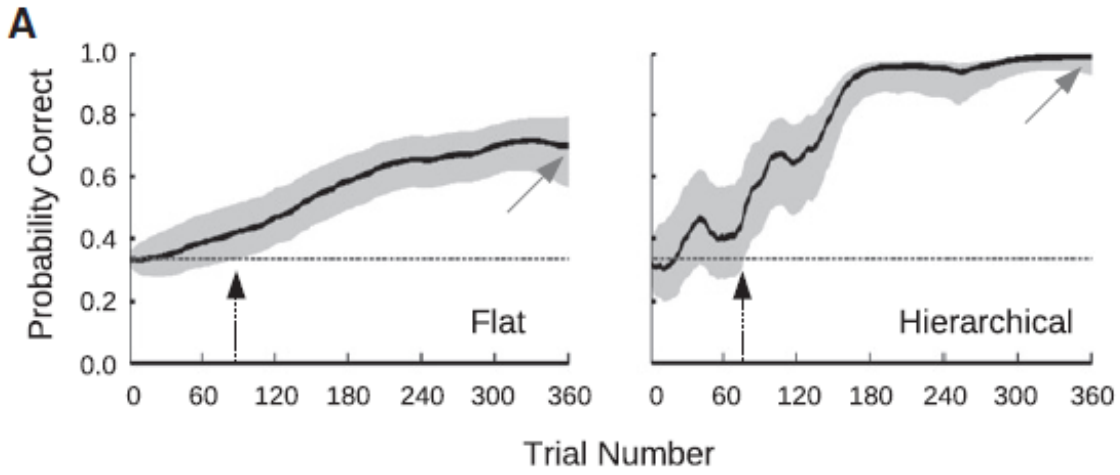


*Figure 19*. Human results from Badre, Kayser, and D'Esposito's (2010). Graphs show the learning curve estimates and 90% confidence interval of the most typical subject for each condition. Graphs taken from Badre Kayser, and D'Esposito's (2010).

that is set up to take advantage of the hierarchical organization of the first condition, and one that merges the three stimulus dimensions and treats them as a single composite dimension with 18 values, representing the second, flat condition.

The hierarchical model used the default model almost exactly. Each stimulus dimension had 3 cells, one for each possible shape and orientation and three PMC cells, each representing one of the responses. The parameters for the plastic non-somatic connections can be found in Tables 2 and 3. Of those parameters, all of the eta and theta terms are free parameters for a total of eight free parameters.

The flat model is not hierarchical and therefore only used one rule-set and one corresponding concrete rule-set cell with a paired thalamic cell. Although technically none of the stimulus features were cues, the concrete rule-set cell must still receive excitation from somewhere so there was one 1 external cue cell, which was activated in every trial. As in the hierarchical model, there were 3 PMC cells for each of the three responses. Each unique combination of shape, orientation and border color was treated as a separate stimulus feature meaning there is a single dimension containing 18 stimulus cells. With only one choice of rule-set the connection weight of the non-somatic inhibition from the abstract rule-set cell was set to zero. Leaning was also disabled for that connection. The connection parameters for the plastic non-somatic connections in this model can be found in Tables 2 and 3. Since there is no rule-set switching, only the eta and theta terms in Table 2 are free (four free parameters). Although there was only one rule-set cell and therefor no need to clear the old rule-set the GPi received input as it did in the hierarchical task to maximize equivalence between the two implementations of

Table 2

*Parameter Values for the Plastic Non-Somatic Connections Between the Concrete Rule-Set Cell and the Stimulus Cell Outputs in the model for the Badre, Kaysre and D'Esposito Tasks*

$$w_{(concrete\ rule)\underset{ns}{\Rightarrow}\left((in\ stim)\underset{s}{\rightarrow}(PMC)\right)}$$

| Parameter | Value in Flat Task | Value in Hierarchical Task |
|---|---|---|
| $\eta_1$ | $3.0\cdot10^{-10}$ | $1.4\cdot10^{-9}$ |
| $\eta_2$ | $1.3\cdot10^{-15}$ | $1.8\cdot10^{-13}$ |
| $\theta_1$ | 34000 | 15000 |
| $\theta_2$ | 2000 | 6000 |
| $w_{max}$ | 4 | 4 |
| $w_{min}$ | 0 | 0 |
| $W_{init}$ | 0.7 | 0.7 |

Table 3

*Parameter Values for the Plastic Non-Somatic Connections Between the Abstract Rule-Set Cell and the Internal Cue Cell Outputs in the Model for the Badre, Kaysre and D'Esposito Tasks*

$$w_{(abstract\ rule)\underset{ns}{\Rightarrow}\left((external\ cue)\underset{s}{\rightarrow}(concrete\ rule)\right)}$$

| Parameter | Value in Flat Task | Value in Hierarchical Task |
|---|---|---|
| $\eta_1$ | 0 | $4.0 \cdot 10^{-10}$ |
| $\eta_2$ | 0 | $2.0 \cdot 10^{-14}$ |
| $\theta_1$ | 15000 | 20000 |
| $\theta_2$ | 5000 | 8000 |
| $w_{max}$ | 4 | 4 |
| $w_{min}$ | 0 | 0 |
| $W_{init}$ | 0.0 | 0.7 |

the model. For both models, a run consisted of 20 repetitions of each of the 18 possible stimuli presented in a random order. Data was collected over 1000 runs.

## Results

Each condition was run 1000 times and the results of each trial were averaged (Figure 20). As a reminder Bardre, Kayser, and D'Espositio used learning curve estimates for their figures (presented in Figure 19), not the means as used in Figure 20. The model was able to learn the associations in both conditions with qualitatively different learning curves. In the flat condition, accuracy rose linearly and reached an accuracy of approximately 0.7 by trial 300. In the human data for this condition, the accuracy started to asymptote at this point while the model's performance continued to rise linearly. An explanatory hypothesis for this difference is that the 18 unique combinations of features caused human subjects to suffer from interference effects. Although the simulation treats the images as 18 unique stimuli, each image actually shared many features with other images and not all of these other images had the same response association. This would cause interference in selecting the correct response. The model does not account for this kind of interference and so the asymptote does not appear in the model's results. In the Hierarchical condition accuracy rose much more steeply until trial 180 and then slowly approached an accuracy of 0.9. To calculate the $r^2$ of the model, subject data was estimated from the graphs in Figure 20 at trials 0, 60, 120, 180, 240, 300, and 360. The results yield an $r^2$ of 0.958. The failure of the model to reach the prefect accuracy achieved by the representative subject is not a universal failing. Over 70% of experimental runs reached an average accuracy of 0.9 or greater in the final 30 trials and 20% never made a mistake in that period.

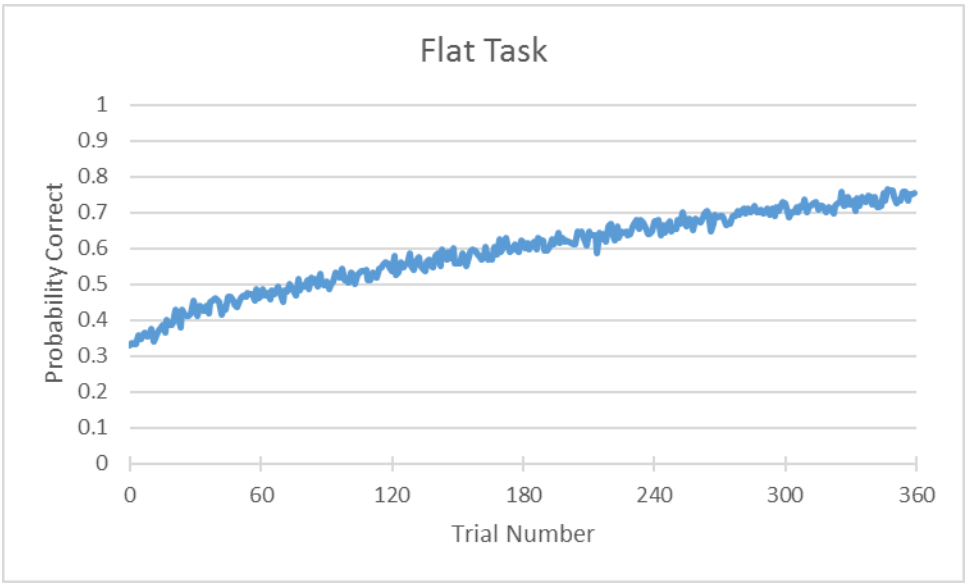


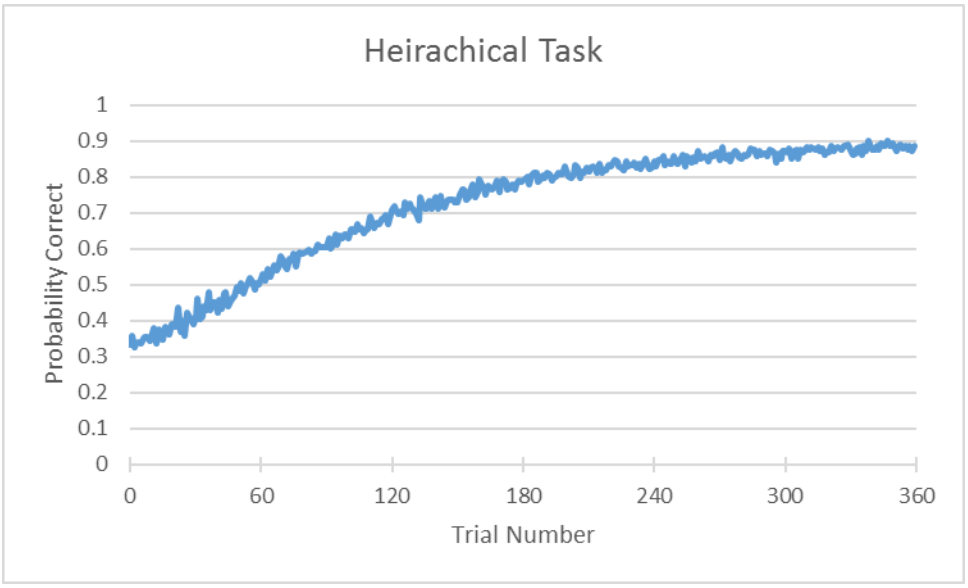


*Figure 20*. Model data from performing the task from Bardre, Kayser, and D'Espositio (2010).

## Discussion

The model demonstrated the ability to concurrently develop associations between stimuli/response pairs and cue/rule-set pairs. This allowed the model to learn correct responses faster when using a hierarchical representation compared to a flat representation. However, in this task each cue indicated a different dimension of the stimulus. Should a cue be incorrectly associated with the rule-set of the wrong dimension there is no way learning at the stimulus-response level will improve performance. In this situation, the model should eventually reverse that association. This is not the case when both rule-sets govern the same dimension, as is the case with the next task, presented in an article by Collins and Frank in 2013. In this task, there is a possibility that the stimulus-response associations of the two rule-sets could interfere with each other, becoming some hybrid of the two intended rule-sets. In such a situation, performance might be only good enough to avoid being unlearned while still being far below perfect. The Collins and Frank task will also investigate the model's ability to associate multiple cues with the same rule-set, thereby generalizing knowledge to novel stimuli.

# EXPERIMENT 2 – COLLINS AND FRANK 2013

The model is not restricted to using only one rule-set for a dimension. Multiple rule-sets can exist governing the same dimension. Neither are rule-sets limited to only using a single cue. Different cue features can be associated with the same rule-set. Thus allows learning to transfer into new contexts. In an experiment presented in a 2013 paper by Collins and Frank, subjects categorized colored shapes. Correct categorization was assigned such that if subjects used a particular task representation learning would transfer when subjects learned a second set of stimuli. Two implementations of the model were tested, one that used the representation conducive to learning transfer and an alternate representation that had been commonly used by subjects. This experiment will demonstrate that the model can learn multiple rule-sets on the same dimension and can use learning transfer to its advantage.

## Task

In the Collins and Frank task, subjects were trained to associate a stimulus that varied in two dimensions, shape and color, with 4 responses (Figure 21 [A]). There were 4 different colors and 2 different shapes. Subjects were trained over two phases. In the first phase two out of the four colors were used resulting in 4 stimuli, red triangles, yellow triangles, red circles, and yellow circles. Each of these stimuli was associated with a different response. There are multiple ways that the set of stimuli and responses could have been represented but most notably, if a hierarchical representation is used, there is no environmental bias toward or against using shape as a cue for rule-sets on the color dimension compared to using color as a cue for rule-sets on the shape dimension. In the second phase there is a shift in color – the presented shapes are now blue or green instead

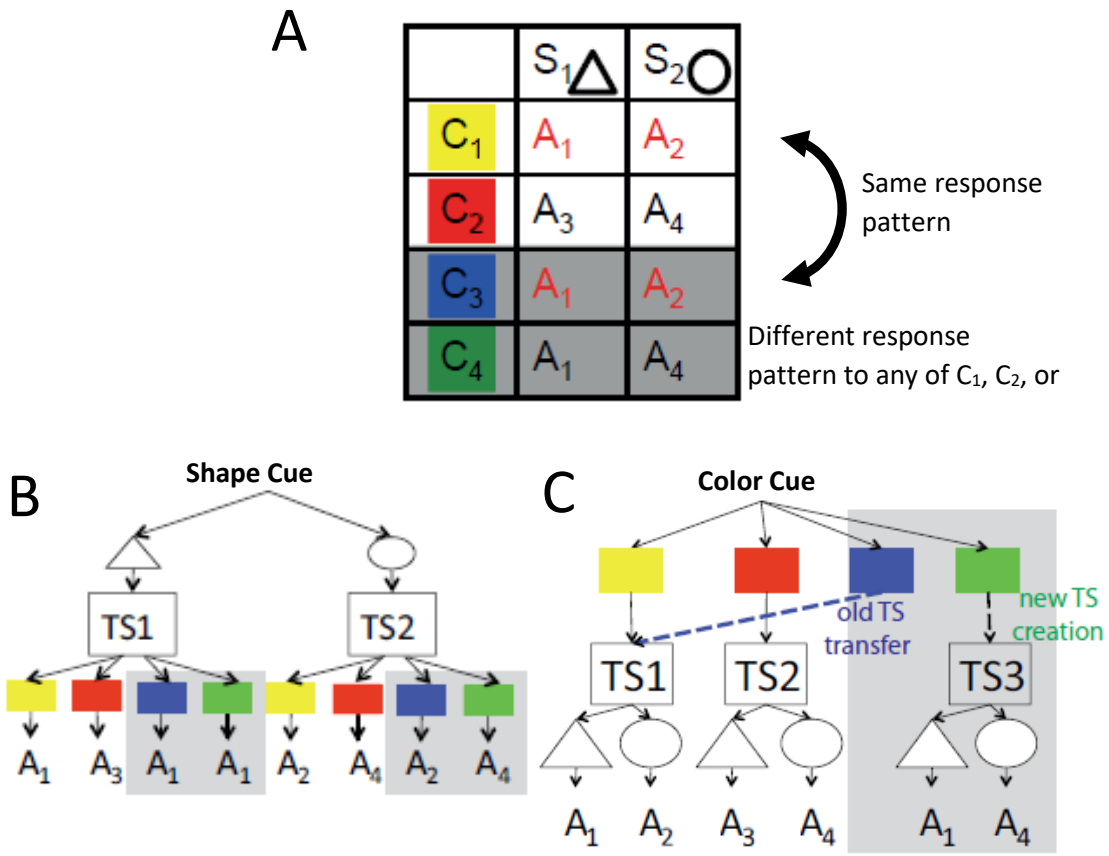


*Figure 21*. A) Stimulus-response (responses denoted as $A_1$, $A_2$, etc.) pairings for the Collins and Frank task. Note how yellow and blue stimuli have the same response pattern with regards to shape while green stimuli have a unique response pattern. B) The task representation that uses shape to select a task set (TS) and color to select a response. Note how in phase 2 response associations for blue and green stimuli are added to the existing task-sets. No transfer can take place. C) The task representation that uses color to select a task set (TS) and shape to select a response. Note how in phase 2 new task-sets are used for blue and green stimuli allowing blue stimuli to use the task-set already created for yellow stimuli. This enables transfer. Figures taken from Collins and Frank (2013).

of red or yellow. The category associations for blue shapes are the same as for yellow shapes (that is, yellow triangles and blue triangles share a response, and yellow circles and blue circles share a different response) but the correct responses for green shapes follow a new pattern (Figure 21 [A]).

The shift during phase 2 caused a dissociation between subjects that developed a representation that uses color as a cue and subjects that developed a representation that uses shape as a cue. If subjects used shape as a cue, there should have been no difference in the rate at which stimulus-response associations are learned for blue stimuli vs. green stimuli since associations for both needed to be added to the rule-sets cued by shape (Figure 21 [B]). If instead subjects used color as a cue, they would have been able to transfer the stimulus-response pairings learned for yellow stimuli to blue stimuli (Figure 21 [C]). That is, the color blue would have been associated with the same rule-set as the color yellow and thus subjects could have used concrete rules that had already been learned. Green, on the other hand, required a new set of rules and thus a new rule-set, all of which needed to be learned. In this case, the acquisition of the response pairings for the blue stimuli would have occurred faster than it would have for the green stimuli.

**Human Data**

For each subject, Collins and Frank calculated two reaction-time switch-costs, one for switches in stimulus color and one for switches in stimulus shape. They divided subjects into three groups based on the difference between these switch costs. Group 1 was made up of the top third of subjects with the highest color-switch cost minus shape-shift cost. Group 3 was made up of the bottom third of subjects with the lowest color-switch cost minus shape-shift cost. Group 2 was made of the remaining third. Collins and

Frank then examined the learning curves and error patterns across these groups (Figure 22). They looked at the types of errors made, whether subjects neglected the stimulus' shape when choosing a response, the stimulus' color, or both (Figure 22 insets). Group 1's learning curve and error patterns matched the predicted pattern for using color as a cue and while Group 3's learning curve and error patterns matched the predicted pattern for using shape as a cue. For subjects in Group 3 green and blue stimuli where learned at statistically similar rates and the types of errors made did not significantly differ between stimuli of those two colors. In Group 1, blue stimuli were learned more quickly than green stimuli and there were significantly more 'neglect color' errors made with green stimuli than with blue. Since Group 2 represented subjects whose task representations were harder to infer and who may have used alternative representations, their results are not considered.

**Task-Specific Model Parameters**

As in the task used by Badre, Kayser and D'Esposito, dynamically altering the problem representation is outside of the model's scope. Thus, as before, two versions of the model were used– one that used shape as the cue and one that used color as the cue.

The shape-cue model used two concrete rule-set cells but unlike the default model, both rule-sets are on the same stimulus dimension, color. This model used four stimulus cells in that dimension each representing a different color. There were four possible responses and thus four PMC cells. The color-cue model used 4 concrete rule-set cells with 4 corresponding thalamus cells and every rule-set cell represented a rule-set on the shape dimension. There were four external cue cells, one for each color. The task was identical to the task for the shape cue model and so again there are four responses and

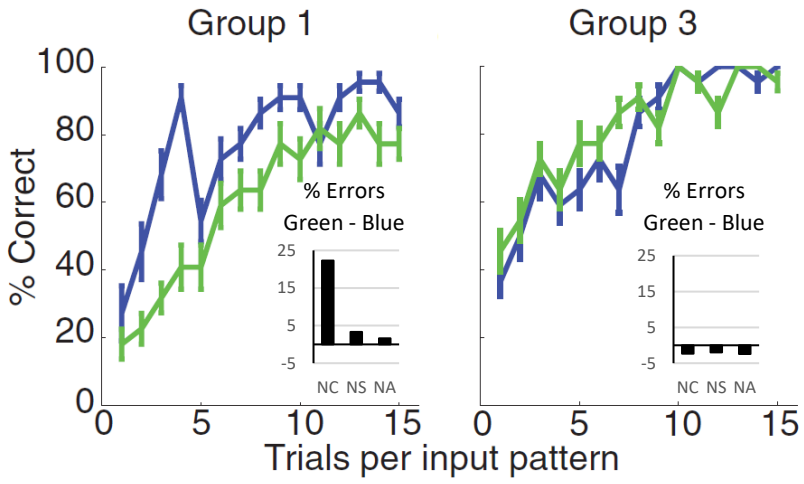


*Figure 22*. Human results for the third of subjects with the highest color-switch cost minus shape-shift cost (Group 1) and the third of subjects with the lowest color-switch cost minus shape-shift cost (Group 3). The graphs show average accuracy (error bars are standard error) and the inserts show green errors minus blue errors by error type: errors made due to neglecting color (NC), errors made due to neglecting shape (NS), and errors made due to neglecting both color and shape (NA). Figures taken from Collins and Frank (2013).

four PMC cells. The single stimulus dimension had only two stimulus cells, one representing circles and one representing triangles. Other parameters are presented in Table 4. For both models the weight of the non-somatic connections from the positive feedback cell to the stimulus-to-PMC connection and cue-to-concrete-rule-cell were set to 1.0 at the beginning of a run. The parameters for the plastic non-somatic connections from the abstract rule-set cell and concrete rule-set cells are presented in Tables 5 and 6. There were eight free parameters, the four eta terms and the four theta terms. For both models, a run consisted of 15 repetitions of each of the four stimuli that were colored red or yellow, presented in random order. This was followed by 15 repetitions of each of the four stimuli that were colored blue or green, again presented in random order. Data was collected over 1000 runs.

## Model Results

The mean results for the models across 1000 repetitions of the task are shown in Figure 23. The model for the Shape Cue condition achieved a maximum performance accuracy of 0.845 for stimuli of both colors. The errors made by this model are also the same across all colors and error types (see insets of Figure 23). In the Color Cue model, the blue stimuli reached an accuracy of 0.807 and the green stimuli reached an accuracy of 0.749. The $r^2$ of the model results was 0.937. While errors caused by neglecting shape or neglecting both were the same between the two colors, there were far more neglect color errors made with the green stimuli than with the blue stimuli.

## Discussion

The model learned the task well and showed itself capable of concurrently learning multiple rule-sets on the same dimension. Take note of the faster acquisition of

Table 4

*Fixed Connection Weights in the Model for the Collins and Frank Tasks That Differ From the Default Values*

| Weight Term | Mean |
|---|---|
| $w_{(pos)\underset{ns}{\rightarrow}\left((stim)\underset{s}{\rightarrow}(PMC)\right)}$ | 1.0 |
| $w_{(pos)\underset{ns}{\rightarrow}\left((external\ cue)\underset{s}{\rightarrow}(concrete\ rule)\right)}$ | 1.0 |

Table 5

*Parameter Values for the Plastic Non-Somatic Connections Between the Concrete Rule-Set Cell on the Stimulus Cell Outputs in the Model for the Collins and Frank Tasks*

$w_{(concrete\ rule)\underset{ns}{\Rightarrow}\left((stim)\underset{s}{\rightarrow}(PMC)\right)}$

| Term | Value |
|---|---|
| $\eta_1$ | $1.4 \cdot 10^{-9}$ |
| $\eta_2$ | $1.2 \cdot 10^{-13}$ |
| $\theta_1$ | 15000 |
| $\theta_2$ | 6000 |
| $w_{max}$ | 4 |
| $w_{min}$ | 0 |
| $W_{init}$ | 0.7 |

Table 6

*Parameter Values for the Plastic Non-Somatic Connections Between the Abstract Rule-Set Cell on the Internal Cue Cell Outputs in the Model for the Collins and Frank Tasks*

$w_{(abstract\ rule)\underset{ns}{\Rightarrow}\left((internal\ cue)\underset{s}{\rightarrow}(concrete\ rule)\right)}$

| Term | Value |
|---|---|
| $\eta_1$ | $3.0 \cdot 10^{-9}$ |
| $\eta_2$ | $1.6 \cdot 10^{-13}$ |
| $\theta_1$ | 25000 |
| $\theta_2$ | 13000 |
| $w_{max}$ | 4.0 |
| $w_{min}$ | 0.0 |
| $W_{init}$ | 0.7 |

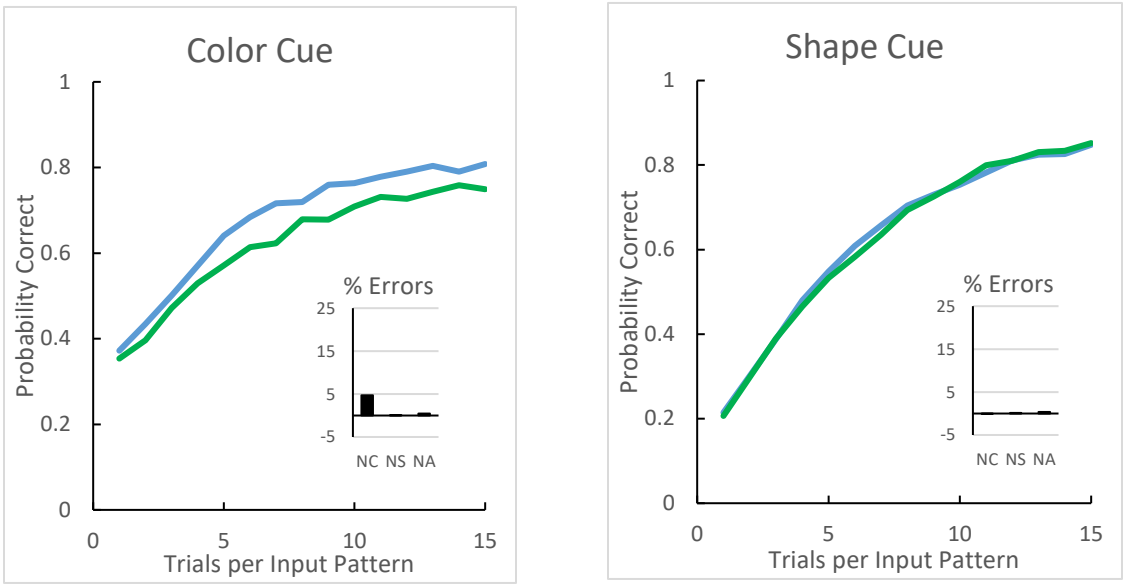


*Figure 23*. Model performance in the Collins and Frank task. ). The graphs show average accuracy and the inserts show green errors minus blue errors by error type: errors made due to neglecting color (NC), errors made due to neglecting shape (NS), and errors made due to neglecting both color and shape (NA).

blue stimuli vs. green stimuli when color was used as the cue. This shows that the model can transfer previously learned rule-sets into new environments instead of having to relearn them. Lastly, the model made more errors by neglecting color when categorizing green stimuli compared to blue stimuli in the color cue condition. These are the same error patterns that were found in human data and so it can be concluded that the model was, in some ways at least, using a method similar to that of human subjects.

The model has shown that it can successfully learn associations between stimuli and responses while simultaneously learning associations between cues and rule-sets. This has been demonstrated in a few different situations while reproducing several key advantages of using rule-sets. Next, the model will be tested on its ability to account for reaction-time data in task-switching experiments.

# EXPERIMENT 3 – ALPORT, STYLES, AND HSIEH 1994

The use of spiking neurons in this model allows it to generate response time data without adding any additional mechanism. Alport, Styles, and Hseih (1994) presents the first of two response time data sets that the model will reproduce. The data in question demonstrates a finding that alternating between rule-sets increases subject response time substantially. If the model can reproduce this data it would show that the model responds on the same timescale as human subjects and that time-pressured task switching has a similar effect on the model.

## Task

In their 1994 paper, Alport, Styles, and Hseih explored a number of manipulations of a two-choice decision task. They presented subjects with a list of 7 items. There were between 1 and 9 identical numerical digits displayed that could range from 1 to 9 in value. Five was excluded as both a possible numerosity and from the possible value of the individual digits. There were many conditions that were manipulated by Alport, Styles, and Hseih but the most relevant for this new model is an experiment where subjects had to determine if a relevant number was greater than or less than five for each item. For some lists the relevant number for every item was always the value of the digits that made up the item, for others it was always the number of digits, and for some the relevant number alternated between value and count for each trial. It should be noted that the next stimulus was presented immediately upon a subject responding to the prior stimulus. The subjects were informed ahead of time if the relevant dimension would be consistent or would alternate and what the relevant dimension of the first stimulus would be. The time taken to complete each list was recorded and compared.

## Human Data

Subjects took about 4 seconds on average to complete a list in the baseline condition (≈0.6 seconds per item) and about 10 seconds on average to complete a list in the alternating condition (≈1.4 seconds per item) (Figure 24). This showed that subjects took 2.5 times as long to respond to items in the alternating condition compared to lists where the task did not change. This is a severe example of switch-costs and the effects of interference from previously used rule-sets.

## Task-Specific Model Parameters

In this experiment the subjects were fully aware of what response should be made for each stimulus prior to the start of the experiment. Learning was not part of the task. Therefore, the weights of the non-somatic inhibition from the abstract and concrete rule-set cells were initialized at extreme values for the task and all learning was disabled. In addition, although only one of the two stimulus dimensions was relevant on any given trial, there was no cue. Subjects had to remember which dimension was relevant or, in the alternating condition, keep track of which dimension was used in the previous trial. While the model could retain the most recently used rule-set across trials, using information from previous trials to select a different rule-set goes beyond the scope of the model. Therefore, each stimulus was treated as though it included a cue. However instead of this cue representing some part of the stimulus, it represented the knowledge of the dimension that should be used in the current trial. This knowledge is assumed to be maintained in brain areas outside of the model and made available to areas within the model when required. Long term memory is a likely candidate for the area in question.

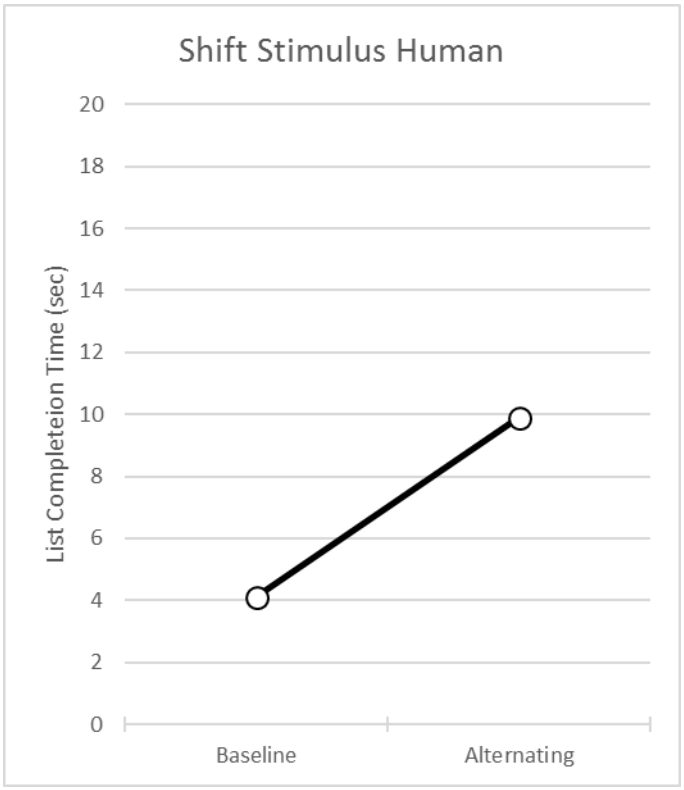


*Figure 24*. Human data from Alport, Styles, and Hseih (1994). Reproduced using figures in Alport, Styles, and Hseih (1994, shift-stimulus condition).

Movement from one item in a list to the next was controlled by the subject, who was instructed to complete the list as quickly as possible. To simulate this, the 500ms burn-in period at the beginning of each trail was removed. In the alternating condition it was assumed that the interference from the previous task causes a large delay in the generation/recall of the cue. Since there is no learning in this task, this delay and the response threshold (see below) were the only two free parameters. In this case, the cue was delayed by 850ms. This is a long time but, as Spector & Biederman (1976) showed, switch costs are particularly severe when correct task selection necessitates that subjects keep track of previous tasks. In addition, the presentation of the stimulus was delayed by the same amount so as to coincide with the presentation of the cue. The justification for this was the limited attentional resources of the subjects and their knowledge that processing the stimulus would be of no use without knowing which dimension was relevant. The period of time for which the cue and stimulus were delayed did count toward the calculated response time.

The Alport, Styles, and Hsieh task permitted only two responses so there were two corresponding PMC cells. Each of the stimulus dimensions had 8 cells corresponding to the numbers 1 through 9 excluding 5. The cue presented was determined by the condition being tested. Since learning had been disabled, the feedback cells do not send any signals to other cells in the model during the trials. The response threshold was set at 5000 to account for the greater time pressure on the subject.

A run consisted of 22 trials that used digit value as the relevant dimension, 22 trials that used digit count as the relevant dimension, and 22 trials where the relevant dimension alternated between those two. Stimuli were chosen randomly from the pool of

possibilities. The mean response time for each condition was calculated and multiplied by 7 to get the approximate time to complete a 7-item list. Data was collected over 1000 runs. Table 7 lists the connection strengths of every non-plastic synapse that varies from the default model. Tables 8 and 9 give the parameters for the plastic connections of the model.

## Model Results

The model took about 10 seconds to complete a list in the Alternating condition and 4 seconds in the Baseline condition (Figure 25). When compared to the human data, the $r^2$ coefficient is 1.0. This is an uninformative and trivial result of the fact that the data set has only two points. Therefore the root mean squared error (RMSE) is presented to give a more informative measure of fit. The RMSE is 0.195. This is a near perfect match to the human data that captures both the main effect and the individual data points.

## Discussion

The model reproduces the results almost exactly. The response time matches the human data in both the baseline and the alternating condition. This provides a good indication that this is a useful model of response time for categorization tasks. In addition, the model is able to simulate the effect of internally directed task switching by using a delayed cue presentation. An alternate implementation would be to inhibit the motor area (increase the response threshold) due to the uncertainty caused by interference from the alternate task. This ambiguity might be investigated using a task similar to this with the addition of severe time constraints (see Chapter 8).

Table 7

*Fixed Connection Weights in the Model for the Alport, Styles, and Hsieh Tasks That Differ From the Default Values*

| Weight Term | Mean |
|---|---|
| $w_{(stim)\underset{s}{\rightarrow}(PMC)}$ | 50 |
| $w_{(pos)\underset{ns}{\rightarrow}\left((stim)\underset{s}{\rightarrow}(PMC)\right)}$ | 0 |
| $w_{(neg)\underset{ns}{\Rightarrow}\left((stim)\underset{s}{\rightarrow}(PMC)\right)}$ | 0 |
| $w_{(pos)\underset{ns}{\rightarrow}\left((external\ cue)\underset{s}{\rightarrow}(concrete\ rule)\right)}$ | 0 |
| $w_{(neg)\underset{ns}{\Rightarrow}\left((external\ cue)\underset{s}{\rightarrow}(concrete\ rule)\right)}$ | 0 |

Table 8

*Parameter Values for the Plastic Non-Somatic Connections Between the Concrete Rule-Set Cell on the Stimulus Cell Outputs in the Model for the Alport, Styles, and Hsieh Tasks*

$$w_{(concrete\ rule)\underset{ns}{\Rightarrow}\left((in\ stim)\underset{s}{\rightarrow}(PMC)\right)}$$

| Term | Value |
|---|---|
| $\eta_1$ | 0 |
| $\eta_2$ | 0 |
| $\theta_1$ | 10000 |
| $\theta_2$ | 2000 |
| $w_{max}$ | 4 |
| $w_{min}$ | 0 |
| $W_{init}$ | 0 for correct connection<br>4 for all others |

Table 9

*Parameter Values for the Plastic Non-Somatic Connections Between the Abstract Rule-Set Cell on the Internal Cue Cell Outputs in the Model for the Alport, Styles, and Hsieh Tasks*

$$w_{(abstract\ rule)\underset{ns}{\Rightarrow}\left((internal\ cue)\underset{s}{\rightarrow}(concrete\ rule)\right)}$$

| Term | Value |
|---|---|
| $\eta_1$ | 0 |
| $\eta_2$ | 0 |
| $\theta_1$ | 35000 |
| $\theta_2$ | 2000 |
| $w_{max}$ | 4.0 |
| $w_{min}$ | 0.0 |
| $W_{init}$ | 0 for correct connection<br>4 for all others |

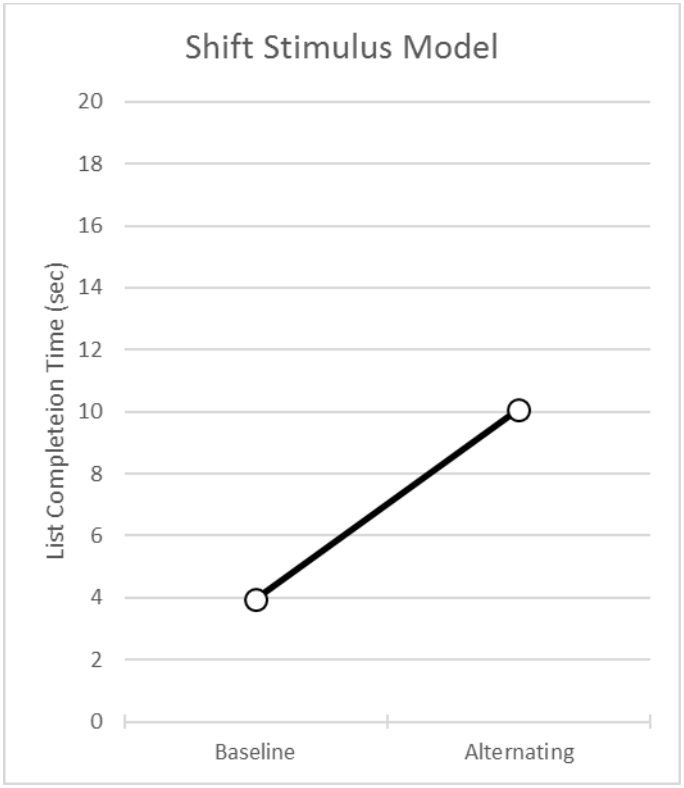


*Figure 25*. Model data for the Alport, Styles, and Hseih (1994) task.

Along with task-switching, priming has a well-established effect on reaction time (see Chapter 1). Ideally, the model will be able to reproduce both. This will be tested using the task and subject data from van ‘t Wout, Lavric, and Monsell (2015).

# EXPERIMENT 4 – VAN 'T WOUT, LAVRIC, AND MONSELL 2015

Response time can be reduced through priming as shown in Sudevan & Taylor (1987). In 2015 van 't Wout, et al. expanded upon this by manipulating the number of potential rule-sets that subjects might have to use and examining the interaction with priming. The effect of the number of rule-sets was minimal in all cases but their data presents a clear example of the effect of priming and thus is a good source of human data with which to test the model. In this experiment, the cue was presented in advance of the stimulus with a long or short interval between the two. With more time to process the cue, subjects respond more quickly. Can the model reproduce the effect of priming?

## Task

In the task used in van 't Wout, Lavric, and Monsell's 2015 paper, subjects were presented stimuli featuring stylized animals or trees. Both animals and trees varied in five dimensions (e.g. head size or body pattern for the animals and fruit shape or trunk thickness for the trees). Each dimension had only two possible values, which corresponded to one of the two possible responses. Subjects were given a verbal cue prior to the presentation of the stimulus that indicated which feature should guide the response. Subjects were trained for an extended period prior to data collection to minimize the effect of learning on the results. Once the extensive training was complete, subjects were tested under two pairs of conditions. In one pair, the first condition limited the relevant dimension to one of three dimensions while in the other condition any of the five dimensions of the stimulus could have been relevant. Subjects in the so-called three-task condition were aware that they were in this condition and understood the effect of the condition on the task. The second pair of conditions varied the interval between the

onset of the verbal cue and the presentation of the stimulus (Cue-Stimulus Interval - CSI) using values of either 100ms or 1300ms. All four combinations of these conditions were tested.

## Human Data

The mean response time for each condition was plotted above (Figure 26). The most relevant effect for the current model is the effect of manipulating CSI. Subjects that were given longer CSIs responded about 450ms after the stimulus was presented. Subjects that saw the shorter, control CSI responded after around 700ms. There is also a small effect of rule-set number: a larger pool of tasks slows reaction time slightly. However, this effect was not found to be significant in either condition.

## Task-Specific Model Parameters

As in the Alport, Styles, and Hsieh experiment the subjects were fully aware of and well versed in the associations between cues/stimuli and rule-sets/responses. So as before, the weights of the non-somatic inhibition from the abstract and concrete rule-set cells are set to the extreme values for the task and learning is disabled. For this task the burn-in time was increased to 1400ms and the overall runtime was increased to 3500ms to compensate. This reflects the 1400ms presentation of the fixation dot before the presentation of the stimulus in the subjects' task. In this task, only the stimulus was presented at the end of the burn-in time, and the point at which the cue was presented varies based on the experimental condition. Although van 't Wout, Lavric, and Monsell described the cue as presented 100ms or 1300ms before the stimulus, remember that it was a verbal cue. Speech has a time component and it arguably makes more sense to present the cue to the model at the off-set of the cue rather than the on-set. Van 't Wout,

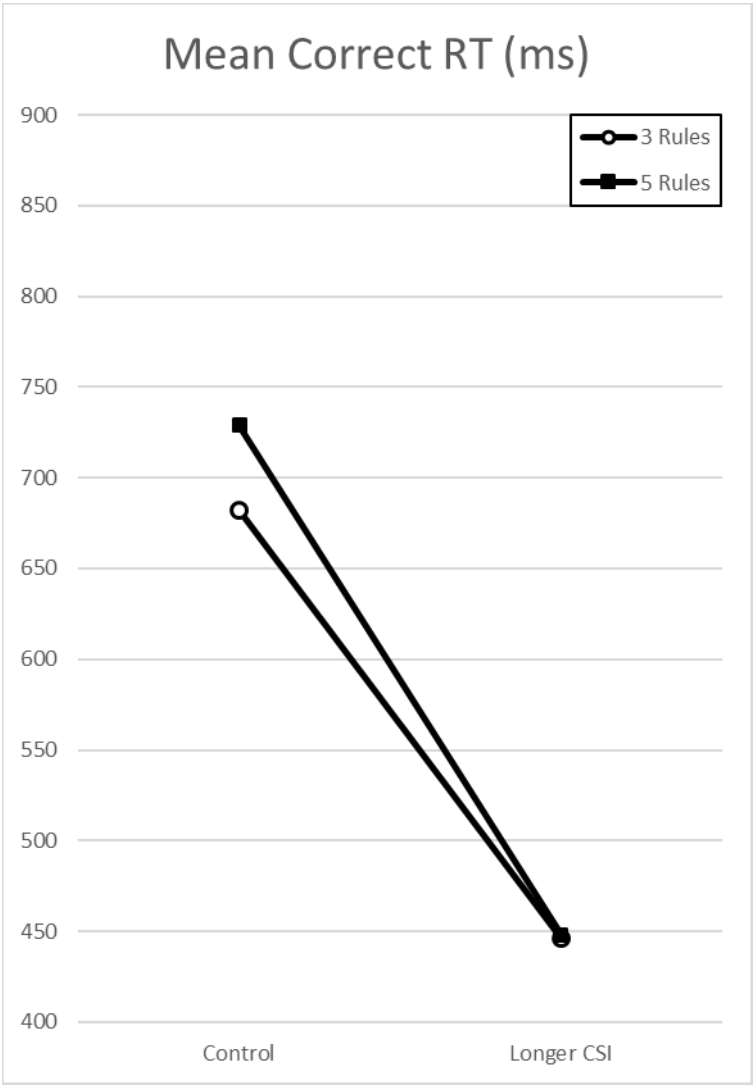


*Figure 26*. Human data from van ‘t Wout, Lavric, and Monsell (2015). Reproduced using figures from van ‘t Wout, Lavric, and Monsell (2015).

Lavric, and Monsell state that all of the verbal cues last 350ms and so in the model the cue was presented either 950ms before the stimulus presentation or 250ms afterward. In the case where the cue was presented after the stimulus, the presentation of the stimulus to the model was delayed to coincide with the presentation of the cue. This was done for the same reasons stated in the Alport, Styles, and Hsieh experiment: attentional limitations and necessity of the cue for useful stimulus processing.

The van ‘t Wout, Lavric, and Monsell task used two models, one representing the 3-rule-set condition, and one representing the 5-rule-set condition. The 3-rule-set model used three concrete rule-set cells with three corresponding thalamus cells and three external cue cells. Since the subjects were aware that certain rule-sets and cues were not going to occur, the cue cells and concrete rule-set cells corresponding to those cues and rule-sets were not included. All five stimulus dimensions were represented, even though two of them have no concrete rule-set cell associated with them. Each stimulus dimension contained 2 cells, each representing one of the possible values. In the 5-rule-set condition all five concrete rule-set cells/thalamus/cue cells were implemented. As before, there were five stimulus dimensions, each with two stimulus cells. All of the following parameters were the same between both models. The models used two PMC cells corresponding to the two possible responses. As with the model used in the Alport, Styles, and Hsieh task, the response threshold was reduced to 5000 to reflect time pressure but otherwise is determined as in the default model. Since there is no learning and the delay of the cue was determined by the task, this was the only free parameter. The response time is calculated relative to the external presentation of the stimulus, i.e. starting from the end of the burn-in time.

A run consisted of 576 trials that used only 3 possible relevant dimensions (6 of each possible stimulus and relevant dimension combination) with short CSI, 576 trials that used only 3 possible relevant dimensions with long CSI, 640 trials that used all 5 possible relevant dimensions (4 of each possible stimulus and relevant dimension combination) with short CSI, and 640 trials that all used 5 possible relevant with long CSI. The mean response time for correct trials was calculated for each condition using 1000 runs.

Table 10 lists the connection strengths of type every non-plastic synapse in the model that is different from the default. Tables 11 and 12 give the parameters for the plastic connections of the model.

## Model Results

The model reproduced the main effect of longer mean response time in short CSI conditions. The magnitude of the difference, a 250ms increase is reproduced as well (Figure 27). However, the model results show a longer response time than the human subjects show for both conditions. Although the model is slower than the human subjects, the pattern of results is well reproduced with an r2 of 0.99. The insignificant effect of higher number of rule-set slightly increasing reaction time was reproduced as well which is only notable because the model was not designed to reproduce it.

## Discussion

The relevant effects are well reproduced by the model. However, for both conditions the results show a longer response time than the human subjects. It should be pointed out that the model accounted for the faster response times in the Longer CSI condition in the same way that it accounted for the fast responses in the control condition

Table 10

*Fixed Connection Weights in the Model for the van 't Wout, Lavric, and Monsell Task That Differ From the Default Values*

| Weight Term | Mean |
|---|---|
| $w_{(stim)\underset{s}{\rightarrow}(PMC)}$ | 50 |
| $w_{(pos)\underset{ns}{\rightarrow}\left((stim)\underset{s}{\rightarrow}(PMC)\right)}$ | 0 |
| $w_{(neg)\underset{ns}{\Rightarrow}\left((stim)\underset{s}{\rightarrow}(PMC)\right)}$ | 0 |
| $w_{(pos)\underset{ns}{\rightarrow}\left((external\ cue)\underset{s}{\rightarrow}(concrete\ rule)\right)}$ | 0 |
| $w_{(neg)\underset{ns}{\Rightarrow}\left((external\ cue)\underset{s}{\rightarrow}(concrete\ rule)\right)}$ | 0 |

Table 11

*Parameter Values for the Plastic Non-Somatic Connections Between the Concrete Rule-Set Cell on the Stimulus Cell Outputs in the Model for the van 't Wout, Lavric, and Monsell Task*

$$w_{(concrete\ rule)\underset{ns}{\Rightarrow}\left((in\ stim)\underset{s}{\rightarrow}(PMC)\right)}$$

| Term | Value |
|---|---|
| $\eta_1$ | 0 |
| $\eta_2$ | 0 |
| $\theta_1$ | 10000 |
| $\theta_2$ | 2000 |
| $w_{max}$ | 4 |
| $w_{min}$ | 0 |
| $W_{init}$ | 0 for correct connection<br>4 for all others |

Table 12

*Parameter Values for the Plastic Non-Somatic Connections Between the Abstract Rule-Set Cell on the Internal Cue Cell Outputs in the Model for the van 't Wout, Lavric, and Monsell Task*

$$w_{(abstract\ rule)\underset{ns}{\Rightarrow}\left((internal\ cue)\underset{s}{\rightarrow}(concrete\ rule)\right)}$$

| Term | Value |
|---|---|
| $\eta_1$ | 0 |
| $\eta_2$ | 0 |
| $\theta_1$ | 35000 |
| $\theta_2$ | 2000 |
| $w_{max}$ | 4.0 |
| $w_{min}$ | 0.0 |
| $W_{init}$ | 0 for correct connection<br>4 for all others |

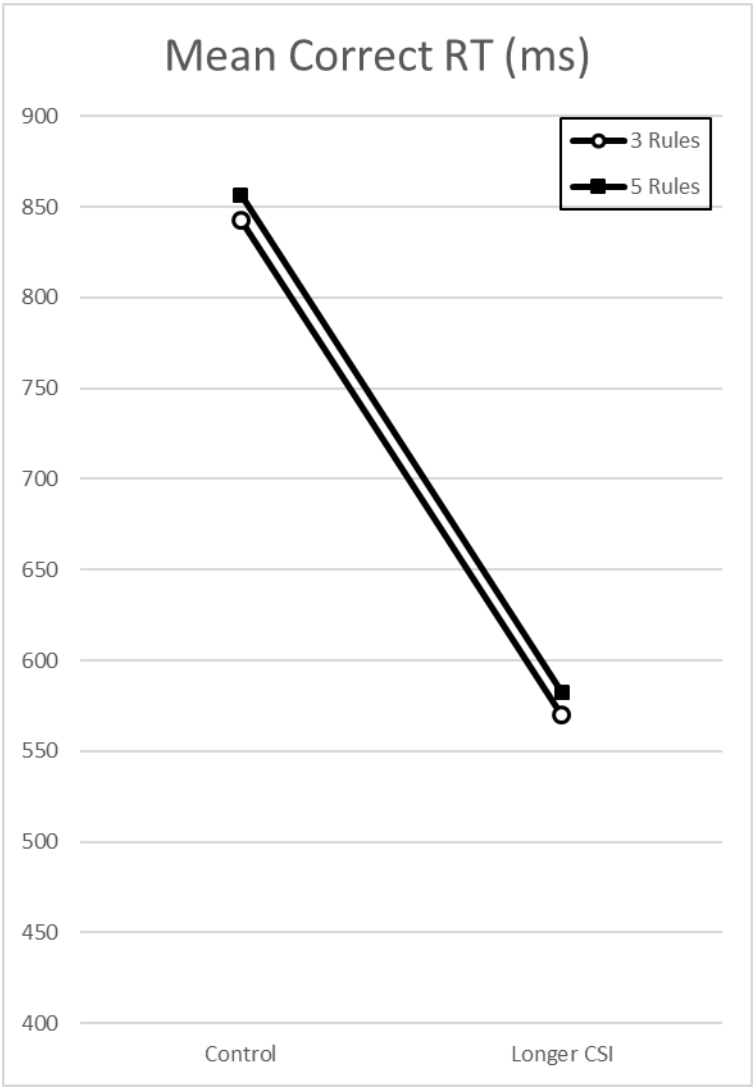


*Figure 27*. Model data for the van ‘t Wout, Lavric, and Monsell (2015) task.

of the Alport et al. experiment. However, an analysis of the response times shows that humans do not take the same amount of time to respond to these two tasks. Responding to seven consecutive stimuli in 4 seconds as seen in the results from Alport et al. (1994) yields an average of 582ms per item. This is the same reaction time that the model gives for the long CSI condition in van ‘t Wout et al. (2015). However the human data in van ‘t Wout et al. (2015) shows a response time of 450ms for that condition. Therefore, there must be some additional factor that this model did not account for which hastens reaction time in the long CSI condition of van ‘t Wout et al. (2015) but not the control condition of Alport et al. (1994). A possible explanation is that priming lowers the firing threshold of motor neurons, effectively decreasing the response threshold that the model should have had.

# CONCLUSIONS

It is sign of good design when a model can be applied to situations it was not specifically designed for. This model has shown that it can perform a variety of tasks and reproduce learning curves, error-type ratios, and reaction-time measures. Its design also avoids faults common in other models. Here is a review of the most important features and accomplishments of the model that demonstrate its versatility and plausibility.

## Key Features

Concrete rules are associations between stimulus features and responses and are implemented as inhibitory gates on the connections between feature representations and responses. These gates are plastic and can change based on reward and feedback. Weak inhibition (open gate) represents a positive association between the stimulus feature and the response. Concrete rule-sets are sets of complementary rules that are often specific to a task (i.e., a task set) and are stored in working memory. Concrete rule-sets are implemented in the posterior PFC as sources of the gating inhibition and therefore can gate every connection between the features and the responses. Each rule-set set can have a completely different pattern of gating. Concrete rule-sets are maintained in working memory through mutual parallel excitatory loops with the thalamus. This maintenance can be suppressed by means of an inhibitory signal from the basal ganglia. Abstract rules gate the connections between stimulus cues and concrete rule-sets, thereby associating cues and concrete rule-sets. The abstract rule-set is implemented in the same way as concrete rule-sets, albeit located in the anterior PFC, creating a recursive design.

## Accomplishments

The model respects several biological truths that often go ignored, most notably the learning mechanism and the integration of stimulus features and task-sets. Learning requires a reward signal that can keep up the feedback from trial to trial. The typical vector for this is the neurotransmitter dopamine (Schultz 2015). However, dopamine in the PFC does not have the temporal resolution required for learning and the basal ganglia is an impractical location for all PFC-related learning (see Chapter 2). Instead, this model uses feedback activated cells that facilitate advantageous Hebbian Learning at the inhibitory gates (Helie et al., 2015). These gates solve another common problem: combinatorial explosion. Both stimulus and task-set must be taken into account when selecting a response. Representing all possible combinations of these two already large dimensions with neurons results in an exponentially growing population of cells. By using synapse-level gating the explosion only results in a large number of axon terminals, which are far smaller than full neurons.

The model was tested against experimental results from several papers, covering both learning-curve data and response-time data. Running the model on the task used in Badre et al. (2010) replicated the finding that hierarchical task representations using multiple rule-sets facilitated learning better than flat representations with only a single concrete rule-set. The other set of learning curve data comes from Collins and Frank (2013) and explored the ability of rules learned in one context to be applied to another context. The model performed well, learning new rule-sets when required and using previously learned rule-sets when applicable.

Switch-cost refers to the additional time it takes to complete a task when the previous task in the sequence was different compared to when the previous task was the same as the current one. The model reproduced data from Alport et al. (1994) where reaction time for trials that alternated between two tasks was 2.5 times longer than trials where the task did not change. On switch trials, if subjects are given early knowledge of the new task and time to prepare, the switch cost is mitigated. This priming effect is observable in van ‘t Wout et al. (2015) and the model was able to replicate the data presented in that paper.

As seen, the model fits the biological constraints while being flexible enough to replicate a variety of empirical data.

## Implications and Predictions

The biggest predictions made by this model are about the role of the basal ganglia and the effect of storing task-sets in working memory. For example, if salient, attention-grabbing distractions were occasionally paired with stimuli presented in a task-switching experiment the model predicts that the distractor would enter working memory, ousting the task-set. When the subject refocuses on the task, the task-set would have to be re-generated. This would increase reaction time but moreover there would not be a difference between the reaction time for switch trials and stay trials.

Another, more speculative prediction involves the basal ganglia. The GPe is part of the indirect pathway of the basal ganglia while the GPi is part of the indirect and the direct pathway. Both pathways originate in the striatum, which inhibits both the GPe and the GPi. In the direct pathway the GPi is inhibited, freeing the thalamus. In the indirect pathway the GPe is inhibited, freeing the GPi to inhibit the thalamus. Putting this together

with the model's hypothesis that the thalamus maintains working memory creates the prediction that both pathways work together to constrain which items can enter working memory. Signals along the indirect pathway send general inhibition to the thalamus, making most items impossible to maintain in working memory. At the same time, the direct pathway is activated, sending more localized signals, protecting a few select WM items and allowing them to be maintained. These protected items would presumably be those that are relevant to the current situation.

## Extensions, Improvements, and Future Work

The model was shown to be robust and versatile but it is by no means perfect. Tasks like the WCST and intradimensional/extradimensional task involve rule-set switching based only on task performance without the use of cues. Even one of the primary inspirations for this model the HICL model (Helie et al., 2015) was designed to perform un-cued rule-set switching. Yet this model does not include a mechanism for switching rule-sets based on recent performance. Some of the pieces are already in place though. Un-cued task switching requires that the switches be infrequent so that the agent can discover and adapt to the change. This makes the thalamic maintenance mechanism of the model very useful. When a rule-set switch causes the performance to drop, the current rule-set should be cleared. This can be accomplished using a performance monitoring system that can send a reset signal. The biggest challenge is to prevent the abandoned rule-set from being reselected when selecting a new rule-set. Some attempts at this were made using cue cells that represented some internal classification of the environment but none made it into the final version.

A particularly ambitious way to add performance monitoring would be to incorporate a full model of memory. Not only would this allow for performance monitoring, it would also a much more meaningful exploration of cross-task interference in task-switching and priming studies. A confidence mechanic could also be useful should these phenomena be revisited.

Another potential improvement to the model is the addition of a concrete rule-set cell (or similar) that represents the situation where a rule-set has not been selected yet. Consider how the model behaves without such a cell. When the activation of the rule-set cells are still low and a winner has not been chosen none of the rule-set cells are active enough to substantially gate any of the connections from the stimulus cells to the response cells. This means every response cell is receiving strong excitatory signals from the active feature cells. The addition of a “no-rule-set” cell allows that cell to be active in the absence of a winning rule-set cell and to send blanket inhibition to every stimulus/response connection. That way activation of response cells is prevented until a rule-set has been selected. Future work could be devoted to exploring situations where the cue is presented after the other stimulus features.

Consider a task similar to the one used in the hierarchical condition of Badre et al. (2010) but the cue, border color, is presented a period of time after the 3D shape. Surely subjects will not respond until the cue is presented, thus plainly showing the need for some sort of placeholder rule-set. Yet that is not the only piece of information such an experiment can provide. Will subjects develop the same hierarchical task representation with border as the cue or will the delay change the representation. What other representations are there? Could the shape and orientation be the cues and the border

color be the stimulus feature? In such a case, would the dimensions be separated resulting in two simultaneously active rule-sets or would the dimension be combined resulting in 9 different rule-sets? Would priming effects occur? Exploring these possibilities would offer interesting insights into how versatile or limited rule-sets and the response selection system can be.

Another direction would be to broaden the scope of the model into the realm of animal studies. Ideally very little of the actual design of the model would need to be change but some capabilities may need to be scaled back to bring it more in line with the documented cognitive limitations of animal subjects. This would open up data from single cell recordings in addition to the behavioral data already modeled. The use of spiking neurons in the model means that the model would not need to be altered to produce this kind of data.

# LIST OF REFERENCES

Alexander, G. E., DeLong, M. R., & Strick, P. L. (1986). Parallel organization of functionally segregated circuits linking basal ganglia and cortex. *Annual Review of Neuroscience*, *9*(1), 357-381.

Altmann, E. M., & Gray, W. D. (2008). An integrated model of cognitive control in task switching. *Psychological Review*, *115*(3), 602-639.

Alport, A., Styles, E. A., & Hsieh, S. (1994). 17 Shifting Intentional Set: Exploring the Dynamic Control of Tasks.

Aron, A. R., Behrens, T. E., Smith, S., Frank, M. J., & Poldrack, R. A. (2007). Triangulating a cognitive control network using diffusion-weighted magnetic resonance imaging (MRI) and functional MRI. *Journal of Neuroscience*, *27*(14), 3743-3752.

Ashby, F. G., Ell, S. W., Valentin, V. V., & Casale, M. B. (2005). FROST: A distributed neurocomputational model of working memory maintenance. *Journal of Cognitive Neuroscience*, *17*(11), 1728-1743.

Ashby, F. G., & Helie, S. (2011). A tutorial on computational cognitive neuroscience: Modeling the neurodynamics of cognition. *Journal of Mathematical Psychology*, *55*(4), 273-289.

Badre, D., Hoffman, J., Cooney, J. W., & D'Esposito, M. (2009). Hierarchical cognitive control deficits following damage to the human frontal lobe. *Nature Neuroscience*, *12*(4), 515-522.

Badre, D., Kayser, A. S., & D'Esposito, M. (2010). Frontal cortex and the discovery of abstract action rules. *Neuron*, *66*(2), 315-326.

Badre, D., & Wagner, A. D. (2004). Selection, integration, and conflict monitoring: assessing the nature and generality of prefrontal cognitive control mechanisms. *Neuron*, *41*(3), 473-487.

Bamber, D. (1969). Reaction times and error rates for "same"-"different" judgments of multidimensional stimuli. *Perception & Psychophysics*, *6*(3), 169-174.

Buckner, R. L. (2003). Functional–anatomic correlates of control processes in memory. *Journal of Neuroscience*, *23*(10), 3999-4004.

Bunge, S. A. (2004). How we use rules to select actions: A review of evidence from cognitive neuroscience. *Cognitive, Affective, & Behavioral Neuroscience*, *4*(4), 564-579.

Bunge, S. A., & Zelazo, P. D. (2006). A brain-based account of the development of rule-set use in childhood. *Current Directions in Psychological Science*, *15*(3), 118-121.

Cass, W. A., & Gerhardt, G. A. (1995). In vivo assessment of dopamine uptake in rat medial prefrontal cortex: Comparison with dorsal striatum and nucleus accumbens. *Journal of Neurochemistry*, *65*(1), 201-207.

Chklovskii, D. B. (2004). Synaptic connectivity and neuronal morphology: Two sides of the same coin. *Neuron*, *43*(5), 609-617.

Collins, A. G., & Frank, M. J. (2013). Cognitive control over learning: Creating, clustering, and generalizing task-set structure. *Psychological Review*, *120*(1), 190-229.

Crone, E. A., Wendelken, C., Donohue, S. E., & Bunge, S. A. (2005). Neural evidence for dissociable components of task-switching. *Cerebral Cortex*, *16*(4), 475-486.

Dove, A., Pollmann, S., Schubert, T., Wiggins, C. J., & von Cramon, D. Y. (2000). Prefrontal cortex activation in task switching: An event-related fMRI study. *Cognitive Brain Research*, *9*(1), 103-109.

Duncan, J. (2001). An adaptive coding model of neural function in prefrontal cortex. *Nature Reviews Neuroscience*, *2*(11), 820-829.

Freedman, D., Riesenhuber, M., Poggio, T., & Miller, E. K. (2003). A comparison of primate prefrontal and inferior temporal cortices during visual categorization. *Journal of Neuroscience*, *23*(10), 5235-5246.

Fuster, J. M. (2008). *The prefrontal cortex*. London, England: Academic Press.

Grant, D. A., & Berg, E. (1948). A behavioral analysis of degree of reinforcement and ease of shifting to new responses in a Weigl-type card-sorting problem. *Journal of Experimental Psychology*, *38*(4), 404-411.

Helie, S., Ell, S. W., Filoteo, J. V., & Maddox, W. T. (2015). Criterion learning in rule-based categorization: Simulation of neural mechanism and new data. *Brain and Cognition*, *95*, 19-34.

Helie, S., Roeder, J. L., & Ashby, F. G. (2010). Evidence for cortical automaticity in rule-based categorization. *Journal of Neuroscience*, *3*0(42), 14225-14234.

Hübner, R., Futterer, T., & Steinhauser, M. (2001). On attentional control as a source of residual shift costs: Evidence from two-component task shifts. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *27*(3), 640-653.

Izhikevich, E. M. (2007). *Dynamical systems in neuroscience*. Cambridge, MA: MIT Press.

Kida, S., Josselyn, S. A., de Ortiz, S. P., Kogan, J. H., Chevere, I., Masushige, S., & Silva, A. J. (2002). CREB required for the stability of new and reactivated fear memories. *Nature Neuroscience*, *5*(4), 348-355.

Koch, I., & Allport, A. (2006). Cue-based preparation and stimulus-based priming of tasks in task switching. *Memory & Cognition*, *34*(2), 433-444.

Koechlin, E., Ody, C., & Kouneiher, F. (2003). The architecture of cognitive control in the human prefrontal cortex. *Science*, *302*(5648), 1181-1185.

Logan, G. D., & Bundesen, C. (2003). Clever homunculus: Is there an endogenous act of control in the explicit task-cuing procedure? *Journal of Experimental Psychology: Human Perception and Performance*, *29*(3), 575-599.

Miller, E. K., Freedman, D. J., & Wallis, J. D. (2002). The prefrontal cortex: Categories, concepts and cognition. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, *357*(1424), 1123-1136.

Monsell, S., & Mizon, G. A. (2006). Can the task-cuing paradigm measure an endogenous task-set reconfiguration process? *Journal of Experimental Psychology: Human Perception and Performance*, *32*(3), 493-516.

Newell, A. (1992). SOAR as a unified theory of cognition: Issues and explanations. *Behavioral and Brain Scienc*es, *15*(3), 464-492.

O'Doherty, J., Kringelbach, M. L., Rolls, E. T., Hornak, J., & Andrews, C. (2001). Abstract reward and punishment representations in the human orbitofrontal cortex. *Nature Neuroscience*, *4*(1), 95-102.

O'Reilly, R. C., Noelle, D. C., Braver, T. S., & Cohen, J. D. (2002). Prefrontal cortex and dynamic categorization tasks: Representational organization and neuromodulatory control. *Cerebral Cortex, 12*(3), 246-257.

Reynolds, J. R., Braver, T. S., Brown, J. W., & Van der Stigchel, S. (2006). Computational and neural mechanisms of task switching. *Neurocomputing, 69*(10-12), 1332-1336.

Rogers, R. D., & Monsell, S. (1995). Costs of a predictible switch between simple cognitive tasks. *Journal of Experimental Psychology*: *General, 124*(2), 207-231.

Rumelhart, D. E., & Zipser, D. (1985). Feature discovery by competitive learning. *Cognitive Science, 9*(1), 75-112.

Rushworth, M. F. S., Hadland, K. A., Paus, T., & Sipila, P. K. (2002). Role of the human medial frontal cortex in task switching: A combined fMRI and TMS study. *Journal of Neurophysiology, 87*(5), 2577-2592.

Schneider, D. W., & Logan, G. D. (2005). Modeling task switching without switching tasks: A short-term priming account of explicitly cued performance. *Journal of Experimental Psychology: General, 134*(3), 343-367.

Seamans, J. K., & Robbins, T. W. (2010). Dopamine modulation of the prefrontal cortex and cognitive function. In K. A. Neve (Ed.), *The dopamine receptors* (pp. 373-398). Totowa, NJ: Humana Press.

Shepherd, G. M. (2004). Introduction to synaptic circuits. In G. M. Shepherd (Ed.), *The synaptic organization of the brain* (5th ed., pp. 1-38). Oxford, England: Oxford University Press.

Spector, A., & Biederman, I. (1976). Mental set and mental shift revisited. *The American Journal of Psychology*, *89*(4), 669-679.

Stroop, J. R. (1935). Studies of interference in serial verbal reactions. *Journal of Experimental Psychology*, *18*(6), 643-662.

Sudevan, P., & Taylor, D. A. (1987). The cuing and priming of cognitive operations. *Journal of Experimental Psychology: Human Perception and Performance*, *13*(1), 89-103.

van't Wout, F., Lavric, A., & Monsell, S. (2015). Is it harder to switch among a larger set of tasks?. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *41*(2), 363.

Wallis, J. D., & Miller, E. K. (2003). From rule-set to response: Neuronal processes in the premotor and prefrontal cortex. *Journal of Neurophysiology*, *90*, 1790-1806.

Zanolie, K., Van Leijenhorst, L., Rombouts, S. A. R. B., & Crone, E. A. (2008). Separable neural mechanisms contribute to feedback processing in a rule-learning task. *Neuropsychologia*, *46*(1), 117-126.