Criterion learning in rule-based categorization: Simulation of neural mechanism and new data

Sebastien Helie[a], Shawn W. Ell[b], J. Vincent Filoteo[c], W. Todd Maddox[d]

[a] Department of Psychological Sciences, Purdue University

[b] Department of Psychology, University of Maine & Maine Graduate School of Biomedical Sciences and Engineering

[c] VA San Diego Healthcare System & University of California, San Diego

[d] Department of Psychology, University of Texas, Austin

In perceptual categorization, rule selection consists of selecting one or several stimulus-dimensions to be used to categorize the stimuli (e.g, categorize lines according to their length). Once a rule has been selected, criterion learning consists of defining how stimuli will be grouped using the selected dimension(s) (e.g., if the selected rule is line length, define 'long' and 'short'). Very little is known about the neuroscience of criterion learning, and most existing computational models do not provide a biological mechanism for this process. In this article, we introduce a new model of rule learning called Heterosynaptic Inhibitory Criterion Learning (HICL). HICL includes a biologically-based explanation of criterion learning, and we use new category-learning data to test key aspects of the model. In HICL, rule selective cells in prefrontal cortex modulate stimulus-response associations using pre-synaptic inhibition. Criterion learning is implemented by a new type of heterosynaptic error-driven Hebbian learning at inhibitory synapses that uses feedback to drive cell activation above/below thresholds representing ionic gating mechanisms. The model is used to account for new human categorization data from two experiments showing that: (1) changing rule criterion on a given dimension is easier if irrelevant dimensions are also changing (Experiment 1), and (2) showing that changing the relevant rule dimension and learning a new criterion is more difficult, but also facilitated by a change in the irrelevant dimension (Experiment 2). We conclude with a discussion of some of HICL's implications for future research on rule learning.

Keywords: rule-based categorization, criterion learning, prefrontal cortex, pre-synaptic inhibition, Hebbian learning.

1 Introduction

Rule-guided behavior is essential in adapting one's actions to the ever changing environment (Bunge & Souza, 2007; Miller & Cohen, 2001). For example, rule-use allows for a direct generalization of performance to new stimuli or new situations (Helie & Ashby, 2012). As a result, rule learning is essential for simple everyday tasks such as classifying new animals as friends or foes, or learning the meaning of traffic signs. Once a rule has been learned, rules can be applied in a variety of contexts. Rule learning is to be contrasted with associative learning (i.e., gradually learning that a stimulus is predictive of a particular response, e.g., Erickson & Kruschke, 1998; Maddox & Ashby, 2004). Differences between rule learning and associative learning may be the result of different brain circuits being recruited (Ashby, Alfonso-Reese, Turken, & Waldron, 1998). Specifically, associative learning is thought to rely primarily on the basal ganglia while rule learning uses a working memory network relying primarily on the prefrontal cortex (PFC). While much computational work has been devoted to the biological network responsible for associative learning (e.g., Ashby & Crossley, 2011; Ashby, Ennis, & Spiering, 2007; for a review, see Helie, Chakravarthy, & Moustafa, 2013), the biological network responsible for rule learning remains largely unexplored by computational modelers.

Rule learning involves at least two different cognitive operations, namely rule selection and criterion learning. In perceptual categorization, rule selection consists of selecting one or several stimulus dimensions to be used to categorize the stimuli (e.g., categorize lines according to their length). Once a rule has been selected, criterion learning consists of defining how stimuli will be grouped using the selected dimension(s) (e.g., if the selected rule is line length, define 'long' and 'short'). Ashby and colleagues (1998) have proposed a biologically-based model of how rules can be selected or switched away from, but the rule criterion is generally pre-inserted in the model, and no biological account of criterion representation or criterion learning has been proposed.

In this article, we propose a new biologically-based model of criterion learning and use category-learning data to test key components of the model. The new model is called Heterosynaptic Inhibitory Criterion Learning (HICL). The key aspect of HICL is that rule selective cells in lateral PFC are used to modulate stimulus-response associations using pre-synaptic inhibition (Helie & Ashby, 2009). Each neuron is implemented using two differential equations to calculate its membrane potential at every millisecond (Izhikevich, 2007). In HICL, criterion learning is implemented by a new type of heterosynaptic error-driven Hebbian learning at GABAergic (i.e., inhibitory) synapses that uses feedback to drive cell activation above/below gating thresholds. HICL is used to account for new human behavioral data that were collected and simulated to explore the effect of intra-dimensional (ID) and extra-dimensional (ED) shifts on criterion learning. Specifically, Experiment 1 (ID) tests the ability of participants to consecutively learn more than one criterion on the same stimulus dimension while Experiment 2 (ED) tests the ability of participants to learn more than one criterion by consecutively selecting new stimulus dimensions and learning new criteria. The model was used to simulate the new experiments with a single set of parameters. Additional results and predictions made by the computational cognitive neuroscience model will also be discussed.

The remainder of this article is organized as follows. First, we briefly review research on the cognitive neuroscience of rule-guided behavior to establish some biological constraints that HICL must satisfy. Second, we discuss how rules are represented and learned in HICL before simulating a toy problem to explore how the information is propagated in the model as well as how rule criteria can be learned. Third, we present results from two new human behavioral experiments that were designed to test the new model. In these experiments, each participant was asked to learn more than one criterion, thus providing more criterion learning data to test the model. The new data are simulated with HICL. Fifth, we conclude with a discussion of how HICL can impact future cognitive neuroscience research on rule learning.

2   Rules in the brain

Early clinical studies have shown that damage to the prefrontal areas impairs performance in a rule-based card sorting task (see Bunge, 2004; Fuster, 2008 for reviews). As a result, most neurobiological studies of rule learning and rule use have focused on the role of the PFC (e.g., Asaad, Rainer, & Miller, 1998; Miller & Cohen, 2001; White & Wise, 1999; for a review, see Bunge, 2004). Rule-selective cells have been found in the lateral PFC of rhesus monkeys in a rule-based task (Wallis & Miller, 2003), and Helie, Roeder, and Ashby (2010) found that the BOLD signal in the PFC was related to rule learning in humans. In addition, PFC lesions have been associated with impaired rule shifting (e.g., Dias et al., 1996; Owen et al., 1993). Hence, there is strong indication that lateral PFC plays a critical role in supporting rule-guided behavior.

However, the role of lateral PFC in rule application is not restricted to rule representation. For example, Wallis and Miller (2003) also found stimulus selective cells in the PFC for discrete categorical stimuli (i.e., where each stimulus was the only category member). This result was further explored by Freedman et al. (2003) using continuous-valued stimuli and the results show that, while individual stimuli are typically represented in higher visual areas (e.g., inferotemporal cortex), the PFC represented visual stimuli in a categorically informative way. For example, Freedman et al. used a morphed animal that could continuously vary between a cat and a dog. Most stimulus-selective cells in the PFC would selectively fire to cats or dogs, without much differentiation within the cat and dog categories. In addition, Freedman et al. showed that varying irrelevant stimulus dimensions did not significantly affect the visual representation in the PFC.[1] Hence, the lateral PFC also plays an important role in producing the visual representation used for the application of learned rules, and this representation already contains some rule-relevant information (e.g., relevant and irrelevant stimulus features).

To summarize, electrophysiological data collected from monkeys and fMRI data from humans producing rule-guided behavior suggest that lateral PFC contains a categorical representation of the stimuli in addition to rule-selective cells. These important results were taken into consideration in the design of the proposed rule learning model.

---

[1] But see Folstein, Palmeri, & Gauthier (2013) for evidence of category-learning-related changes in visual cortex for some morphing spaces.

3   A neural model of rule representation

In addition to the biological results described above, one important behavioral constraint when implementing rules in a network is flexibility (Ashby, Ell, & Waldron, 2003). Specifically, unlike associative learning, the stimulus-response associations can be changed flexibly in rule-based categorization (i.e., without affecting accuracies), which makes it unlikely that rule learning would depend exclusively on weighted connections between stimulus and response representations (as in, e.g., backpropagation networks). Keeping this constraint in mind, Helie and Ashby (2009) proposed that rules could be coded in the brain using pre-synaptic inhibition (Shepherd, 2004). The proposed model of rule implementation includes three populations of neurons: (1) rule cells, (2) categorical stimulus cells and, (3) motor planning cells. The first two populations of cells are located in lateral PFC while the last cell population is located in the premotor cortex. The rule cells are maintained actively in working memory, and each population of rule cells represents a different stimulus dimension (Helie & Ashby, 2009). The Helie and Ashby model of rule implementation has been used to simulate electrophysiological data collected in monkeys performing rule-guided behavior and will be used as a starting point for rule implementation in HICL. This model is conceptually described below.

When selecting a rule in perceptual categorization, one essentially selects a stimulus dimension (Ashby et al., 1998). For example, if categorizing visual stimuli with various shapes and colors, selecting a rule would mean selecting a dimension (e.g., stimulus color), and there would be a different population of rule cells representing each stimulus dimension (e.g., shape, color).

The selected rule cell population affects the categorical stimulus representation in lateral PFC (Freedman et al., 2003). In lateral PFC, stimuli are represented using only the dimension relevant for application of the selected rule. If the selected dimension is discrete (e.g., stimulus shape), then each population of neurons represents a different shape (and ordering is not meaningful). Staying with the previous example, if the selected rule is shape, then the stimulus representation in lateral PFC will only represent the stimulus shape. Accordingly, a blue triangle and a red triangle would have the same representation in lateral PFC, but a different representation if the selected rule would have been color. However, if the selected dimension is continuous (e.g., stimulus size), then each cell population is a fuzzy measurement on the selected dimension, and neighboring cells have slightly bigger (smaller) measurements on the selected dimension. Note, however, that objects are still represented as a whole in inferotemporal cortex, so that a blue triangle and a red triangle are represented differently in specialized visual areas (Folstein et al., 2013; Freedman et al., 2003).

The last population of neurons, motor planning cells, is located in premotor cortex and corresponds to category assignments. For example, if the task is to produce a different motor response indicating whether the stimulus is a member of the 'A' or 'B' category, than one cell population corresponds to the 'A' response while the other corresponds to the 'B' response. Note that the model response corresponds to the motor plan – not the actual motor action (e.g., kinematics). Motor actions are assumed to be located further downstream and are outside the scope of the proposed model.

Figure 1a shows the connectivity between the three cell populations included in the model. Because rules need to allow for flexible stimulus—response mappings (Ashby et al., 2003), every categorical stimulus cell population is connected to all task-relevant motor planning

cell populations using a fixed connection weight. The fixed connection weight is the same for all population pairs, so these connections do not allow for categorization because every stimulus equally activates every response. Instead, Helie and Ashby (2009) showed that categorization can be obtained by using pre-synaptic inhibition (Shepherd, 2004). In pre-synaptic inhibition, a pre-synaptic terminal is post-synaptic to another terminal. This is illustrated in Figure 1b, where the pre-synaptic terminal of the categorical stimulus cell is post-synaptic to the terminal of the rule cell. The rule cell activity is relayed through an inhibitory (GABA) neuron (not represented), so its role is to gate the excitatory (glutamate) activity flowing from the categorical stimulus cells to the motor cells. The gate provided by the rule cell (represented by the grey rectangles in Figure 1a) is continuous, and can affect each synapse differently between the categorical stimulus cells and the motor planning cells. For example, Figure 1a has four stimulus cells and two motor planning cells, so there are eight synapses, each with their own amount of pre-synaptic inhibition (represented by negative connection weights). At the psychological level, the rule cell allows for closing the gate and preventing the circuit from initiating responses that should not be made.

INSERT FIGURE 1 ABOUT HERE

In addition to the connectivity described above, the circuit includes lateral inhibition between the motor planning cells (also through GABAergic interneurons), and (selective) increased activation of the rule cells after a response has been made to turn off the motor planning cell once the rule has been applied and a response has been made. The former implements a soft-winner take-all for category selection, and the latter ensures that the motor planning cells are 'turned off' after a response has been made.

4   Formalization of the neural model

The neural model described above was implemented using a neurocomputational model. All the units in the network were formalized using the Izhikevich (2007) simple model. In the model implementation, each unit represents a group of tightly connected neurons that fire together, so each unit should be understood as a sparse neural representation (Bowers, 2009). Specifically, we used the parametrization corresponding to pyramidal cells[2], the main projection neurons in cortex:

$$100\frac{dV(t)}{dt} = 0.7[V(t)+60][V(t)+40] - U(t) + I(t) + \varepsilon(t)$$

$$\frac{U(t)}{dt} = 0.03\{-2(V(t)+60) - U(t)\}$$

(1)

where $V(t)$ is the membrane potential (in mV) at time t, $I(t)$ is the input at time t, and $\varepsilon(t)$ is Gaussian white noise at time t (with variance σ_ε^2), and $U(t)$ is an abstract regulatory term that is meant to describe slow recovery after an action potential is initiated. For instance, $U(t)$ could represent activation in the K^+ current or

---

[2] Pyramidal cells are glutamatergic, so all inhibitory connections are relayed through GABA interneurons that are not explicitly simulated in the model.

inactivation in the Na$^+$ current, or some combination of both. When $V(t)$ reaches 35 mV, a spike is recorded and $V(t)$ is reset to -50 mV. In addition, $U(t)$ is reset to $U(t) + 100$.

Before connecting the network, we need to specify a model of neural propagation. In this article, we focus on modeling the temporal delays of spike propagation and the temporal smearing that occurs at the synapse. A standard solution is to use an α-function (e.g., Rall, 1967). The output of a unit at time t is:

$$f[V(t)] = \sum_{s \in S} \left[\frac{t-s}{\lambda} e^{\left[1 - \left(\frac{t-s}{\lambda} \right) \right]} \right]^+ \tag{2}$$

where λ is a constant that can be fixed to model any desired temporal delay, S is the set of spike times and $[h(t)]^+$ equals $h(t)$ when $h(t) > 0$, and 0 when $h(t) \leq 0$. For example, $S = \{5, 17, 29\}$ means that a spike was recorded at 5, 17, and 29 milliseconds (msecs). Equation 2 is referred to as the *output* of the unit.

The input to the selected rule unit is a constant square function that is equally active for the entire trial duration so that $I_{rule}(t) = amp_{rule}$ for all t. This constant input represents the fact that the rule unit is maintained active in working memory. The selected stimulus unit is also activated using a square function that is constant so that $I_{stim}(t) = amp_{stim}$ when the stimulus is presented and $I_{stim}(t) = 0$ when the stimulus is not presented. The input to motor planning unit A is:

$$I_{motor,A}(t) = \sum_i \left[w^{stim \rightarrow motor} \times f[V_{stim,i}(t)] - w^{rule \rightarrow motor(i,A)} \times \left(f[V_{rule}(t)] + w^{motor \rightarrow rule} \times g[V_{motor,A}(t)] \right) \right]^+$$
$$- \sum_{j \neq A} w^{motor \rightarrow motor} \times f[V_{motor,j}(t)] \tag{3}$$

where $I_{motor,A}(t)$ is the input to motor planning unit A at time t, $f[V_{stim,i}(t)]$ is the output of stimulus unit i at time t, $f[V_{rule}(t)]$ is the output of the selected rule unit at time t, $f[V_{motor,j}(t)]$ is the output of motor planning unit j at time t, $g[V_{motor,A}(t)] = f[V_{motor,A}(t)]$ if a response has been selected and 0 otherwise[3], $w^{stim \rightarrow motor}$ is the excitatory connection between the stimulus and motor units, $w^{rule \rightarrow motor(i,A)}$ is the pre-synaptic inhibition from the selected rule unit to the {stimulus(i), motor(A)} synapse, $w^{motor \rightarrow rule}$ is the excitatory connection from the motor unit to the rule unit axon (to turn off the motor unit after the rule has been applied to produce a response), and $w^{motor \rightarrow motor}$ is the lateral inhibition between the motor units. In words, the first line of Eq. 3 is the activation from the categorical stimulus units, which is moderated by the pre-synaptic inhibition from the rule unit. The second line is the lateral inhibition received from the other motor planning units.

---

[3] This input comes from a motor region downstream from the model to inform the model that a response has been made. However, we use the motor unit activation because (1) they are assumed to be directly connected to the motor response and (2) the motor response is not included in the model.

For the decision rule, we used the following simple decision model: At every time step (msec), if $\int_0^t f\left[V_{motor,j}(t)\right]dt \geq \tau$, set the model response to the most active motor planning unit and set the response time to t. Otherwise, calculate another msec of model time (where τ is the response threshold). In all the simulations included in this article, $\tau = 40,000$.

5 Synaptic plasticity

Sections 3-4 reviewed the biological model of rule representation of Helie and Ashby (2009). However, the connections between the rule unit and the axons of the categorical stimulus units needed to be fixed a priori. In this section, we introduce a new learning algorithm to iteratively adjust the weights of the rule unit to allow HICL to learn a rule criterion using trial-and-error. The major difficulties are (1) to introduce a trial-and-error learning algorithm without using dopamine as a reinforcement signal and (2) designing a biologically-realistic learning algorithm for inhibitory (GABA) synapses. We will now elaborate on these two problems in turn, before describing the proposed solutions.

5.1 Trial-and-error learning in cortex

Learning in the brain is often described as a form of reinforcement learning, and dopamine (DA) typically plays the role of the reinforcement signal (Glimcher, 2011). However, a necessary feature of any reinforcement learning training signal is high temporal resolution (Helie, Ell, & Ashby, 2015). If the first response is correct then DA must be released into the relevant synapses quickly, before the critical traces disappear. But after the correct synapses have been strengthened, it is also essential that excess DA be quickly cleared from the synapse. If it is not, and the next response is an error, then the residual DA will strengthen inappropriate synapses – namely, those responsible for producing the incorrect response. This would undo the beneficial learning that occurred following correct responses, and prevent learning. For example, DA is quickly cleared from synapses by dopamine active transporter (DAT) in the striatum, and as a result, the temporal resolution of DA is high enough for DA to serve as an effective reinforcement training signal. Unlike the striatum, however, DAT concentration in cortex is low (e.g., Seamans & Robbins, 2010). As a result, cortical DA levels change very slowly. For example, the delivery of a single food pellet to a hungry rat can increase DA levels in prefrontal cortex for approximately 30 minutes (Feenstra & Botterblom, 1996). Thus, although DA is likely to facilitate long-term potentiation in cortex, it operates too slowly to serve as a cortical reinforcement training signal (Lapish et al., 2007). This is because the first rewarded behavior in a training session is likely to cause cortical DA levels to rise, and the absence of DAT will cause DA levels in cortex to remain high throughout the training session. As a result, all synapses that are activated during the session are likely to be strengthened, regardless of whether the associated behavior is appropriate or not. Instead, it is thought that cortical long-term potentiation follows Hebbian learning algorithms – that is, all (quasi)-simultaneously active synapses are strengthened, regardless of whether or not the resulting behavior was rewarded (Doya, 2000).

In computational models, Hebbian learning is generally homosynaptic and described using the product of pre- and post-synaptic activity (Ashby & Helie, 2011; Proulx & Helie, 2005). One way to insert feedback into a Hebbian learning algorithm without using DA is to

make it heterosynaptic by including a third population of neurons that increases post-synaptic activity following a correct response and decreases post-synaptic activity after an incorrect response. This form of heterosynaptic plasticity will increase or decrease (respectively) the amount by which the synaptic strength is adjusted. This is consistent with the observation that most observed forms of synaptic plasticity at GABA synapses are heterosynaptic (Castillo, Chiu, & Carroll, 2011). Populations of neurons responding to positive or negative feedback have been observed in orbitomedial PFC (O'Doherty et al., 2001), and orbitomedial PFC is highly connected to lateral PFC (Fuster, 2008). In addition, Stuber and colleagues (2010) have observed that at least some DA-producing neurons co-release glutamate. The glutamate could be used to modulate heterosynaptic plasticity. This can be accomplished indirectly by driving the activity in the feedback neurons already present in orbitomedial PFC, or by directly contributing to heterosynaptic plasticity in lateral PFC. HICL is consistent with all three of these possibilities, so we adopt the former for the present simulations (but this position could be revised as more is learned). Adding the orbitomedial PFC → lateral PFC connections to Figure 1a completes the HICL model. The resulting diagram is shown in Figure 2. As can be seen, the feedback units increase (or decrease through GABAergic interneurons) the input to the premotor units using a global signal – that is, the activation sent to all the premotor units is either increased or decreased by the same amount depending on the feedback. Hence, the input to premotor unit A becomes:

$$I_{motor,A}(t) = \sum_i \left[w^{stim \to motor} \times f\left[V_{stim,i}(t)\right] - w^{rule \to motor(i,A)} \times \left(f\left[V_{rule}(t)\right] + w^{motor \to rule} \times g\left[V_{motor,A}(t)\right]\right)\right]^+$$
$$- \sum_{j \neq A} w^{motor \to motor} \times f\left[V_{motor,j}(t)\right] + w^{pos \to motor} \times f\left[V_{pos}(t)\right] - w^{neg \to motor} \times f\left[V_{neg}(t)\right]$$

(4)

where $f[V_{pos}(t)]$ and $f[V_{neg}(t)]$ are the output of the positive and negative feedback units (respectively), $w^{pos \to motor}$ and $w^{neg \to motor}$ are the connection weights, and the rest of the symbols are as described in Eq. 3 above. Note that Eq. 4 is identical to Eq. 3 except for the addition of the feedback unit terms at the end. As such, motor input is simply increased or decreased by a global constant following feedback.

INSERT FIGURE 2 ABOUT HERE

5.2   Learning at inhibitory (GABA) synapses

Changes in synaptic efficiency of GABA synapses can be either pre- or post-synaptic, but the most prevalent types of plasticity are all heterosynaptic (i.e., they involve a third population of neurons). In the proposed learning model, the synapses to be modified are inhibitory connections between the rule unit and the axons of the categorical stimulus units (for pre-synaptic inhibition). The third population of neurons making this a heterosynaptic learning algorithm is the newly added population of feedback neurons located in orbitomedial PFC (as in Figure 2). Importantly, the conditions for long-term potentiation and long-term depression at inhibitory synapses are the opposite of those for excitatory synapses. Specifically, strong pre- and post-synaptic activity leads to long-term depression while weak pre- and post-synaptic activation leads to long-term potentiation (Castillo et al., 2011). Mathematically, the new proposed learning algorithm can be described by:

$$w^{rule \to motor}(n+1) = w^{rule \to motor}(n) - \eta_{LTD} \int f\left[V_{rule}(n,t)\right]dt \times \left\{\int f\left[I_{motor}(n,t)\right]dt - \theta_1\right\}^+ \times w^{rule \to motor}(n)$$
$$+ \eta_{LTP} \int f\left[V_{rule}(n,t)\right]dt \times \left\{\theta_1 - \int f\left[I_{motor}(n,t)\right]dt\right\}^+ \times \left\{\int f\left[I_{motor}(n,t)\right]dt - \theta_2\right\}^+ \times \left[w_{max} - w^{rule \to motor}(n)\right]$$

(5)

where $w^{rule \rightarrow motor}(n)$ is a connection weight between the rule unit and a synapse between a categorical stimulus unit and premotor unit at trial n, w_{max} is the maximum allowable weight, η_{LTP} and η_{LTD} are learning rates for long-term potentiation and depression (respectively), and θ_1 and θ_2 are parameters controlling the opening and closing of ionic gates (e.g., NMDAR receptors) or brain-derived neurotrophic factors (with $\theta_1 > \theta_2$). All the other symbols are as described before. The first term of Eq. 5 corresponds to conditions necessary to produce long-term depression (i.e., post-synaptic activation above θ_1) while the second term corresponds to conditions necessary for long-term potentiation (i.e., post-synaptic activation between θ_1 and θ_2). If post-synaptic activation is below θ_2, the synapse is not modified. Because the feedback units directly affect post-synaptic activity, their effect is to drive the post-synaptic activation above (for correct feedback) or below (for incorrect feedback) θ_1. Hence, the feedback units contribute to determining if the connection weight will be strengthened or reduced after each response. Note that Eq. 5 provides a general formalization of synaptic plasticity at inhibitory synapses as described by Castillo and colleagues (2011). When more research becomes available about synaptic plasticity at inhibitory synapses, than the physical role of θ_1 and θ_2 may be specified with more precision.

5.3 A toy problem

To show the new learning algorithm's capacity to learn categorization criteria, we simulate a simple toy problem where in each trial the model is given a numerical value between 1 and 100 and needs to classify all values below 50 as "A" and the remaining as "B" (so the ideal criterion is $x = 50.5$). The model architecture was as shown in Figure 2. The model included 10 categorical stimulus units each defined by a Gaussian distribution of mean $\mu = [5...95]$ by increment of 10 and a common standard deviation $\sigma = 10$ (see Figure 5a). The model also included one rule unit, one correct feedback unit, one incorrect feedback unit, and two response units (one for each category). All the model units were as defined in Eq. 1, and the parameter values related to the simulation are listed in Table 1. The model was simulated for 1,500 trials each having a duration of 2,800 msecs. The input to the rule unit was a fixed square function for the duration of the trial ($I_{rule} = 2,000$), the input to the stimulus units was $I_{stim} = 0$ for the first 500 msecs of each trial and $I_{stim} = 125,000$ afterward, and the input to the feedback unit was $I_{feedback} = 5,000$ after a response had been produced by the model and $I_{feedback} = 0$ before a response was made. Note that only one of the feedback units received input in each trial, corresponding to the accuracy feedback of the given trial.

Table 1. Parameter values for the simulations.

Parameter	Toy problem	Human experiments
General (Eqs. 1-2)		
σ_ε^2	2,000	0
$\lambda_{glutamate}$	60	60
λ_{GABA}	30	30
Weights (Eqs. 3-4)		
$w^{stim->motor}$	65	65
$w^{motor->motor}$	400	400
$w^{pos->motor}$	50	100
$w^{neg->motor}$	150	100
$w^{motor->rule}$	200	300
Learning (Eq. 5)		
η_{LTD}	1×10^{-14}	1×10^{-15}
η_{LTP}	4×10^{-24}	1×10^{-23}
θ_1	1×10^7	8×10^6
θ_2	1×10^5	1×10^6
$w^{rule->motor}(0)$	$5 + u(2)$	$30 + u(0.5)$
w_{max}	300	550

Note. All the parameters are as defined in the text. $u(x)$ is a uniform random number $[0...x]$.

Figures 3 and 4 show how the activity is propagated in an example trial. A trial went as follows: (1) The rule unit became activated by the square function (Figure 4, top-left). (2) 500 msecs later, a stimulus was randomly selected (e.g., 26.23) and presented to the model by activating all the stimulus units according to their receptive field (see Figure 3). (3) Activity was transferred to the motor planning units until one of the units reached the response threshold. In this example, motor planning unit 'A' reached the response threshold after about 750 ms (Figure 4, middle-left). (4) The unit that first reached the response threshold produced a response (in this case 'A'), and the appropriate feedback unit (in this case positive) received activation from the square function (Figure 4, bottom-left). (5) The trial ended after 2,800 msecs and the learning algorithm was applied.

INSERT FIGURE 3 ABOUT HERE
INSERT FIGURE 4 ABOUT HERE

The process described above was repeated 1,500 times. The results of the simulation are shown in Figure 5. As can be seen in Figure 5b, the model accuracy was initially at chance, but quickly increased to about 90% correct. This rapid improvement in accuracy was caused by the learning algorithm. Figure 5c shows the initial (random) rule weights in input space. The *x*-axis

shows the categorical stimulus units ordered from 1…10 (corresponding to the receptive fields in Figure 3) and the *y*-axis shows the weight magnitude. The grey line shows the weights affecting motor planning unit "A" synapses while the black line shows the weights affecting motor planning unit "B" synapses. The two lines cannot be distinguished in Figure 5c, but Figure 5d shows the same weights after 1,500 trials of training. As can be seen, the connections from the first 5 categorical stimulus units are highly inhibited by the "B" weights (black) but not much by the "A" weights (grey). As such, stimuli that activate maximally these five units will be prevented from producing a "B" response by the rule unit, and instead output "A". The reverse can be said of the last 5 categorical stimulus units: They are highly inhibited by the "A" weights but not much by the "B" weights. As such, stimuli maximally activating these units output a "B" response. Stimuli between 45 and 55 fall between categorical stimulus units 5 and 6 and will produce an "A" or "B" response depending on the randomly generated noise. Note that the response criterion is not represented explicitly in HICL. Instead, the transition from the "A" response to the "B" response happens where the two lines cross, roughly $x \sim 55$ in this simulation. The discrepancy between the optimal criterion and the learned criterion accounts for the ~10% error rate and could be reduced by decreasing the noise in the model. However, importantly, the criterion learned by the model is an epiphenomenon and does not have a biological meaning (although it may still have a psychological meaning as discussed in Section 7.2). In the following section, we present some new empirical data that were collected to test the model.

<div align="center">INSERT FIGURE 5 ABOUT HERE</div>

6   Criterion shifts in rule-based categorization

While the rule implementation of HICL can account for single-cell recordings (as shown in Helie & Ashby, 2009), and its learning algorithm can gradually learn a criterion (as shown in Section 5.3), HICL has yet to be assessed by comparing its performance to human data. In this section, data from two human behavioral experiments are used to test the new model. One way to test for criterion learning is to have participants learn a categorization criterion (e.g., large shapes vs. small shapes), and then at some point during the experiment change the categorization criterion so that participants need to learn a new criterion on either the same stimulus dimension (referred to as an intra-dimensional shift) or a different stimulus dimension (referred to as an extra-dimensional shift). Experiments where more than one criterion is learned (both simultaneously and in succession) are ideal to test HICL because they provide for more criterion learning data from each participant to test the model. Towards this end, we take as our starting point a 4-category perceptual classification task (thus requiring three rule criteria). During pre-shift training, each participant completed 300 trials. On each trial, a single line was presented, and the participant's task was to classify the line into one of four categories based on the line length. A scatterplot of the pre-shift stimuli and category structures is displayed in Figure 6A. The different symbols denote stimuli from different categories with each stimulus having a unique length and unique orientation. The use of a large number of unique stimuli reduces the possibility that participants memorize particular stimuli and puts the emphasis on rule learning as opposed to explicit memorization.

<div align="center">INSERT FIGURE 6 ABOUT HERE</div>

Because each participant in each experiment was asked to consecutively learn more than one criterion, the model needed to be enhanced to include a rule switching/selection mechanism. Rule selection/switching is not the focus of the current model, and the neurobiology of these processes is unclear at this point (Paul & Ashby, 2013). Hence, we made the simplifying assumption that rule switching/selection is controlled by a random walk model (Ratcliff, 1978) that accumulates "surprise" (i.e., how different the new stimulus is from previously experienced stimuli). Using a random walk model is a natural choice given the abundance of data suggesting that random walk models of decision making are neurally plausible (e.g., Heekeren, Marrett, Bandettini, & Ungerleider, 2004) and their widespread use in cognitive psychology (e.g., Ratcliff, 1979; Roe et al., 2001).

In a random-walk model, the magnitude of change following new evidence is called the 'drift rate'. Using surprise as the drift rate stems from the universal generalization gradient proposed by Shepard (1987). Specifically, Shepard observed that the probability of generalization from learned stimuli to new stimuli decreases exponentially with distance in psychological stimulus space. We hypothesize that this should also apply to rules in perceptual categorization, so that the estimated probability (or surprise) that a rule applies to a new stimulus is exponentially related to its distance from the previously experienced stimuli that are already covered by the rule (Tenenbaum & Griffith, 2001). For these reasons, surprise is used to modulate confidence in the current rule, and rule changes happen when a decision boundary is crossed by the confidence score. We speculate that these processes result from interactions between the caudate nucleus and the anterior cingulate cortex (Brown & Braver, 2005), but the random walk model should be considered a prosthetic part of HICL that should be replaced by a neurobiologically plausible alternative (implementing the same function of surprise accumulation) as more is learned about rule selection in the brain. The random-walk model used in the simulations is formalized in the Appendix.

With the current implementation, HICL makes four general predictions: (1) Participants can learn categorization criteria; (2) Participants will switch their categorization criteria more quickly if the new stimuli are different from the (original) training stimuli (i.e., "surprising"); (3) Changing the visual representation of the stimuli on the pre-shift rule irrelevant dimension only will not affect categorization performance if the dimension is still irrelevant after the shift; (4) However, changing the stimulus on the pre-shift rule irrelevant dimension will produce a rule shift if the dimension becomes relevant post-shift.

Prediction (1) is trivial and many experiments have already shown that participants can learn categorization criteria (e.g., Busemeyer & Myung, 1992; Kubovy & Healy, 1977). In HICL, criterion learning is achieved using the algorithm described in Section 5. Prediction (4) is less trivial but is consistent with previous work on knowledge restructuring that suggests that two conditions need to be met for participants to change their response strategy (Kalish, Lewandowsky, & Davies, 2005; Little, Lewandowsky, & Heit, 2006): they need to make errors, and an alternative strategy needs to be available. In HICL, the alternative strategy is to select a new stimulus dimension and learn a new criterion. The fact that a previously irrelevant dimension is now becoming relevant will generate the errors required for rule shifting.

Unlike Predictions (1) and (4), predictions (2) and (3) are new. Prediction (2) follows from the surprise score being used to control the drift rate of the random walk model (as formalized in Eq. A1). More different ("surprising") stimuli will thus have a large effect on the

confidence score and allow for faster rule shifting (as shown in Eq. A2). Prediction (3) follows from biological constraints suggesting that only the rule-relevant stimulus dimension is represented in PFC (Freedman et al., 2003), and therefore a change in the rule irrelevant dimension only will not affect the model performance. These latter two predictions are tested in the following experiments.

6.1 Experiment 1: Intra-Dimensional Shifts in Rule-Based Category Learning

Experiment 1 examines the effects of several forms of intra-dimensional shift on four-category unidimensional rule-based perceptual classification. In addition to testing whether participants can learn the categorization criteria, the goal of these manipulations is to vary the similarity between the training and test stimuli and see their effect on criterion shifting. All participants were trained on the Pre-Shift/No-Shift category structure (Figure 6A). Following training, participants were tested on one of the four category structures plotted in Figure 6. The No-Shift(ID) condition was a control condition in which the category structure did not change from pre-shift training to post-shift test. In the Length-Shift(ID), Orientation-Shift(ID), and Both-Shift(ID) conditions, however, the category structures shifted along the length and/or orientation dimensions. More specifically, during training (and test in the No-Shift(ID) condition), the optimal strategy was to set decision criteria along the line length dimension, to ignore line orientation, and to respond "A" to lines shorter than 80 pixels, "B" to lines between 80 and 100 pixels, "C" to lines between 100 and 120 pixels, and "D" to lines greater than 120 pixels. In the Length-Shift(ID) condition, the categories were shifted along the length dimension such that the optimal strategy was to respond "A" to lines shorter than 180 pixels, "B" to lines between 180 and 200 pixels, "C" to lines between 200 and 220 pixels, and "D" to lines greater than 220 pixels. In the Orientation-Shift(ID) condition, the decision rule on length was unchanged from training to test, but the lines were considerably steeper than those presented during training. In the Both-Shift(ID) condition, the decision rule changed as it did in the Length-Shift condition, but the orientation values also become steeper (as they did in the Orientation-Shift(ID) condition).

The model predicts no change in accuracy in the Orientation-Shift(ID) condition (because the categorical stimulus units in PFC are insensitive to changes along the irrelevant dimension). In addition, the Both-Shift(ID) condition should have a smaller switch cost (i.e., decrease in accuracy from training to test) than the Length-Shift(ID) condition because shifts in the decision criteria are more likely when the training and test stimuli are more dissimilar (because of increased surprise scores).

6.1.1 Methods

6.1.1.1 Participants

Ninety-six participants completed the study and received course credit for their participation at University of Texas Austin. All participants had normal or corrected to normal vision. Each participant served in only one condition. To ensure that only participants who performed well above chance were included in the post-shift performance analyses, a learning criterion of 40% correct (25% is chance) during the final pre-shift block was applied. All but 8

participants met the performance criterion (No Shift(ID): N = 22; Length Shift(ID): N = 23; Orientation Shift(ID): N = 23; Both Shift(ID): N = 20).

6.1.1.2   Stimuli and Stimulus Generation

The stimuli for each condition are displayed in Figure 6 and were generated by drawing 75 random samples from each of four bivariate normal distributions along the stimulus dimensions of length and orientation. In all conditions the variance along the length dimension was 25 pixels, the variance along the orientation dimension was 36 degrees and the covariance was zero. The training stimuli used in all conditions, and during test for the No-Shift(ID) condition, were sampled from category distributions with means along the length dimension of 70, 90, 110, and 130 and along the orientation dimension of 30 for categories A – D, respectively. The test stimuli used in the Length-Shift(ID) condition were sampled from category distributions with means along the length dimension of 170, 190, 210, and 230 and along the orientation dimension of 30 for categories A – D, respectively. The test stimuli used in the Orientation-Shift(ID) condition were sampled from category distributions with means along the length dimension of 70, 90, 110, and 130 and along the orientation dimension of 60 for categories A – D, respectively. The test stimuli used in the Both-Shift(ID) condition were sampled from category distributions with means along the length dimension of 170, 190, 210, and 230 and along the orientation dimension of 60 for categories A – D, respectively. Optimal accuracy was 95%. The training and test stimuli were randomized separately for each participant.

6.1.1.3   Procedure

Participants were randomly assigned to one of the four conditions: No Shift(ID), Length Shift(ID), Orientation Shift(ID), or Both Shift(ID). Each condition consisted of three 100-trial pre-shift training blocks using the stimuli in Figure 6A, followed by three 100-trial post-shift test blocks with a participant controlled rest period between each block.

At the start of the experiment, participants were told that they were to categorize lines on the basis of their length and orientation, that there were four equally-likely categories, and that high levels of accuracy could be achieved. At the start of each trial, a fixation point was displayed for 1 second and then the stimulus appeared. The stimulus remained on the screen until the participant generated a response by pressing the "Z" key for category A, the "W" key for category B, the "/" key for category C, or the "P" key for category D. If any button other than one of these four was pressed, an "invalid key" message was displayed. Following the response, the word "correct" was presented if their response was correct or the word "incorrect" was presented if their response was incorrect, along with the correct category label. Once feedback was given, the next trial was initiated. Participants were given no instructions regarding the post-shift manipulation.

6.1.2   Results

The learning curves for all four conditions across the three pre- and three post-shift blocks are shown in Figure 7A. To verify that there were no pre-shift performance differences across conditions, we examined pre-shift accuracy across the four conditions. A 4 Condition × 3 Pre-shift block ANOVA was conducted and revealed a main effect of block [$F(2, 168) = 39.31$, $p < .001$, $\eta^2 = .319$], but no effect of condition [$F(3, 84) = .611$, ns, $\eta^2 = .021$] and no interaction

[F(6, 168) = .82, ns, η^2 = .028]. Most importantly, there were no differences across conditions in the final pre-shift block [F(3, 84) = .670, ns, η^2 = .023]. Thus, pre-shift training performance was equated across conditions.

INSERT FIGURE 7 ABOUT HERE

Our main performance measure is the <u>cost</u> associated with the introduction of the post-shift condition. We quantified the cost by subtracting accuracy in the first post-shift block from accuracy in the final pre-shift block. The larger the cost the greater the immediate impact of the post-shift condition on performance. The cost was significantly larger than zero in the Length-Shift(ID) [t(23) = 11.91, p < .001] and Both-Shift(ID) [t(21) = 11.89, p < .001] conditions, and did not differ from zero in the No-Shift(ID) [t(21) = .52, ns] and Orientation-Shift(ID) conditions [t(22) = .51, ns]. A 2 Length (shift vs. no shift) × 2 Orientation (shift vs. no shift) ANOVA was conducted on the performance costs. This revealed a significant effect of Length shift [F(1, 84) = 205.92, p < .001, η^2 = .710] and a significant interaction [F(1, 84) = 4,44, p < .05, η^2 = .050], but a non-significant effect of Orientation shift [F(1, 84) = 1.73, ns, η^2 = .020]. As expected, a shift along the length dimension led to a significant increase in the shift cost (.27) relative to no shift along the length dimension (shift cost = 0.00). Shifts along the irrelevant dimension reduced the size of the performance cost when the relevant dimension also shifted [Both Shift(ID) vs. Length Shift(ID)]. Specifically, the shift cost was significantly smaller in the Both-Shift(ID) condition (.23) than in the Length-Shift(ID) condition (.30) (p < .05).

The Both-Shift(ID) condition is analogous to ID shift conditions in the visual discrimination literature and ID shifts typically yield smaller costs than ED shifts in visual discrimination (Dias et al., 1996; Owen et al., 1993). The fact that the Length-Shift(ID) condition yielded a larger cost suggests that ID shift conditions might exist that yield comparable costs to traditional ED shift conditions. We address the issue of ED shift directly in Experiment 2.

6.1.3   Brief Summary

The predictions made by the model were supported by the results of Experiment 1. First, changing the stimuli only on the rule irrelevant dimension did not affect the participants' performance. This is because the rule irrelevant dimension is not represented in HICL, and the categorization criteria remain unchanged. Second, changing the stimuli on the rule relevant dimension did affect performance, but participants were able to learn the new criteria in both conditions. Third, changing only the rule relevant dimension yielded a higher cost, because the post-shift stimuli were still relatively similar to the pre-shift stimuli (hence a low surprise score). However, changing the stimuli on both the relevant and irrelevant dimensions resulted in brand new stimuli that had nothing in common with the pre-shift stimuli (and a high surprise score). As a result, fourth, the cost was much smaller in the Both-Shift(ID) condition. Numerical simulation of Experiment 1 is presented in section 6.3, but first we explore the effect of ED shift in the 4-category categorization experiment.

6.2   Experiment 2: Extra-Dimensional Shifts in Rule-Based Category Learning

Experiment 2 examines the effects of several forms of ED shifts on four-category unidimensional rule-based perceptual classification. This test is essential to ensure that the new model can not only learn new criteria on a selected dimension (ID shift), but also shift to a different dimension and learn new criteria (ED shift). The pre-shift training items are identical to

those from Experiment 1, and the shifts are to the same locations in the length-orientation stimulus space. All participants were trained on the Pre-Shift category structures (Figure 8A). Following training, participants were tested on one of the five category structures plotted in Figure 8. A Control condition was included in which the category structures did not change from pre-shift training to post-shift test (Figure 8A). In the No-Shift(ED), Length-Shift(ED), Orientation-Shift(ED), and Both-Shift(ED) conditions (Figures 7B – 7E), however, there was an ED shift. With the exception of the No-Shift(ED) condition, the category structures shifted along the length and/or orientation dimensions. More specifically, in the No-Shift(ED) condition (Figure 8B), participants were tested on category structures with stimuli sampled from the same part of the length-orientation space as sampled during pre-shift training, but with categories that were rotated in such a way that the optimal strategy was to attend to orientation while ignoring length and to respond "A" to lines shallower than 24 degrees, "B" to lines between 24 and 30 degrees, "C" to lines between 30 and 36 degrees, and "D" to lines greater than 36 degrees. In the Length-Shift(ED) condition the lengths shifted toward longer lines, but the values along the orientation dimension and thus the post-shift decision rule were identical to those in the No-Shift(ED) condition. In the Orientation-Shift(ED) condition the range of length values remained unchanged from those used during training, but the orientation values shifted toward steeper angled lines, and thus the decision rule shifted such that the optimal strategy was to respond "A" to lines shallower than 54 degrees, "B" to lines between 54 and 60 degrees, "C" to lines between 60 and 66 degrees, and "D" to lines greater than 66 degrees. In the Both-Shift(ED) condition the length and orientation values shifted and thus the decision rule shifted in the same way that it shifted in the Orientation-Shift(ED) condition.

INSERT FIGURE 8 ABOUT HERE

Here, we predict that the No-Shift(ED) and the Orientation-Shift(ED) conditions will be most difficult. This is because some of the test stimuli will be in the same category as the training stimuli (e.g., there is a small range of stimuli that would be categorized as 'A' during training and test), which produces fewer errors and maintains confidence in the pre-shift rule. In addition, No Shift(ED) will be worst because the surprise score on error trials will be smallest. In contrast, the Length-Shift(ED) and Both-Shift(ED) conditions will be easier, because the test stimuli will all change categories, thus producing more errors and reducing confidence in the pre-shift rule. In addition, Both-Shift(ED) will be easiest, because the stimuli are more perceptually different and yield a larger surprise score (similar to Experiment 1).

6.2.1   Methods

6.2.1.1   Participants

One hundred and twenty-four participants from University of Texas Austin completed the study and received course credit for their participation. All participants had normal or corrected to normal vision. Each participant served in only one condition. To ensure that only participants who performed well above chance were included in the post-shift performance analyses, a learning criterion of 40% correct (25% is chance) during the final pre-shift block was applied. All but 5 participants met the performance criterion (Control: N = 24; No Shift(ED): N = 22; Length Shift(ED): N = 24; Orientation Shift(ED): N = 25; Both Shift: N = 24).

6.2.1.2 Stimuli and Stimulus Generation

The stimuli for each condition are displayed in Figure 8. The training stimuli and the test stimuli in the Control condition were identical to those from Experiment 1. The stimuli used in the No-Shift(ED) condition were generated by drawing 75 random samples from each of four bivariate normal distributions along the stimulus dimensions of length and orientation. The variance along the orientation dimension was 2.25 degrees, the variance along the length dimension was 400 pixels and the covariance was zero. The category distribution mean along the length dimension was 100 pixels for all categories. The category distribution means along the orientation dimension were 21, 27, 33, and 39 degrees for categories A – D, respectively. In the Length-Shift(ED) condition the orientation means were identical to those in the No-Shift(ED) condition, but the mean along the length dimension was 200 pixels. In the Orientation-Shift(ED) condition the orientation means were shifted toward steeper angles (i.e., 51, 57, 63, and 69 for categories A – D, respectively), whereas the length mean remained at 100 pixels. Finally, in the Both-Shift(ED) condition, the length and orientation means were shifted.

6.2.1.3 Procedure

The procedure was identical to that from Experiment 1.

6.2.2 Results

The learning curves for all five conditions across the three pre- and three post-shift blocks are shown in Figure 9A. A 5 Condition $\times$ 3 Pre-shift block ANOVA was conducted to determine whether there were any pre-shift performance differences. The block effect was significant F(2, 228) = 56.23, p < .001, η^2 = .330], but the condition effect [F(4, 114) = .61, ns, η^2 = .021] and the interaction [F(8, 228) = 1.16, ns, η^2 = .039] were both non-significant. Most importantly, there were no differences across conditions in the final pre-shift block [F(4, 114) = 1.38, ns, η^2 = .046]. Thus, pre-shift training performance was equated across conditions.

<div align="center">INSERT FIGURE 9 ABOUT HERE</div>

The Control condition was included to verify that performance had reached asymptote at the end of pre-shift training. An ANOVA on accuracy in the final pre-shift block and in the three post-shift blocks was non-significant [F(3, 69) = .43, ns, η^2 = .018], suggesting no additional learning in the Control condition. These data are not included in any subsequent analyses.[4] As expected, the cost was significantly larger than zero in all conditions [No Shift(ED): t(21) = 13.07, p < .001; Length Shift(ED): t(23) = 11.17, p < .001; Orientation Shift(ED): t(24) = 14.39, p < .001; Both Shift(ED): t(23) = 11.54, p < .001]. As in Experiment 1, we begin with a 2 Length (shift vs. no shift) $\times$ 2 Orientation (shift vs. no shift) ANOVA on the performance costs. The main effect of length shift was significant [F(1, 91) = 22.30, p < .001, η^2 = .197] and suggested that a larger cost was associated with <u>no</u> shift along the previously relevant length dimension (shift cost = .40) than was associated with a length shift (shift cost = .28). The main effect of orientation [F(1, 91) = 1.21, ns, η^2 = .013], and the interaction were non-significant [F(1, 91) =

---

[4] Post-shift learning was also examined and revealed non-significant main effects of length and orientation shifts, and a non-significant interaction.

0.00, ns, η^2 = .000]. Although the interaction was non-significant, the No-Shift(ED) cost (shift cost = .42) was significantly larger than the Both-Shift(ED) cost (shift cost = .26) (p < .001), as suggested by Figure 9B.

6.2.3   Brief Summary

The most important finding from this study was that, as predicted, the cost was smaller when there was no shift on the length dimension then when the stimuli shifted along the length dimension, regardless of whether there was a shift along the orientation dimension. This is because all the stimuli changed categories when shifting was on the pre-shift rule relevant dimension, which yielded more post-shift errors (and a faster change in rule). In contrast, not all stimuli changed categories when there was no shift along the length dimension, which slowed down rule change and yielded a higher cost. Overall, No Shift(ED) was hardest because in addition to producing fewer errors, the test stimuli were similar to the training stimuli (thus surprise scores were smaller). In contrast, Both Shift(ED) was easiest, because more post-shift errors were produced, and each error had a higher surprise score. This difference was statistically meaningful. The next subsection simulates the experiments and numerically estimates the magnitude of the predicted effects.

6.3   Simulating criterion learning with the new model

The previous two experiments showed that (1) human participants can learn categorization criteria, and (2) that having new stimuli that differ on the rule-relevant dimension increases the likelihood of a rule switch. These two results held for both ID (Experiment 1) and ED (Experiment 2) shifts. These results are consistent with the HICL model because rule criteria are learned using feedback-driven Hebbian learning, and only the rule-relevant dimension is considered when learning a criterion. New rules are selected based on the number of errors produced (confidence) and the similarity between the new stimuli and the old stimuli (surprise). This subsection proceeds with simulating the previous experiments.

In the model, selecting a rule is equivalent to selecting a stimulus dimension on which to learn a criterion (Ashby et al., 1998). For example, with the line stimuli described above, a participant might select line length. If line length is the selected rule, than only this measure is represented by the categorical stimulus units in the PFC (as described in Section 3). In contrast, if the selected rule is line orientation, then only the rotation angle is represented by the categorical stimulus units in the PFC. Hence, each stimulus has a different categorical stimulus representation in PFC depending on the selected rule. The rule is stochastically selected according to its salience, and the distribution of salience is initially uniform. So, on trial one, the simulated participant randomly picks one of the dimensions and tries to learn a criterion. The model is simulated for the same number of trials as the humans (600) using the same stimulus distributions (see Figures 6 and 8). Each trial is simulated for 2,800 msecs. Each condition in each experiment was simulated 100 times.

The simulations include 10 categorical stimulus units each defined by a Gaussian distribution. The Gaussian distributions were designed to give the model the same acuity on both dimensions (i.e., similar distribution overlap). For the line length, the (univariate) Gaussian distributions had a mean μ = [60...240] with an increment of 20 and a common standard deviation σ = 8. For the line orientation, the (univariate) Gaussian distributions had a mean μ =

[18...72] with an increment of 6 and a common standard deviation $\sigma = 2.4$. The simulations also included two rule units (one for each stimulus dimension), one correct feedback unit, one incorrect feedback unit, and four response units (one for each category). All the unit models were as defined in Section 4, and the parameter values related to the simulation are listed in Table 1. The input to the rule unit was a fixed square function for the duration of the trial ($I_{rule} = 5,000$), the input to the stimulus units was $I_{stim} = 0$ for the first 500 msecs of each trial and $I_{stim} = 125,000$ afterward, and the input to the feedback cell was $I_{feedback} = 5,000$ after a response had been produced by the model and $I_{feedback} = 0$ before a response was made (these value are the same as in the toy problem). Note that only one of the feedback units received input in each trial, corresponding to the accuracy feedback of the given trial.

A trial went as follows: (1) The selected rule unit becomes activated by the square function. (2) 500 msecs later, a stimulus is randomly selected and presented to the model by activating all the stimulus units according to their receptive field. (3) Activity gets transferred to the motor planning units until one of the units reaches the response threshold. (4) The unit that first reaches the response threshold produces a response, and the appropriate feedback unit (positive or negative) receives activation from the square function. (5) The trial ends after 2,800 msecs and the learning algorithm is applied. (6) The rule switching/selection mechanism (random walk) is updated (i.e., surprise, confidence, rule salience). The simulation results for both experiments are shown in Figure 10. As can be seen, the model was able to learn the rule criteria in both experiments. The qualitative pattern of results was reproduced with a common set of parameters, and the quantitative results were also well reproduced. The r^2 for blocked-data in Experiments 1 and 2 was $r^2 = 0.883$ and $r^2 = 0.940$ (respectively).

INSERT FIGURE 10 ABOUT HERE

In Experiment 1, the criteria for the Both-Shift(ID) condition were learned more quickly following the shift than the criteria for the Length-Shift(ID) condition, and the Orientation Shift(ID) had no effect on performance. The first result is explained by the surprise score accumulating more slowly in the Length-Shift(ID) condition (because the new stimuli are more similar to the old stimuli than in the Both-Shift(ID) condition), thus reducing the drift rate in the random walk and increasing the number of trials required to shift the rule. The later result is accounted for by the absence of a representation of the rule-irrelevant dimension in the categorical stimulus units in the PFC, so the model does not notice the orientation shift and no errors are made.

In Experiment 2, the criteria for the Length-Shift(ED) and the Both-Shift(ED) conditions were learned more quickly following the shift than the criteria for the Orientation-Shift(ED) and the No-Shift(ED) conditions. These results are accounted for by the number of post-shift errors, which was smaller in the No-Shift(ED) and Orientation-Shift(ED) conditions because some of the stimuli remained in the same category after the shift. A smaller number of errors slowed down the decrease in the confidence of the pre-shift rule. In addition, the drift rate for rule change was affected by the similarity of the post-shift stimuli to the pre-shift stimuli: Similar pre- and post-shift stimuli produced a small amount of surprise, which reduced the drift rate in the random walk model and increased the number of errors required before the rule was changed. These results are consistent with the results obtained in the human experiments (see Figures 7 and 9) and support the psychological relevance of the proposed neurobiological model of criterion learning. Note that none of the parameters were optimized, and that a single set of

parameter values was used for all the conditions in both experiments (so the data had 54 degrees of freedom). Hence, these results were produced by the proposed architecture, and were not the result of excessive parameter tuning.

7   General discussion

This article introduces a new computational cognitive neuroscience model of criterion learning of rule-guided behavior called HICL. HICL is the first model to provide a biological circuit for criterion learning that can be simulated using spiking neurons (Izhikevich, 2007). HICL represents rules using pre-synaptic inhibition (Shepherd, 2004) and proposes a new feedback-driven Hebbian learning algorithm for inhibitory (GABA) synapses. The new learning algorithm is heterosynaptic and relies on a population of feedback neurons located in orbitomedial PFC to open and close ionic channels in the pre- and post-synaptic neurons. The rule implementation model had already been used to account for electrophysiology data collected in monkeys (Helie & Ashby, 2009). In this article, new data were collected with humans showing that learning new criteria on a given dimension (ID shift) was easier if the irrelevant dimension also changed, making the stimuli more perceptually different (Experiment 1). Another experiment showed that changing the relevant rule dimension and learning new criteria (ED shift) is easier if the stimuli are shifted on the rule-relevant dimension (Experiment 2). Taken together, this integration of neurocomputational and behavioral work helps to elucidate an important component of rule-guided behavior.

The proposed model implements three key ideas: (1) learning rule criteria is feedback-driven, (2) rule changes mostly happen when errors are made, and (3) fewer errors are needed for criterion learning following a rule change if the new stimuli are perceptually different (or surprising). Experiments 1 and 2 were designed specifically to test these key ideas. First, participants had to learn several categorization criteria using feedback alone. Second, some conditions were designed to produce more errors at test than others (e.g., when there was a change on the training rule dimension). Third, some conditions were designed so that the test stimuli were more perceptually different from the training stimuli (e.g., when the test stimuli differed from the training stimuli on both dimensions). The model was able to account for all three effects in the data, and did so in a biologically-realistic way. The model was also able to account for additional behavioral and biological constraints, such as flexible stimulus-response mapping (Ashby et al., 2003), neuron firing rates and timing (Helie & Ashby, 2009), the slow time course of cortical DA (Lapish et al., 2007), and synaptic plasticity at inhibitory synapses (Castillo et al., 2011).

7.1   Neurocomputational modeling of rule-guided behavior

Research on the neural substrates of rule-guided behavior has typically focused on the PFC (Badre, Keyser, & D'Esposito, 2010; Bunge & Wallis, 2007; Christoff, Keramatian, Gordon, Smith, & Mädler, 2009). Very few attempts have been made, however, to develop neurocomputational models that specify how rule-guided behavior operates. Those models that have been proposed generally focus on how the PFC and basal ganglia interact to support rule learning (e.g., Ashby et al., 1998) and/or rule generalization (e.g., Collins & Frank, 2013). For example, the COVIS model of category learning proposed by Ashby and colleagues includes an explicit, hypothesis-testing system that assumes that a PFC-anterior cingulate network mediates

the selection of new rules whereas the basal ganglia mediates disengagement from previously attended rules. Importantly, however, rule selection in COVIS (and in the present work) is synonymous with selecting a stimulus dimension. COVIS did not provide a neural account of how criterion learning is achieved once a rule is selected. Moreover, the vast majority of the studies investigating the neural substrates of rule-guided behavior have used tasks with minimal criterion learning demands (e.g., binary-valued stimulus dimensions). Thus, perhaps it is not surprising that detailed neurocomputational accounts of criterion learning have been lacking.

7.2   Implications for criterion learning research

Cognitive psychologists have long been interested in the mechanics of criterion learning. Generally speaking, cognitive models of criterion learning have focused on how the criteria are updated in response to trial-by-trial information (Busemeyer & Myung, 1992; Erev, 1998; Kac, 1962; Kubovy & Healy, 1977; Maddox, 2002; Treisman & Williams, 1984). However, one important implication of the current model is that decision criteria are not explicitly represented: At the biological level, a criterion is an epiphenomenon that emerges from the gradual learning of rule weights responsible for the pre-synaptic inhibition of stimulus-response associations. Specifically, what the model is learning is what response not to make. This is interesting because previous accounts of the role of the PFC have proposed that a major role of the PFC is to bias behavior-related activation using inhibition (Miller & Cohen, 2001). However, the epiphenomenal nature of criteria at the biological level does not imply that the notion of a criterion does not have an important psychological meaning. For example, Maddox (2002) proposed a theory where participants learn to adjust their criterion in perceptual categorization tasks using two mechanisms: (1) competition between reward and accuracy maximization and (2) a flat maxima hypothesis (i.e., criterion placement is based on the steepness of the reward function). These psychological mechanisms can account for the participants' propensity to adjust their categorization criteria (although suboptimally) to category payoffs and base rates. It remains to be seen whether these mechanisms at the psychological level reflect a shift of decision weights or a change in the perceptual representation of the stimuli (Folstein et al., 2013). However, the proposed model may be useful in disentangling these possibilities.

Another interesting implication of the proposed model is that rule-based stimulus-response assignment and criterion learning are one and the same at the biological level. If the stimulus dimensions are discrete, the categorical stimulus representation in the PFC will also be discrete and the same learning algorithm will be used to associate specific dimensional values to responses, without generalizing to other dimensional values. This was the case in the simulation of the "same"-"different" task included in Helie and Ashby (2009), where monkeys saw objects that differed discretely. Future research should focus on testing this prediction and explore how criterion shift is related to changes in stimulus-response assignments (e.g., switch the location of response button after initial training, Ashby et al., 2003).

7.3   Extradimensional and intradimensional set shifting

ED and ID set shifting tasks have been used extensively in the neuropsychological literature to characterize cognitive impairment resulting from numerous neurological disorders (e.g., Parkinson's disease, frontal lobe pathology – Owen et al., 1993). Typically, ED and ID set shifting are studied using visual discrimination tasks in which highly discriminable stimuli (e.g.,

colors, shapes) are associated with different reward outcomes. Initially, the participant's task is to learn the stimulus-reward associations. The aspect of performance that is the central focus, however, is how the participant responds when the task set changes. In an ID shift, the participants need to shift to a new decision criterion on the currently attended dimension. In an ED shift, the participants need to shift to a new decision criterion on a different stimulus dimension (e.g., a previously irrelevant stimulus dimension).

The classic ED and ID set shifting tasks were designed, in part, as a more precise way to investigate the impairments observed on more complex set shifting tasks (e.g., the Wisconsin Card Sorting Task – Grant & Berg, 1948). The use of highly discriminable stimuli maintained continuity between these different set-shifting tasks. As a consequence, however, these set shifting-tasks place minimal demands on criterion learning, potentially obscuring criterion learning deficits. Thus, in addition to extending the ED/ID set shifting literature, the present tasks could provide a valuable tool for investigating criterion learning in neuropsychological populations. This may be particularly important for individuals with Parkinson's disease as these individuals have a consistent impairment in rule-guided behavior that appears to be driven, in part, by a criterion learning deficit (Ell, 2013; Ell, Weinstein, & Ivry, 2010; Filoteo, Maddox, Ing, & Song, 2007). Furthermore, the criterion learning impairment in Parkinson's patients appears to be insensitive to dopaminergic medication status (Ell et al., 2010). Likewise, criterion learning in the new model does not rely on DA for a training signal. Thus, the model proposed in this paper suggests a possible neural mechanism explaining previous findings in patients with Parkinson's disease.

7.4   Limitations and Future Directions

The HICL model focused on how rule criteria are learned. However, before learning a criterion, one must choose a rule. Rule selection and switching was modeled using a simple random walk model with feedback-driven rule saliency learning, but no biological explanation was provided. The rule criterion model of Helie and Ashby (2009) (of which HICL is an extension) assumes that rules are maintained in working memory. If this is the case, then HICL could be connected to an existing biological model of working memory (e.g., FROST: Ashby, Ell, Valentin, & Casale, 2005) and use the working memory model mechanism for cycling through the working memory items for rule switching and selection. Future work should be devoted to such an integration, and especially exploring the trade-off between allowing for sufficient errors to allow for criterion learning, while also allowing for rule switching if the criterion learning is not fruitful (e.g., the exploration/exploitation dilemma).

Relatedly, the PFC is thought to represent rules that are currently in working memory. However, one needs to be able to store rules in long-term memory, and retrieve them when needed. In the current model, a rule cell corresponds to a selected stimulus dimension (to allow for flexible categorical stimulus representation in the PFC) and a set of weights. The current version of the model did not attempt to provide a computational or biological explanation for rule storage and retrieval. Such long-term memory processes probably involve the medial temporal lobes, which are highly connected to the PFC (Eichenbaum, 2004). Rule storage and retrieval is relevant in criterion learning because none of the conditions in either experiment involved the unlearning of the rule criteria learned during the pre-shift period. All the criteria learned during the post-shift period were the results of rule switching. Hence, one prediction of

the model is that participants should be able to revert back to using their pre-shift criteria without having to relearn the criteria. Accordingly, future work should be devoted to exploring how the rule cells can be stored and retrieved using a neuronal circuit of spiking neurons.

8   Acknowledgments

9   References

Asaad, W. F., Rainer, G., & Miller, E. K. (1998). Neural activity in the primate prefrontal cortex during associative learning. *Neuron*, *21*, 1399-1407.

Ashby, F. G., Alfonso-Reese, L. A., Turken, A. U., & Waldron, E. M. (1998). A neuropsychological theory of multiple systems in category learning. *Psychological Review*, *105*, 442-481.

Ashby, F. G., & Crossley, M. J. (2011). A computational model of how cholinergic interneurons protect striatal-dependent learning. *Journal of Cognitive Neuroscience*, *23*, 1549–1566.

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*, 1728–1743.

Ashby, F. G., Ell, S. W., & Waldron, E. M. (2003). Procedural learning in perceptual categorization. *Memory & Cognition*, *31*, 1114-1125.

Ashby, F. G., Ennis, J. M., & Spiering, B. J. (2007). A neurobiological theory of automaticity in perceptual categorization. *Psychological Review*, *114*, 632–656

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

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

Bowers, J. S. (2009). On the biological plausibility of grandmother cells: Implications for neural network theories in psychology and neuroscience. *Psychological Review*, *116*, 220-251.

Brown, J. W., & Braver, T. S. (2005). Learned predictions of error likelihood in the anterior cingulate cortex. *Science, 307,* 1118-1121.

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

Bunge, S. A., & Souza, M. J. (2007). Neural representations used to specify action. In S. A. Bunge & J. D. Wallis (Eds.) *Neuroscience of Rule-Guided Behavior* (pp. 45-65). New York: Oxford University Press.

Bunge, S. A., & Wallis, J. D. (2007). *Neuroscience of rule-guided behavior.* New York: Oxford University Press.

Busemeyer, J. R., & Myung, I. J. (1992). An adaptive approach to human decision making: learning theory, decision theory, and human performance. *Journal of Experimental Psychology: General, 121*, 177-194.

Castillo, P. E., Chiu, C. Q., & Carroll, R. C. (2011). Long-term plasticity at inhibitory synapses. *Current Opinion in Neurobiology*, *21*, 328-338.

Christoff, K., Keramatian, K. Gordon, A. M., Smith, R., & Mädler, B. (2009). Prefrontal organization of cognitive control according to levels of abstraction. *Brain Research, 1286,* 94-105.

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

Dias, R., Robbins, T. W. & Roberts, A. C. (1996). Dissociation in prefrontal cortex of affective and attentional shifts. *Nature*, *380*, 69-72.

Doya K. (2000). Complementary roles of basal ganglia and cerebellum in learning and motor control. *Current Opinion in Neurobiology*, *10*, 1–19.

Eichenbaum, H. (2004). Hippocampus: Cognitive processes and neural representations that underlie declarative memory. *Neuron*, *44*, 109–120.

Ell, S.W..(2013). Targeted training of the decision rule benefits rule-guided behavior in Parkinson's disease. *Cognitive, Affective, and Behavioral Neuroscience, 13,* 830-846.

Ell, S. W., Weinstein, A., & Ivry, R. B. (2010). Rule-based categorization deficits in focal basal ganglia lesion and Parkinson's disease patients. *Neuropsychologia, 48*, 2974-2986.

Erev, I. (1998). Signal detection by human observers: A cutoff reinforcement learning model of categorization decisions under uncertainty. *Psychological Review, 105*, 280-298.

Erickson, M. A., & Kruschke, J. K. (1998). Rules and exemplars in category learning. *Journal of Experimental Psychology: General, 127*, 107-140.

Feenstra, M. G., & Botterblom, M. H. (1996). Rapid sampling of extracellular dopamine in the rat prefrontal cortex during food consumption, handling and exposure to novelty. *Brain Research*, *742*, 17-24.

Filoteo, J. V., Maddox, W. T., Ing, A. D., & Song, D. D. (2007). Characterizing rule-based category learning deficits in patients with Parkinson's disease. *Neuropsychologia, 45*, 305-320.

Folstein, J. R., Palmeri, T. J., & Gauthier, I. (2013). Category learning increases discriminability of relevant object dimensions in visual cortex. *Cerebral Cortex*, *23*, 814-823.

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*, 5235-5246.

Fuster, J. (2008). *The Prefrontal Cortex*. Academic Press.

Glimcher, P. W. (2011). Understanding dopamine and reinforcement learning: the dopamine reward prediction error hypothesis. *Proceedings of the National Academy of Sciences*, *108*, 15647–15654.

Grant, D. A., & Berg, E. A. (1948). 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*, 404-411.

Heekeren, H.Ri., Marrett, S., Bandettini, P.A., Ungerleider, L.G. (2004). A general mechanism for perceptual decision-making in the human brain. *Nature, 431,* 859-62.

Helie, S. & Ashby, G. F. (2009). A neurocomputational model of automaticity and maintenance of abstract rules. *Proceedings of the International Joint Conference on Neural Networks* (pp. 1192-1198). Atlanta, GA: IEEE Press.

Helie, S. & Ashby, F. G. (2012). Learning and transfer of category knowledge in an indirect categorization task. *Psychological Research, 76*, 292-303.

Helie, S., Chakravarthy, S., & Moustafa, A. A. (2013). Exploring the cognitive and motor functions of the basal ganglia: An integrative review of computational cognitive neuroscience models. *Frontiers in Computational Neuroscience, 7*, 174.

Helie, S., Ell, S.W., & Ashby, F.G. (2015). Learning robust cortico-frontal associations with the basal ganglia: An integrative review. *Cortex*, *64*, 123-135.

Helie, S., Roeder, J. L., & Ashby, F. G. (2010). Evidence for cortical automaticity in rule-based categorization. *Journal of Neuroscience, 30*, 14225-14234.

Izhikevich, E. M. (2007). *Dynamical Systems in Neuroscience: The Geometry of Excitability and Bursting*. Cambridge, MA: MIT Press.

Kac, M. (1962). A note on learning signal detection. *IRE Transactions on Information Theory, IT-8*, 126-128.

Kalish, M. L., Lewandowsky, S., & Davies, M. (2005). Error-driven knowledge restructuring in categorization. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, *31*, 846–861.

Kubovy, M., & Healy, A. F. (1977). The decision rule in probabilistic categorization: What it is and how it is learned. *Journal of Experimental Psychology: General, 106*, 427-446.

Lapish, C. C., Kroener, S., Durstewitz, D., Lavin, A., Seamans, J. K. (2007). The ability of the mesocortical dopamine system to operate in distinct temporal modes. *Psychopharmacology*, *191*, 609-625.

Little, D. R., Lewandowsky, S., & Heit, E. (2006). Ad hoc category restructuring. *Memory & Cognition*, *34*, 1398–1413.

Maddox, W. T. (2002). Toward a unified theory of decision criterion learning in perceptual categorization. *Journal of the Experimental Analysis of Behavior*, *78*, 567–595.

Maddox, W. T., Ashby, F. G. (2004). Dissociating explicit and procedural-learning based systems of perceptual category learning. *Behavioral Processes, 66,* 309-332.

Miller, E. K., & Cohen, J. D. (2001). An integrative theory of prefrontal cortex function. *Anuual Review of Neuroscience*, *24*, 167-202.

Muhammad, R., Wallis, J. D., & Miller, E. K. (2006). A comparison of abstract rules in the prefrontal cortex, premotor cortex, inferior temporal cortex, and striatum. *Journal of Cognitive Neuroscience*, *18*, 974-989.

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*, 95–102.

Owen, A. M., Roberts, A. C., Hodges, J. R., Summers, B. A., Polkey, C. E., & Robbins, T. W. (1993). Contrasting mechanisms of impaired attentional set-shifting in patients with frontal lobe damage or Parkinsons's disease. *Brain, 116*, 1159-1175.

Paul, E. J., & Ashby, F. G. (2013). A neurocomputational theory of how explicit learning bootstraps early procedural learning. *Frontiers in Computational Neuroscience*, *7*, 177.

Proulx, R., & Helie, S. (2005). Adaptive categorization and neural networks. In C. Lefebvre & H. Cohen (Eds.) *Handbook of Categorization in Cognitive Science* (pp. 793-815). Oxford: Elsevier.

Rall, W. (1967). Distinguishing theoretical synaptic potentials computed for different soma-dendritic distributions of synaptic input. *Journal of Neurophysiology*, *30*, 1138-1168.

Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, *85*, 59-108.

Roe, R. M., Busemeyer, J. R., & Townsend, J. T. (2001). Multialternative decision field theory: A dynamic connectionst model of decision making. *Psychological Review*, *108*, 370–392.

Seamans, J. K., & Robbins, T. W. (2010). Dopamine modulation of prefrontal cortex and cognitive function. In K. A. Neve (ed.). *The Dopamine Receptors. Second Edition* (pp. 373-398). Springer.

Shepard, R. N. (1987). Toward a universal law of generalization for psychological science. *Science*, *237*, 1317-1323.

Shepherd, G. M. (2004). Introduction to synaptic circuits. In G. M. Shepherd (Ed.). *The Synaptic Organization of the Brain. Fifth Edition* (pp. 1-38). Oxford University Press.

Stuber, G. D., Hnasko, T. S., Britt, J. P., Edwards, R. H., & Bonci, A. (2010). Dopaminergic terminals in the nucleus accumbens but not the dorsal striatum corelease glutamate. *Journal of Neuroscience*, *30*, 8229-8233.

Tenenbaum, J. B., & Griffiths, T. L. (2001). Generalization, similarity, and Bayesian inference. *Behavioral and Brain Sciences*, *24*, 629-641.

Treisman, M., & Williams, T. C. (1984). A theory of criterion setting with an application to sequential dependencies. *Psychological Review, 91*, 68-111.

Wallis, J. D. (2007). Single neuron activity underlying behavior-guiding rules. In S. A. Bunge & J. D. Wallis (Eds.). *Neuroscience of Rule-Guided Behavior* (pp. 23-44). New York: Oxford University Press.

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

White, I. M., & Wise, S. P. (1999). Rule-dependent neuronal activity in the prefrontal cortex. *Experimental Brain Research*, *126*, 315-335.

10 Appendix

The rule selection/switching mechanism was implemented using *surprise* and *confidence*. In each trial, a surprise score is calculated using the city-block distance between the current stimulus and the mean of all the pre-shift stimuli in the experiment:

$$surprise(t) = \sum_i |x_i(t) - \bar{x}_i| \tag{A1}$$

where $x_i(t)$ is the ith dimension of the stimulus presented at time t and $\bar{x}_i$ is the ith dimension of the mean stimulus. The surprise on each trial is used to update the rule confidence. If the response on trial t is correct, the confidence is updated as follows:

$$confidence(t) = min\left[\tau_{rw}, confidence(t-1) + surprise(t)^\gamma\right] \tag{A2a}$$

where τ_{rw} is an absorbing barrier for the random walk and $\gamma = 2.6$ is the rate of the exponential gradient (Shepard, 1987). If the response in trial t is incorrect, the confidence is updated as follows:

$$confidence(t) = confidence(t-1) - surprise(t)^\gamma \qquad \text{(A2b)}$$

In addition, the rule salience is increased or decreased in each trial by $\Delta = 0.04$ depending on feedback (Ashby et al., 1998).[5] The confidence is initialized at 0, and a rule change happens when the confidence reaches $-\tau_{rw}$. Hence, this process corresponds to a win-stay/lose-shift strategy. When the rule changes, a new rule (i.e., stimulus dimension) is stochastically selected using the rule saliences, the rule weights are re-initialized, and the confidence score is reset to 0.[6] For the present simulations, the random walk barrier was set to $\tau_{rw} = 1.75 \times 10^6$.

---

[5] Note that the rule salience is bounded between [0, 1].

[6] The new rule may correspond to the same stimulus dimension that was already used, but the rule weights and confidence score are reset.

11   Figure Captions

(a)   (b)
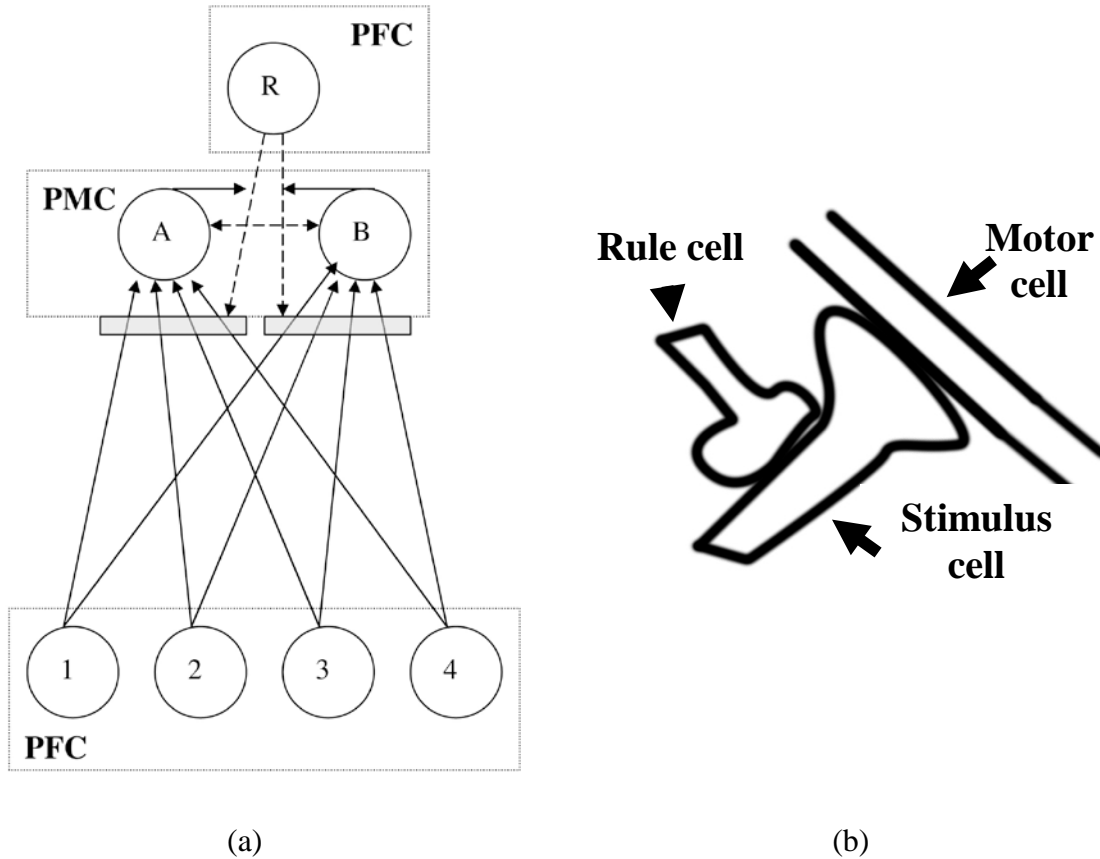
Figure 1. (a) Connectivity for rule representation in HICL. Numbers correspond to categorical stimulus cell populations, 'A' and 'B' correspond to motor plan cell populations, and 'R' represents one rule cell population. Solid lines correspond to excitatory (glutamate) connections while dashed lines correspond to inhibitory (GABA) connections. The grey rectangles represent pre-synaptic inhibition of stimulus activity by the rule cell. (b) A schematic of pre-synaptic inhibition. Here, the rule cell pre-synaptically inhibits (through a GABAergic interneuron) excitation from the categorical stimulus cell before it reaches the motor planning cell. The weight of pre-synaptic inhibition is different at each synapse, so that the rule cell acts as a gating mechanism to implement stimulus selectivity.
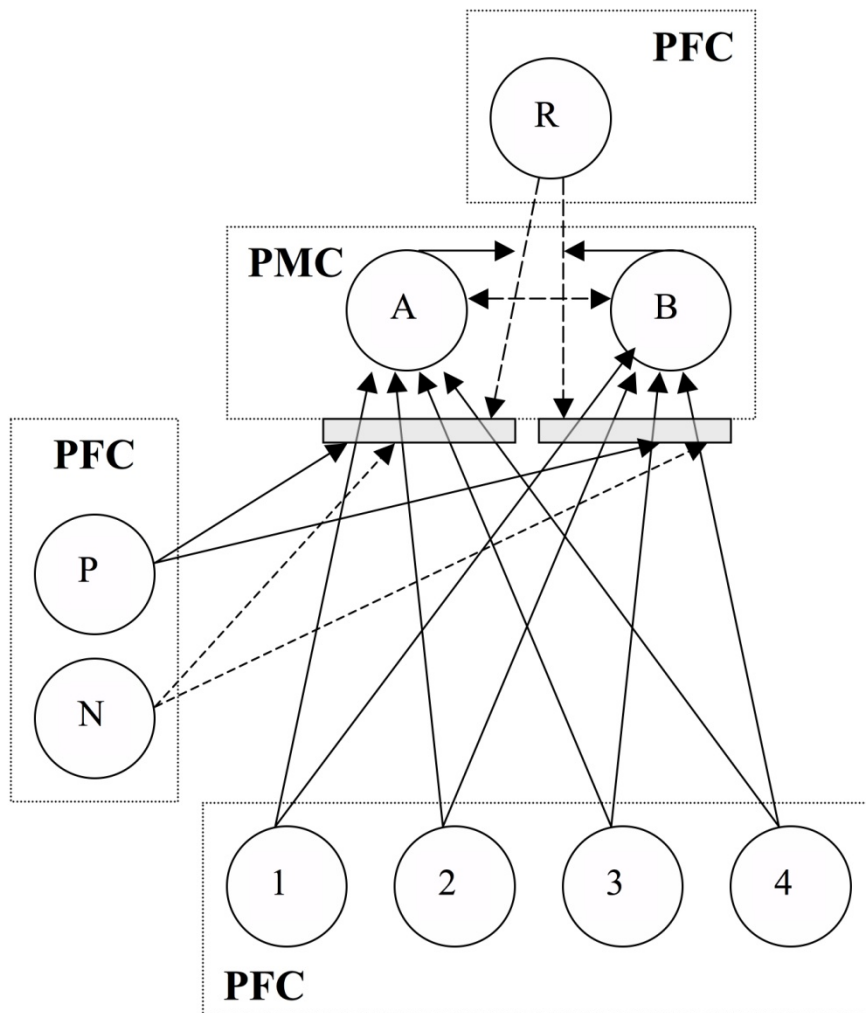
Figure 2. Structural diagram of the HICL model. This architecture is identical to that of Figure 1a except for the addition of feedback units in the prefrontal cortex.
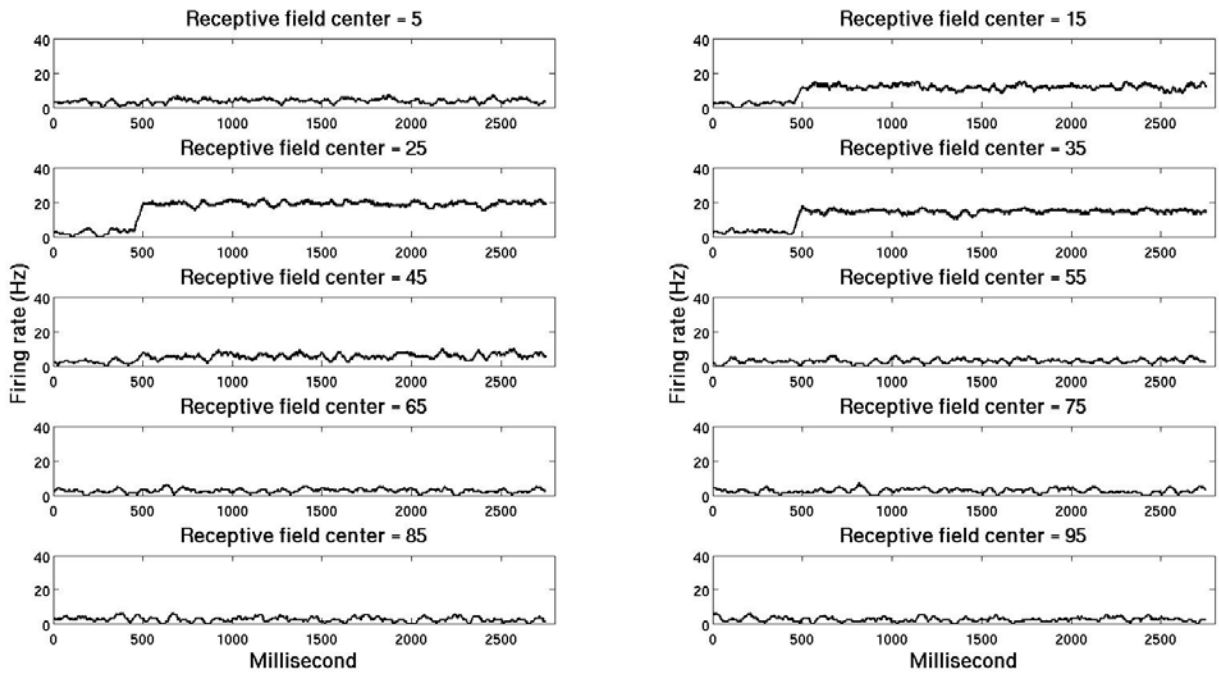
Figure 3. Firing rate of each stimulus categorical unit for a trial where the stimulus was 26.23.
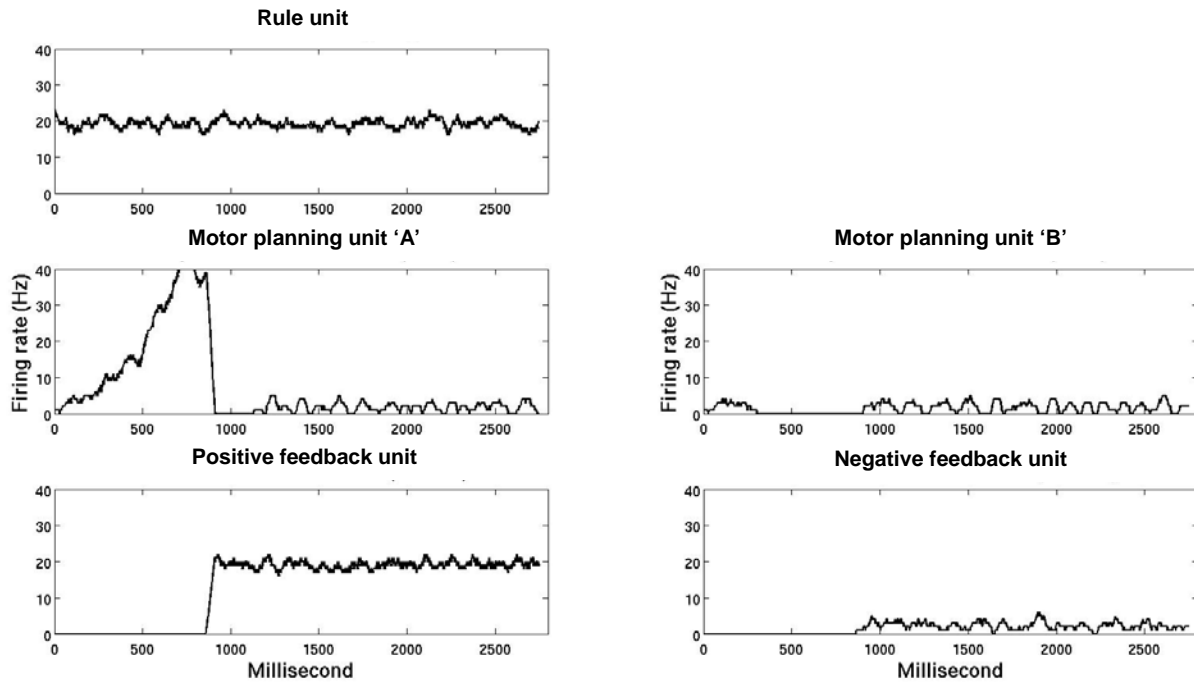
Figure 4. Firing rate of the response, rule, and feedback units for a trial where the stimulus was 26.23. In this case, the model produced the correct 'A' response.
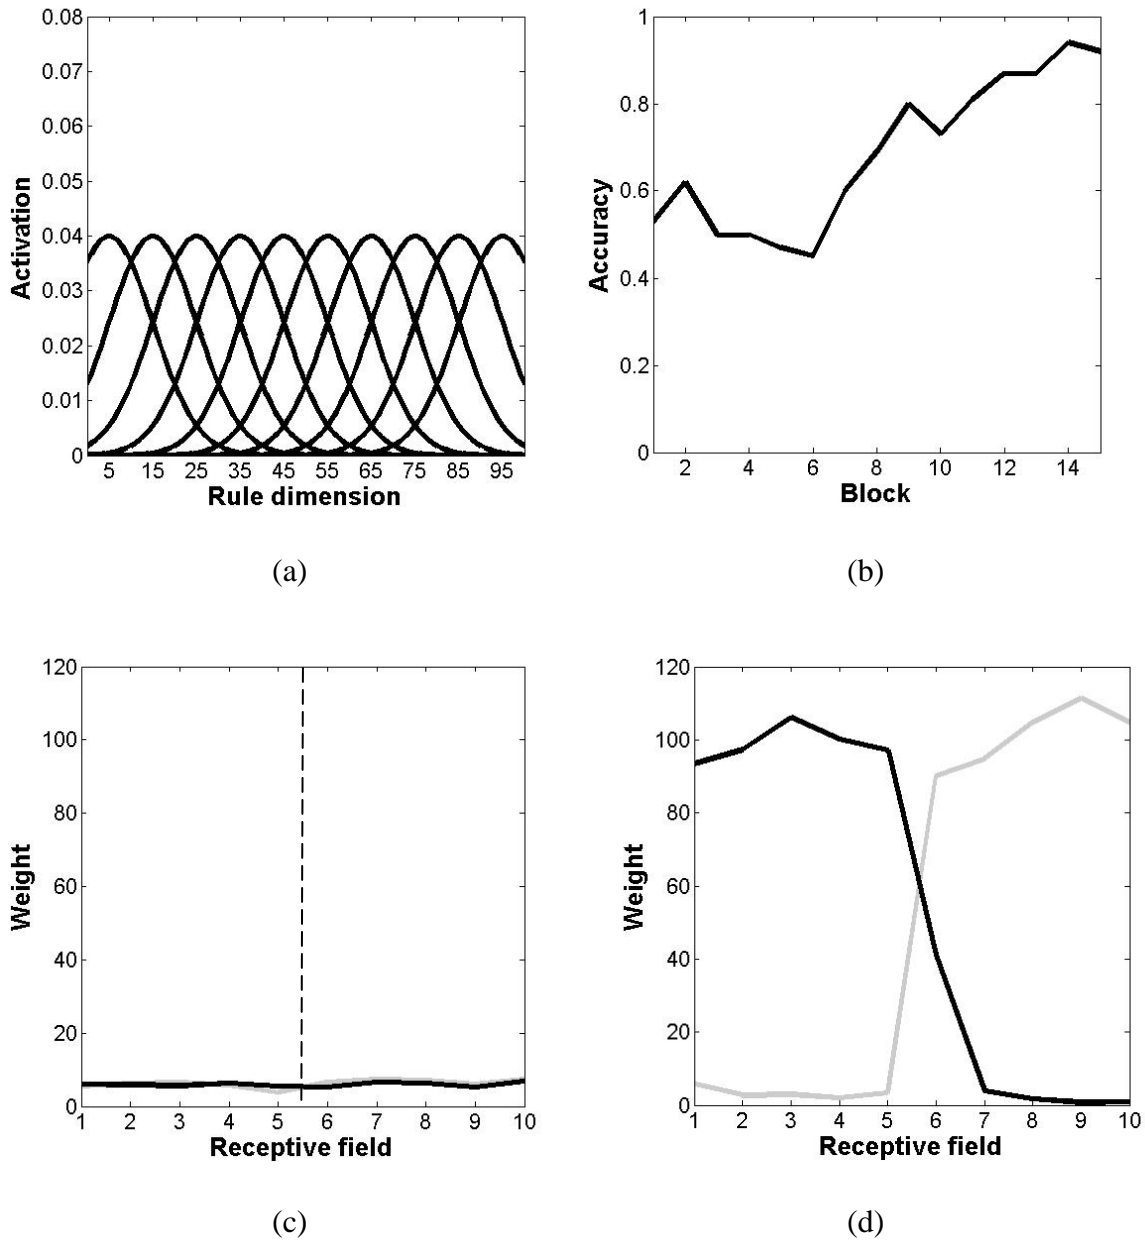
(a)

(b)

(c)

(d)

Figure 5. Simulating a toy problem with HICL. (a) The receptive fields of the categorical stimulus units in PFC. (b) Categorization accuracy for each block of 100 trials. (c) Initial rule weights. (d) Rule weights after 1,500 trials of training. In (c) and (d), the black line shows the "B" weights while the gray line shows the "A" weights. The learned "criterion" is located at the line crossing in (d) ($x \sim 55$).
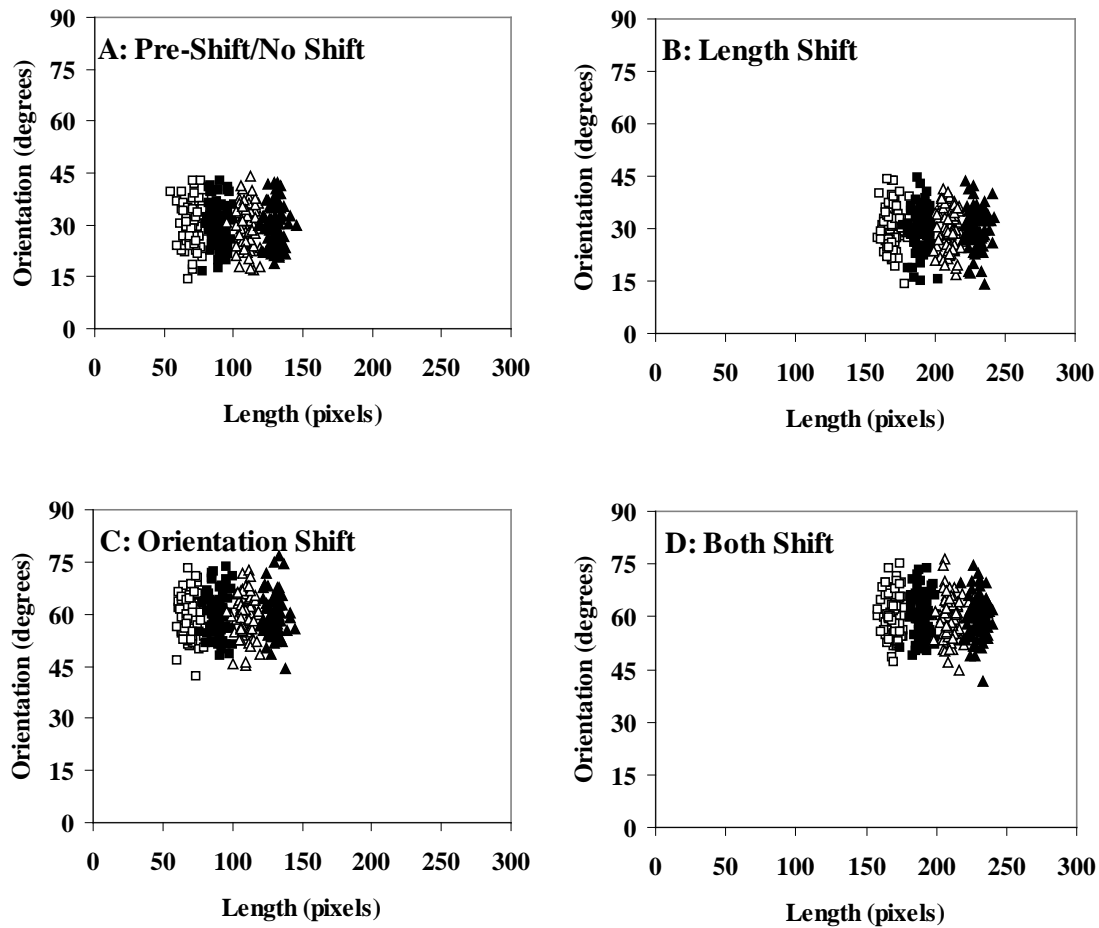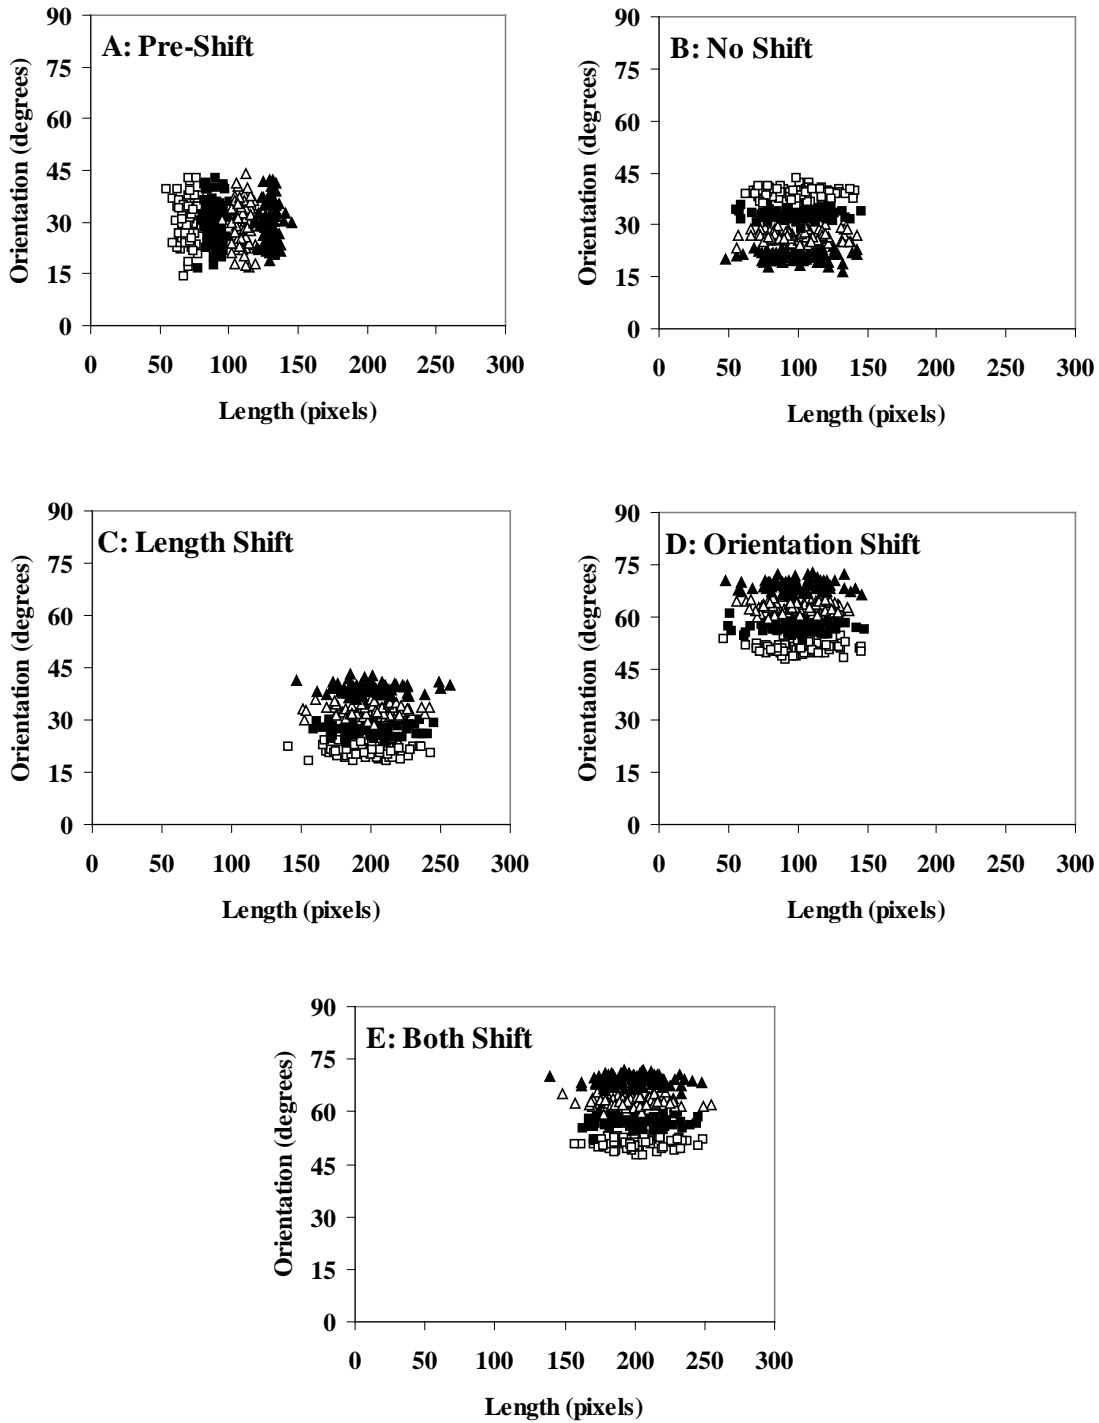
Figure 6. Scatterplots of the stimuli used in the four conditions from Experiment 1 (Intradimensional Shift/ID). Open squares denote stimuli from category A, filled squares denote stimuli from category B, open triangles denote stimuli from category C, and filled triangles denote stimuli from category D.
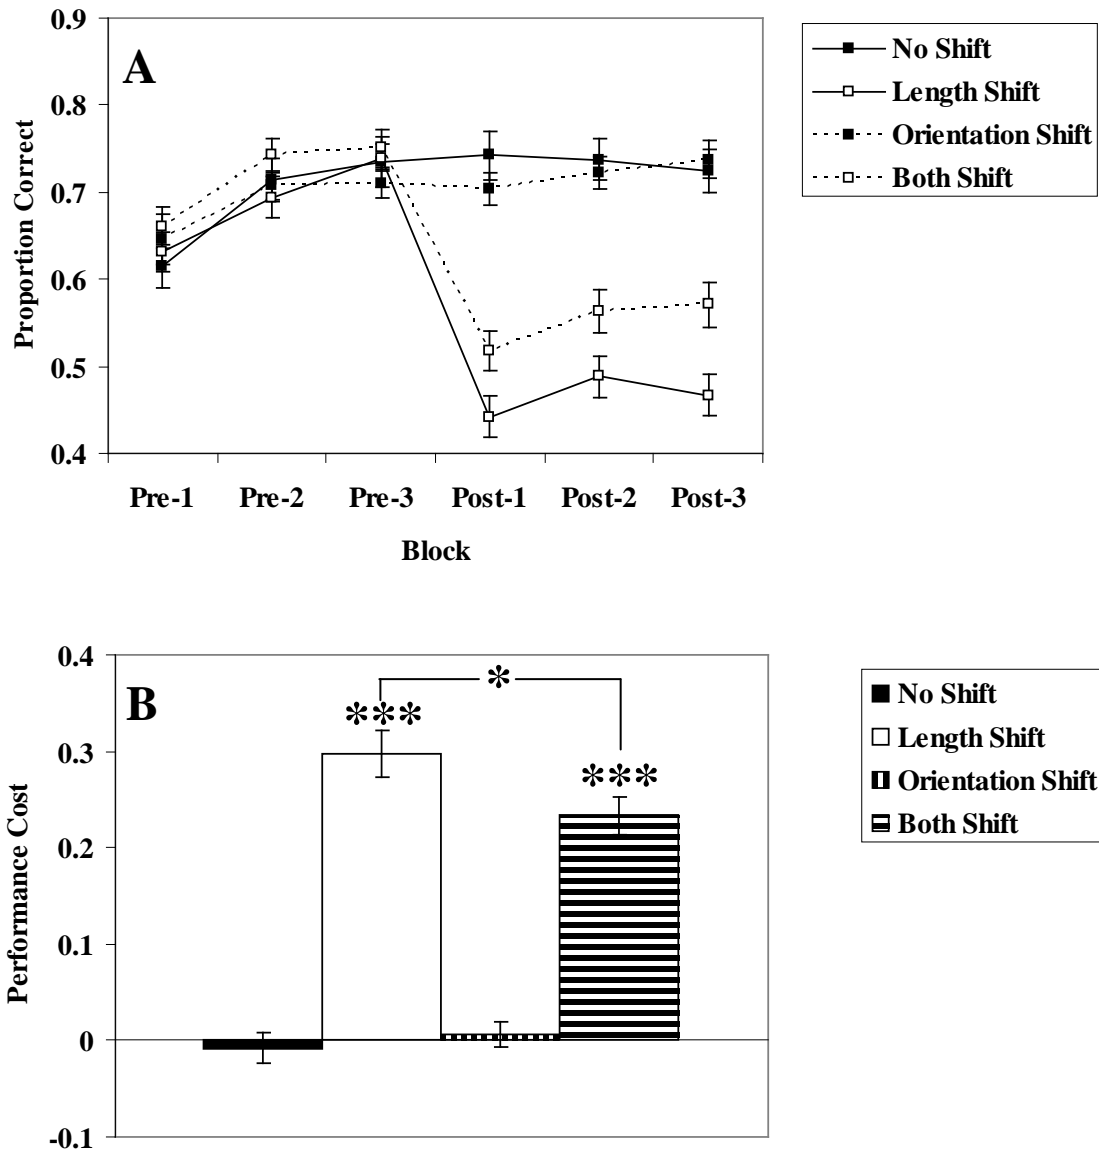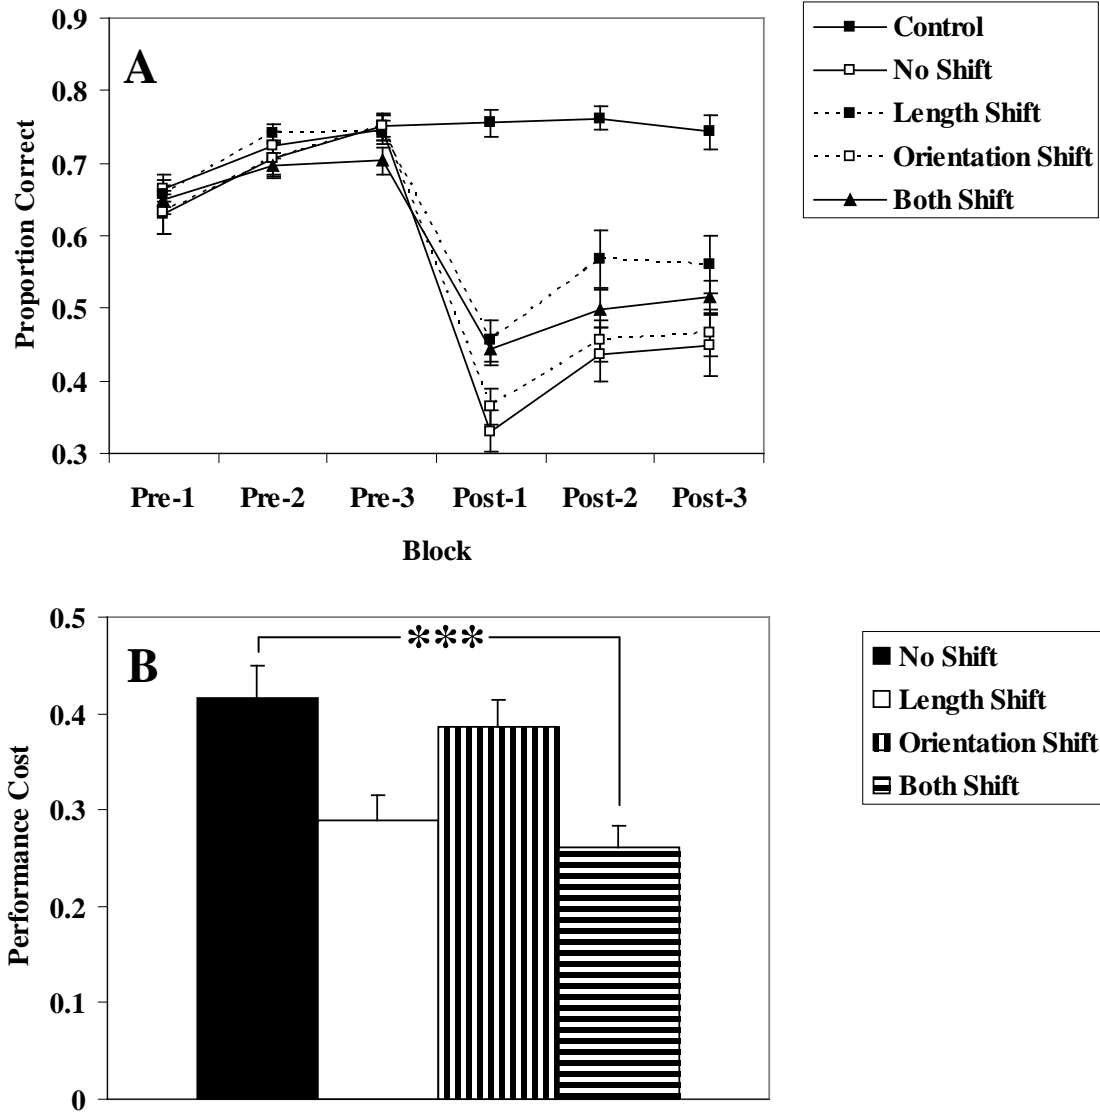
Figure 7. A. Proportion correct (averaged across participants) from Experiment 1 (Intradimensional shift/ID). B. Performance cost determined by subtracting post-shift block 1 performance from pre-shift block 3 performance. Standard error bars included. *** = $p < .001$; * = $p < .05$.

Figure 8. Scatterplots of the stimuli used in the five conditions from Experiment 2 (Extradimensional shift/ED). Open squares denote stimuli from category A, filled squares denote stimuli from category B, open triangles denote stimuli from category C, and filled triangles denote stimuli from category D.

Figure 9. A. Proportion correct (averaged across participants) from Experiment 2 (Extradimensional shift/ED). B. Performance cost determined by subtracting post-shift block 1 performance from pre-shift block 3 performance. Standard error bars included. All switch costs are significantly bigger than 0 ($p < .001$). Also, not shifting the length condition (i.e., as in the no shift and orientation shift conditions) yielded a bigger shift cost than shifting the length condition (i.e., as in the length shift and both shift conditions) ($p < .001$). *** = $p < .001$.
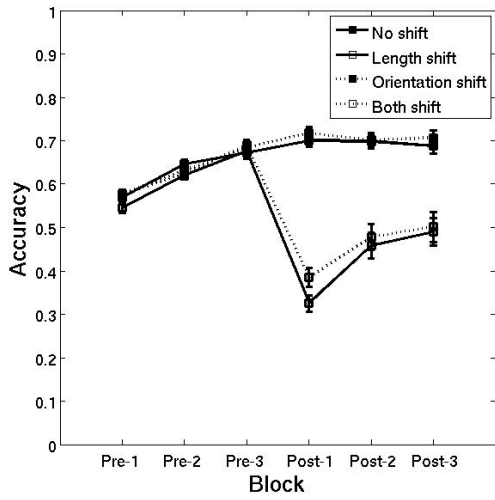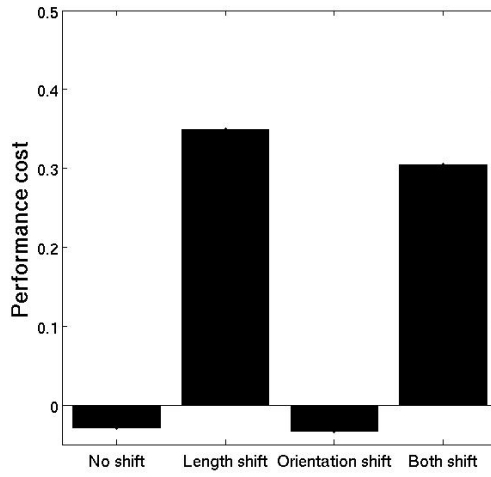
(a)   (b)

(c)   (d)

Figure 10. Simulation results. (a) Learning curves for the simulated Experiment 1 (ID). (b) Switch cost for the simulated Experiment 1 (ID). (c) Learning curves for the simulated Experiment 2 (ED). (d) Switch cost for the simulated Experiment 2 (ED).