



# Interactions of Medial and Lateral Prefrontal Cortex in Hierarchical Predictive Coding

William H. Alexander<sup>1\*</sup> and Thilo Womelsdorf<sup>2</sup>

<sup>1</sup> Center for Complex Systems and Brain Sciences, Florida Atlantic University, Boca Raton, FL, United States, <sup>2</sup> Department of Psychology, Vanderbilt University, Nashville, TN, United States

Cognitive control and decision-making rely on the interplay of medial and lateral prefrontal cortex (mPFC/IPFC), particularly for circumstances in which correct behavior requires integrating and selecting among multiple sources of interrelated information. While the interaction between mPFC and IPFC is generally acknowledged as a crucial circuit in adaptive behavior, the nature of this interaction remains open to debate, with various proposals suggesting complementary roles in (i) signaling the need for and implementing control, (ii) identifying and selecting appropriate behavioral policies from a candidate set, and (iii) constructing behavioral schemata for performance of structured tasks. Although these proposed roles capture salient aspects of conjoint mPFC/IPFC function, none are sufficiently well-specified to provide a detailed account of the continuous interaction of the two regions during ongoing behavior. A recent computational model of mPFC and IPFC, the Hierarchical Error Representation (HER) model, places the regions within the framework of hierarchical predictive coding, and suggests how they interact during behavioral periods preceding and following salient events. In this manuscript, we extend the HER model to incorporate real-time temporal dynamics and demonstrate how the extended model is able to capture single-unit neurophysiological, behavioral, and network effects previously reported in the literature. Our results add to the wide range of results that can be accounted for by the HER model, and provide further evidence for predictive coding as a unifying framework for understanding PFC function and organization.

**Keywords:** medial prefrontal cortex, lateral prefrontal cortex, cognitive control, attention, learning, predictive coding, computational neuroscience

## OPEN ACCESS

### Edited by:

Paul Miller,  
Brandeis University, United States

### Reviewed by:

Mohammad M. Herzallah,  
Al-Quds University, Palestine  
Thomas Wennekers,  
University of Plymouth,  
United Kingdom

### \*Correspondence:

William H. Alexander  
walexander@fau.edu

**Received:** 11 September 2020

**Accepted:** 08 January 2021

**Published:** 03 February 2021

### Citation:

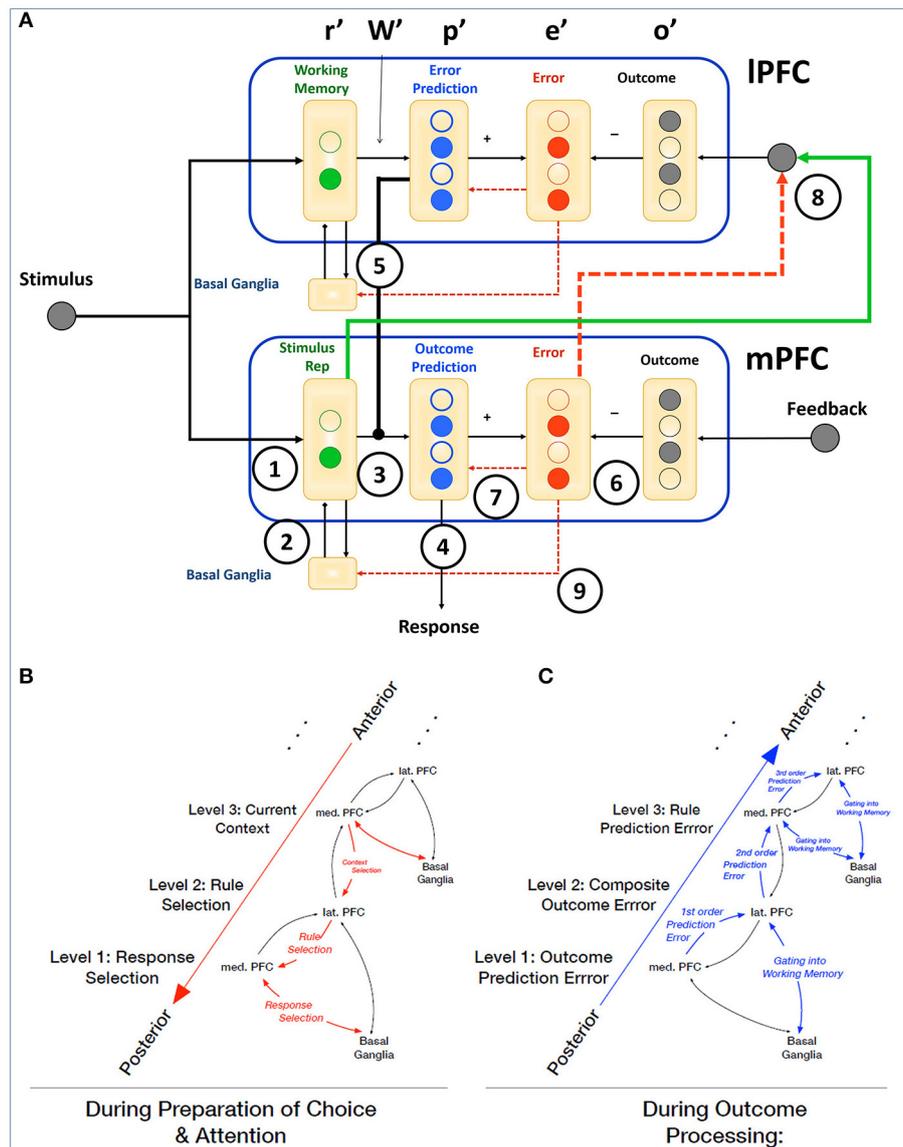
Alexander WH and Womelsdorf T  
(2021) Interactions of Medial and  
Lateral Prefrontal Cortex in  
Hierarchical Predictive Coding.  
*Front. Comput. Neurosci.* 15:605271.  
doi: 10.3389/fncom.2021.605271

## INTRODUCTION

Medial and lateral prefrontal cortex (mPFC/IPFC) are core hubs of the cognitive control and decision-making network in the brain (Cole and Schneider, 2007). The regions are densely and reciprocally connected (Barbas and Pandya, 1989; Barbas and Rempel-Clower, 1997; Náchér et al., 2019), suggesting that their contribution to behavior depends in part on their tightly-coupled interactions during preparation, execution, and monitoring of the consequences of actions. Although these regions have long been the target of focused investigation, it remains an open question as to how they collaborate in an ongoing and interactive fashion to support adaptive behavior (Badre and Nee, 2018).

Independently, both regions have been implicated in a wide range of functions, and these functions appear to suggest a dissociation between the regions along a temporal dimension. Activity in mPFC, especially dorsal anterior cingulate cortex (dACC), is frequently observed in conjunction with behaviorally-salient events, such as error commission (Gehring et al., 1990; Shen et al., 2015) or the appearance of stimuli cueing multiple, potentially conflicting, responses (Brown, 2009). In contrast, activity

in IPFC is typically associated with long-term behavioral contingencies, such as maintaining information over extended periods of time (Sawaguchi and Goldman-Rakic, 1991) and representing the structure of ongoing tasks (Badre and D'Esposito, 2007). This temporal dissociation has been interpreted as reflecting complementary aspects of task performance in which mPFC signals changes in behavioral requirements, including specification of control signals (Shenhav et al., 2013) or selection of action policies



**FIGURE 1 |** Interactions of mPFC and IPFC in the HER model. **(A)** The HER model is organized hierarchically, with each hierarchical level instantiating a computational motif of prediction and prediction error calculations. Information flows between hierarchical levels along top-down and bottom-up pathways, which carry information regarding likely errors in a given context (top-down), and error signals (bottom-up) derived from external feedback (base hierarchical level) conjoined with items maintained in working memory (additional hierarchical levels). Circled numbers indicate governing equations within and between hierarchical levels (see section Methods). Reciprocal connections between mPFC and IPFC correspond to a putative hierarchical rostrocaudal gradient (Badre and D'Esposito, 2009) with bottom-up and top-down pathways governing the adjustment of behavior during ongoing task performance. **(B)** During a trial, top-down connections serve to establish the current context and relevant rules that govern eventual behavioral responses during preparatory periods. **(C)** Following the generation of a response (right), bottom-up error signals are used to update outcome predictions and derive composite "proxy" outcomes that train higher-order representations of rules and contexts.

(Holroyd and Yeung, 2012), and the implementation of appropriate control measures is delegated to LPFC (Botvinick et al., 2001).

How such temporally specific functional contributions of medial and lateral PFC can give rise to a wide variety of cognitive and behavioral effects has been formalized within the framework of hierarchical predictive coding (Alexander and Brown, 2015, 2018). The Hierarchical Error Representation (HER) model states that error signals generated by mPFC are used to train *error representations* in LPFC, and that, once learned, error representations maintained by LPFC serve to contextualize subsequent error calculations carried out by mPFC (Figure 1). Using this basic circuit as a repeating computational motif, the HER model is able to learn complex cognitive tasks in a manner that accords with human behavior, and measures of activity derived from error calculation and representation in the model reproduce qualitative patterns of single-unit and neuroimaging data.

While the HER model suggests how and when mPFC and LPFC might interact, our previous implementation of the model was aimed at the “event” level (Alexander and Brown, 2015): model equations were applied, and activity derived, following the occurrence of salient behavioral events, such as the presentation of task stimuli or performance feedback. The HER model is therefore only able to account for mPFC/IPFC interactions at a temporally coarse level, and is unable to capture the co-evolving development of activity during intra-event periods, nor can it capture aspects of behavior, such as switch costs (Wylie and Allport, 2000), that may manifest through differences in response time. Thus, although the HER model provides a promising framework for understanding the function, organization and interactions of PFC regions (Alexander and Brown, 2018), a significant gap in its explanatory power remains.

To address this limitation and derive novel model predictions, we here incorporate temporal dynamics to the model-unit activity of the HER model. When model units transform inputs to outputs during performing of a reversal learning task we simulate their activation with a non-linear “shunting” equation (Grossberg, 1988). These dynamics reproduce the rise and fall of neuronal activations in cortical circuits. This allowed us to analyze the timing and duration of computations at different hierarchical levels of the model and to relate these dynamics to the behavioral, single-unit, and brain network data of empirical studies.

## METHODS

### Model Equations

Details regarding the operation of the HER model are described in previous publications (Alexander and Brown, 2015). Here we review the equations specifying the model, adopting the notational convention that **BOLD UPPERCASE** variables indicate matrices, **bold lower case** variables indicate column vectors, and *lowercase italicized* variables indicate scalars. Variables marked with a single apostrophe (e.g., **P'**) indicate superior hierarchical levels. **Table 1** lists parameters in the model, their interpretation, and their value for the simulations reported

**TABLE 1** | Model Parameters.

Parameter	Description	Equation	Parameter Values	
			Hierarchical Level 1	Hierarchical Level 2
$\alpha$	Response learning rate	7	0.05	0.02
$\lambda$	Working memory learning rate	10	0.3	0.5
$\beta$	Working memory update gain	2	12	14
$\gamma$	Response selection temperature	4	12	N/A

**TABLE 2** | Model Variables.

Variable	Description	Equation
<b>v</b>	Utility for storing item in WM	(1, 2)
<b>X</b>	Associative WM weights	(2, 10)
<b>s</b>	Stimulus vector	(1, 10)
<b>p</b>	Model-generated predictions	(3, 4, 6)
<b>P</b>	Matricized predictions ( <b>p</b> )	(5)
<b>W</b>	Associated response weights	(3, 5, 7, 9)
<b>r</b>	Working memory representation	(3, 7, 8, 9)
<b>e</b>	Error	(6, 7, 8)
<b>o</b>	Observed response-outcome vector	(6, 8)
<b>e<sub>WM</sub></b>	Backpropagated working memory error	(9, 10)

here. **Table 2** lists model variables, their interpretation, and equations in which they appear.

Generally, the HER model is composed of 2 or more hierarchical levels, and each level of the hierarchy consists of three primary components: a working memory (WM) store **r**, and weight matrix **X** determining the probability that a stimulus **s** will be stored in **r**, and a weight matrix **W** determining how items stored in WM either influence behavior (at the lowest hierarchical level) or modulate the processing of lower hierarchical levels. Whether a stimulus **s** is stored in working memory is determined by the pair of equations:

$$\mathbf{v} = \mathbf{X}^T \mathbf{s} \quad (1)$$

$$\text{and the probability of WM representation } r_i = \frac{\exp^{\beta v_i}}{\exp^{\beta v_i} + \exp^{\beta v_j}} \quad (2)$$

Equation (1) reflects the learned value of storing an item in working memory—at each hierarchical level of the model. Here, **s** is a binary vector of stimulus representations reflecting the current trial, with each element indicating the presence (1) or absence (0) of the corresponding task feature (cf. the task description below). The relative values of maintaining the current contents of WM (indexed by *j*) vs. updating WM with a new representation (indexed by *i*) are compared in Equation (2). Equation (2) is a softmax function calculating the probability

of storing the current stimulus  $s_i$  in WM relative to the current contents of WM and  $\beta$  is a gain parameter governing the degree is a gain parameter governing the degree by which differences in elements of  $\mathbf{v}$  influence the probability of maintaining or updating WM. Intuitively, these equations determine which of multiple “valuable” to store in WM (Equation 1), and transform this value into a probability (Equation 2). WM memory transform this value into a probability (Equation 2). WM memory storage is binary, and only a single WM any time.

The WM representation  $\mathbf{r}$  and associative weights  $\mathbf{W}$  are used to generate the output  $\mathbf{p}$  of that hierarchical level:

$$\mathbf{p} = \mathbf{W}^T \mathbf{r} \quad (3)$$

At the lowest hierarchical level, network output  $\mathbf{p}$  is used to generate a behavioral response  $i$  (out of  $j$  possible responses) using another softmax function over elements of  $\mathbf{p}$ :

$$\text{Probability of response } i = \frac{\exp^{\gamma p_i}}{\sum_j \exp^{\gamma p_j}} \quad (4)$$

where  $\gamma$  is a gain parameter indicating how likely the model is to select the response  $i$  from  $j$  possible responses. The output of higher-order hierarchical levels,  $\mathbf{p}'$ , is used to modulate the processing of lower-order levels output as follows. First,  $\mathbf{p}'$  is reshaped from a vector to a matrix with the same dimensionality as  $\mathbf{W}$ . That is,  $\mathbf{p}'$  is matricized as  $\mathbf{P}'$ . Next, modulated weight matrices are calculated as:

$$\mathbf{W}^{mod} = \mathbf{W} + \mathbf{P}' \quad (5)$$

Note that  $\mathbf{p}'$  is constrained to have the same number of elements as  $\mathbf{W}$  so that  $\mathbf{P}'$  has the same dimensionality following matricization.

## Learning

Learning in the model depends on the calculation of error signals at each level of hierarchy. At the base hierarchical level, errors are computed as:

$$\mathbf{e} = \mathbf{a} \circ (\mathbf{o} - \mathbf{p}) \quad (6)$$

where  $\mathbf{o}$  is a feedback vector indicating the response identity and whether the response was correct or incorrect, and  $\mathbf{a}$  is a filter that is 1 for the index of the generated response and 0 everywhere else. The operator  $\circ$  reflects elementwise multiplication (Hadamard product) of two matrices. Associative weights at each level are updated according to:

$$\mathbf{W}_{t+1} = \mathbf{W}_t + \alpha (\mathbf{e}_t \mathbf{r}_t^T) \quad (7)$$

and  $\alpha$  is the learning rate. To train higher-order hierarchical levels, a proxy outcome signal is composited from the error signal computed at the lower level (Equation 6) and the active WM representation  $\mathbf{r}$  at the lower level. This is done by taking the outer product of the lower-level error and representation vectors:

$$\mathbf{O}' = \mathbf{r} \mathbf{e}^T \quad (8)$$

For computational convenience,  $\mathbf{O}'$  is vectorized as  $\mathbf{o}'$  with the same dimensionality as  $\mathbf{p}'$ .

Training of WM weights  $\mathbf{X}$  is done by backpropagating the output error term  $\mathbf{e}$  to derive an error term for the current WM representation:

$$\mathbf{e}_{WM} = \mathbf{W}_t^T \mathbf{e}_t \circ \mathbf{r}_t \quad (9)$$

$$\mathbf{X}_{t+1} = \mathbf{X}_t + \lambda (\mathbf{e}_{WM} \mathbf{s}_t^T) \quad (10)$$

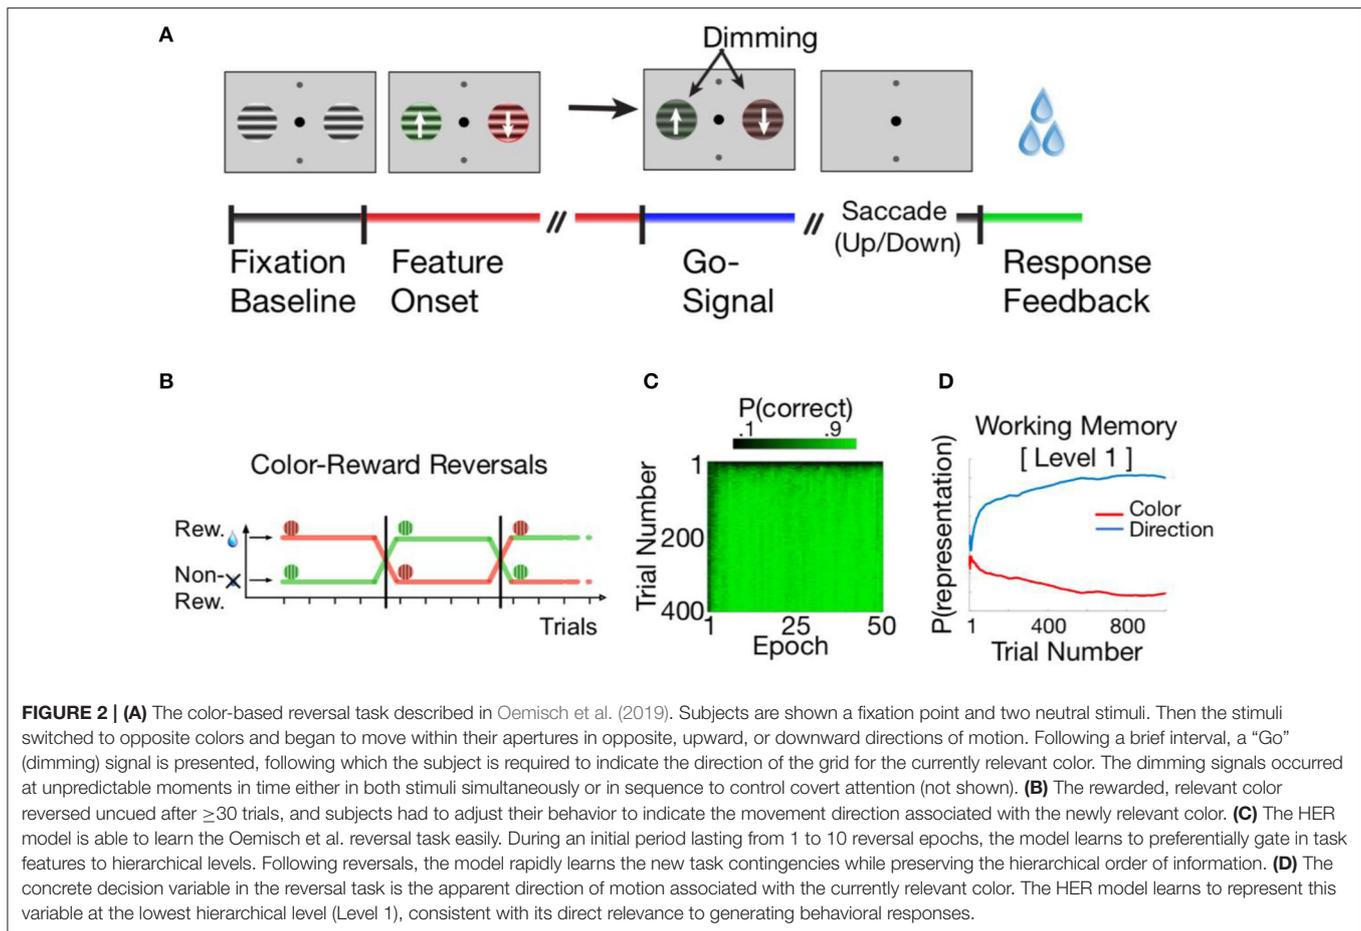
with learning rate  $\lambda$ .

## Simulated Task

The HER model was developed to account for behavior and brain function in tasks requiring the selection and application of rules that govern how to respond to a concrete stimulus (e.g., an arrow cueing response identity). We therefore selected a recent, rule-based task (Oemisch et al., 2019) in which, on each trial, monkeys were presented two stimuli, one on each side of the screen (Figure 2A). Stimuli had two behaviorally relevant features that varied independently from each other, namely color (red/green), and the direction of motion of a grating inside a circular aperture (up/down). The task required that the monkey learn through trial-and-error which of the two colors is currently relevant, and to respond with a saccade in the up- or downward motion direction of the stimulus with the currently relevant color. Once the relevant color has been learned, successful performance of the task requires that, on each trial, the monkey first identifies the *location* (side of the screen) of the appropriately-colored stimulus, after which the identity of the *response* (up/down) corresponding to the direction of the pattern movement is determined. The task described in Oemisch et al. additionally incorporates periodic reversals when the relevant color changes and the monkey must adapt their responses to the new task contingencies (Figure 2B); while the HER model is able to learn such reversals (cf. Figure 2C), this aspect of the task is not further considered here as it does not bear directly on the *within-trial* temporal dynamics of PFC interactions.

## Model Simulations

To simulate real-time dynamics in the HER model, we adopt the approach taken by previous studies of network models of cognitive control [e.g., Yeung et al. (2004)] in which the process used to establish weights suitable for performing a task is carried out independently of simulations of temporal dynamics. In order to establish model weights suitable for performing the Oemisch et al. reversal task, a previously-described version of the (event-level) HER model (Alexander and Brown, 2015) using the same parameters was trained on 20,000 trials of the reversal task, divided into 50 blocks of 400 trials each. On each trial during learning, 3 events were modeled: the onset of task stimuli, the occurrence of a response and feedback, and a neutral cue indicating the start of the intertrial interval (ITI). Stimuli were modeled as a binary vector, with elements corresponding to color and movement direction (four total elements per stimulus), and independent stimulus vectors were used to model each side of the simulated display, for a total of eight task stimulus elements. A 9th element was used to indicate



the ITI. The model was permitted 3 responses, two corresponding to movement direction (up/down), and one neutral response indicating acknowledgment of ITI onset. Each response could be associated with two outcomes (correct/incorrect), for a total of 6 response-outcome predictions. Responses were generated by subtracting, for each response, the prediction of incorrect feedback from the prediction of correct feedback, and passing the values through a softmax function, as described in the methods section. Code for the HER model is available at <https://www.github.com/modelbrains>.

## Temporal Dynamics

The network weights recovered from the training procedure were fixed during the real-time simulations as described above. In order to simulate real-time dynamics in the network, changes in unit activity on each cycle were modeled by a non-linear “shunting” equation (Grossberg, 1988):

$$\Delta E_i = exc_i (\theta - E_i) - (E_i + \eta) + N(0, \sigma) \quad (11)$$

where  $\theta$  is the upper asymptotic activity a unit could achieve (set to a value of 10),  $\eta$  is the lower boundary toward which unit activity decays passively ( $-0.05$ ), and  $N$  is gaussian noise

applied to the signal change with mean 0 and variance  $\sigma = 0.01$ .  $E_i$  is the current activity of unit  $i$ , and  $exc_i$  is the current net excitatory input to the unit, computed as in equations 1, 3, and 6. Specifically,  $exc_i$  is equal to  $\mathbf{v}$  in Equation (1), governing the dynamics of WM update and maintenance. For computing model predictions at all hierarchical levels,  $exc_i$  is equal to  $\mathbf{p}$  in Equation (3). Finally, for computing error (Equation 6),  $exc_i$  is set to  $\mathbf{e}$ , the difference between predictions  $\mathbf{p}$  and observed outcomes  $\mathbf{o}$ .

Models of neuron activity similar to Equation (11) [e.g., Wilson and Cowan (1972)] approximate the dynamics of individual neurons (Hodgkin, 1969) without introducing unnecessary complexity. The approach to simulate activity dynamics with a typical neuronal activation function has advantages over the use of biophysical model neurons when models are complex and the aim is to understand the algorithmic as opposed to the implementational level of a computational problem (Niv and Langdon, 2016). The HER model architecture has not previously been implemented into a biophysical model in order to allow for an analysis of the model mechanisms at the level of the computational units across hierarchy as opposed to the level of biophysical mechanisms underlying these units.

Each simulated trial lasted up to 700 cycles, and on each cycle the processing steps outlined in Equations (1–10) were followed and unit activity updated according to Equation (11). During the first 100 cycles of the trial, no input was presented to the network. On cycle 101, a stimulus vector representing the color and movement direction of the left and right stimuli was presented to the network, after which the network was able to register a response. Responses were generated when the corresponding unit in  $\mathbf{p}$  reached 95% of its upper asymptotic value. Following a response, the stimulus vector was set to 0 and the network was provided feedback for 50 cycles. Following the offset of feedback, the network was run for an additional 250 cycles prior to the beginning of the next trial. Network activity was not re-initialized after each trial. As described in previous work (Alexander and Brown, 2011, 2014, 2015, 2018), Equations (1–10) specify two primary signal types, namely prediction (Equations 1 and 3) and error (Equation 6), and analyses were conducted on unit activity reflecting individual network units recruited during each trial. In the simulations reported here, both types of signals are subject to the temporal dynamics embodied in Equation (11). All model results were derived from 20 simulation runs.

## Granger Causality

We recorded the time course of unit activity for prediction (Equation 3) and error (Equation 6) signals from both hierarchical levels of the model. To determine the average causality between regions for each type of signal, we used the Granger Causality Toolbox (Luo et al., 2013; <https://www.dcs.warwick.ac.uk/~simfeng/causality.html>). The goal of Granger causality is to determine whether the value of a signal  $y$  at time  $t$  is better predicted by the history of another signal,  $x$ , than by its own history. The magnitude of Granger causality is given by the log ratio of the error when considering only the history of  $y$  vs. the history of both  $y$  and  $x$ . When the history of  $x$  does not improve predictions of  $y$  at time  $t$  above only the history of  $y$ , the log ratio is equal to 0 (i.e., the error is the same in both cases).

The length of the signal history to be considered is determined by the analyst; here we selected a time window of 50 model cycles. For causality analyses, trial timing was standardized—even if a response was indicated prior, the model was simulated for 350 cycles following the onset of a stimulus. The average Granger causality magnitude of one hierarchical level on the other for each signal type was estimated for each cycle. This procedure was repeated 20 times, and a two-sample  $t$ -test (Bonferroni corrected) was conducted for the distribution of Granger causality magnitudes at each cycle to determine whether the granger causality of one region on the other was significantly greater. Thus, although both regions could Granger-cause activity in the other at any time point, our analysis indicated whether one region had a significantly higher degree of causality than the other.

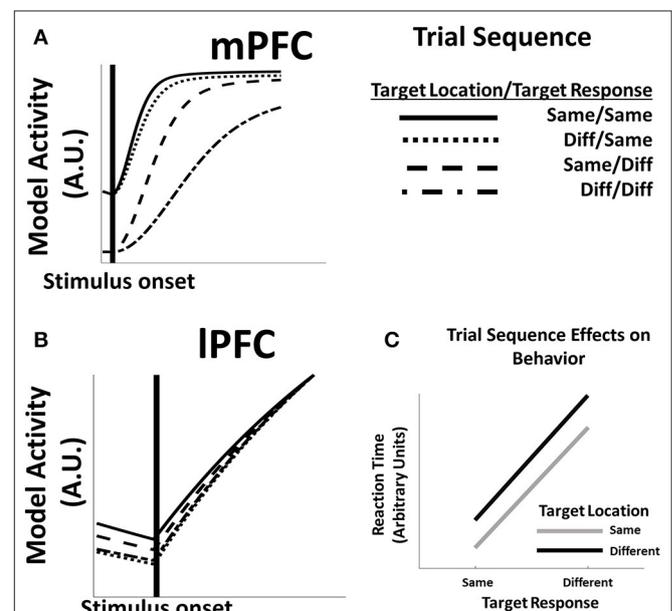
## RESULTS

The HER model was able to learn and perform the reversal task relatively easily. During an initial “burn-in” period (Figure 2C), the model primarily learned to hierarchically segregate relevant

feature dimensions (Figure 2D): the WM “gating” mechanism (Equations 1 and 2) learned to store the “concrete” movement direction at the lowest hierarchical level while color information was stored in the superior hierarchical level. This pattern accords with intuitions regarding how individuals might solve the task (Oemisch et al., 2019): The reward-associated color acts as a feature-reward “rule” that indicates which concrete stimulus (movement direction) should govern the ultimate response. After this initial learning period, the model is able to perform reversals within a short period of time. Because the model has learned stable WM mappings, learning to respond to the opposite color entails relatively rapid changes in top-down modulation of concrete responses, i.e., rather than relearning the task from the ground up. These findings suggest one manner in which a hierarchical representation of information might support rapid and flexible reconfiguration of responses in the face of changing task contingencies.

## Response Preparation

Introducing real-time dynamics allows us to investigate the evolution of predictive activity (Equation 3) in the model during



**FIGURE 3 | (A)** mPFC unit activity in model simulations is influenced by switches or repetitions of feature dimensions between trials. When both feature dimensions (location and response direction) repeat (**same/same**), as well as repetitions of the target direction (**diff/same**), mPFC activity reaches asymptote quickly after stimulus onset, while changes in the target response (**same/diff**, **diff/diff**) produce delays in the evolution of mPFC activity. **(B)** Activity in IPFC likewise shows intertrial effects; however, in this case, delay in the development of IPFC activity is due principally to shifts in the location of the target stimulus (**diff/same**, **diff/diff**), while changes in the more concrete target direction variable (**same/diff**) have relatively little influence on IPFC activity relative to a trial in which all features are the same (**same/same**). **(C)** The delay in model activity following feature switches contributes to changes in model reaction times: reaction times are most rapid on trials in which both feature dimensions repeat, while switches in either or both features result in longer reaction times.

a trial, the influence of previous trial effects, and derive measures of reaction time from model activity. During preparatory periods following the onset of a task stimulus the activity of predictive units in both mPFC (**Figure 3A**) and IPFC (**Figure 3B**) begins to ramp up, and the time course of this ramping up depends on whether task features in the current trial are the same as or different than features in the preceding trial. In IPFC, activity increases more rapidly for trials in which the target location is the same as the previous trial. This rapid increase is a consequence of IPFC's role in representing color information needed to identify the location of the target stimulus during a trial: when the location remains the same, lingering activity in the appropriate IPFC representations are rapidly re-activated. In contrast, mPFC activity is sensitive to changes both in the location of the target stimulus, as well as changes in the cued direction; both factors influence the development of mPFC activity. The sensitivity of mPFC in changes to both stimulus features is due in part to lingering activity within MPFC from the previous trial. As with the effect of location repetitions on IPFC activity, repetitions in the motion direction of the target stimulus re-activate mPFC representations that were already partially active from the previous trial. However, activity in mPFC is additionally modulated by activity in IPFC (Equation 5); changes in the location of the target stimulus that influence the evolution of IPFC activity thus have downstream effects on mPFC activity.

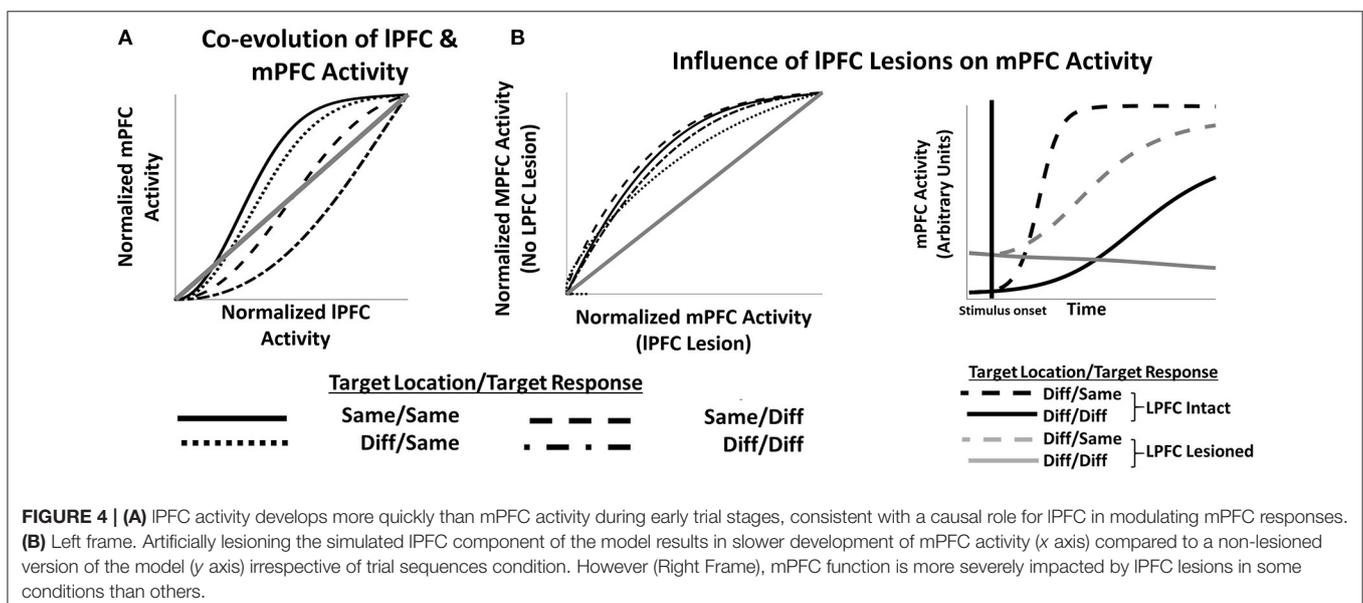
The sequence-dependent development of activity directly influences the speed at which a response is generated (**Figure 3C**). Responses in the model are ultimately generated by mPFC function (Equation 4): the longer activity in mPFC takes to reach a critical threshold, the slower the response time. Because the evolution of mPFC activity in the model depends on whether task features for the current trial are the same, the model reproduces standard trial sequence effects wherein feature repetition facilitates responding, while feature switches interfere with responses (Fecteau and Munoz, 2003). The HER model

further suggests why some feature switches may produce greater interference effects than others. Specifically, switches of the feature that most directly drives response, i.e., the motion direction of the stimulus, results in a greater RT difference than changes in the more abstract feature dimension, i.e., the color of the stimulus in this task.

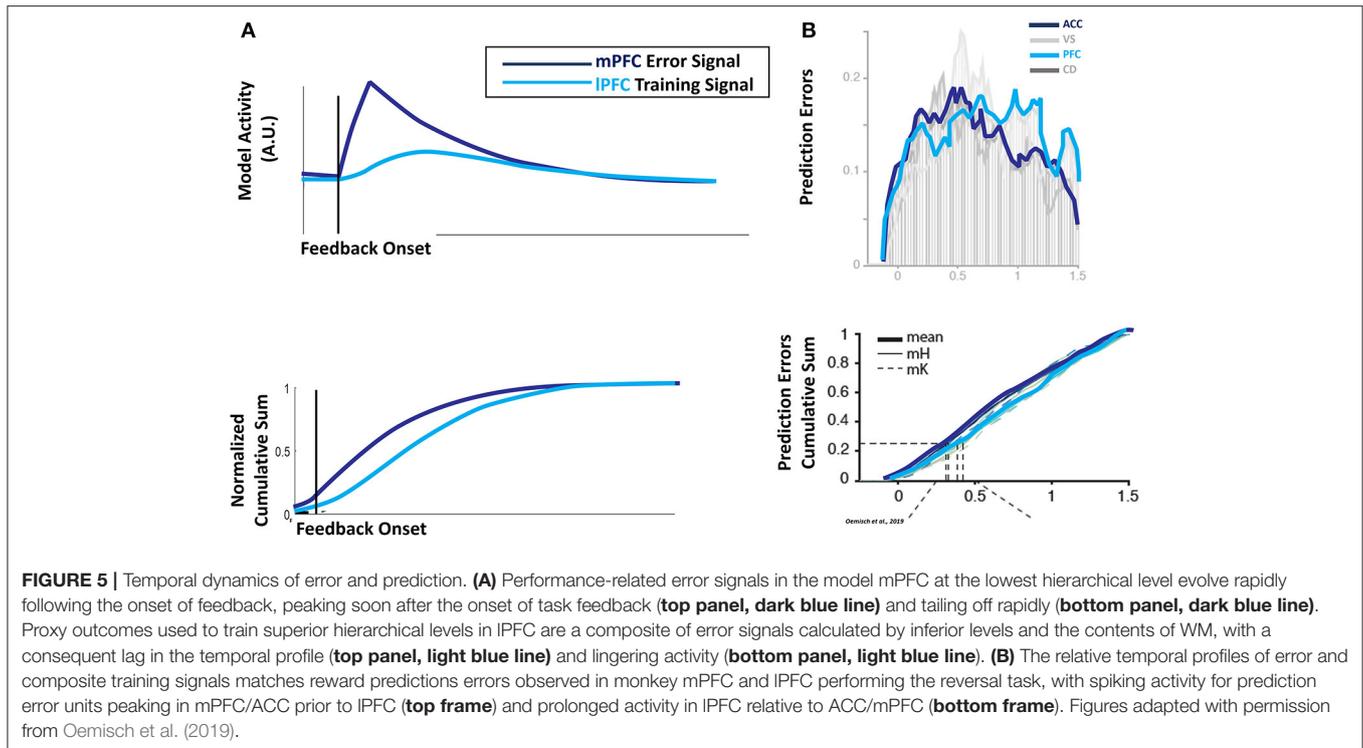
Furthermore, the model suggests how IPFC lesions might affect the development of mPFC activity (and consequent behavioral performance) on tasks that require integrating multiple features (**Figure 4**). The evolution of both LFPC and mPFC activity is affected by whether the abstract task variable (location) in the current trial is the same or different than in the previous trial (**Figure 4A**), and IPFC activity tends to lead mPFC activity, suggesting a causal role of IPFC activity on mPFC activity. This causal relationship can be directly queried by artificially lesioning IPFC in the network (e.g., by multiplying  $P'$  by zero in Equation 5). mPFC activity following lesions to IPFC (**Figure 4B**, left frame) is slower to evolve than when IPFC is intact. When considering only trials in which the location of the target stimulus has changed (**Figure 4B**, right frame), IPFC lesions more profoundly impact mPFC function when both target location and target response identity have changed relative to changes only in the target location.

## Feedback Processing

Following a response, the model receives feedback indicating whether the selected response was correct or incorrect. While equations used to compute errors in the model (Equation 6), like those used to calculate prediction unit activity, apply at every moment in the simulations, error-related activity is most prominent following feedback delivery, during which ongoing predictive activity is compared to an experienced outcome. During learning in the event-level model, comparison of feedback and concrete outcomes occurs only at the lowest hierarchical level; at superior hierarchical levels, "outcomes" are



**FIGURE 4 | (A)** IPFC activity develops more quickly than mPFC activity during early trial stages, consistent with a causal role for IPFC in modulating mPFC responses. **(B)** Left frame. Artificially lesioning the simulated IPFC component of the model results in slower development of mPFC activity (x axis) compared to a non-lesioned version of the model (y axis) irrespective of trial sequences condition. However (Right Frame), mPFC function is more severely impacted by IPFC lesions in some conditions than others.



derived from WM representations at the inferior level combined with the results of the feedback comparison process (i.e., the error signal, Equation 8). These “proxy” outcomes constitute a higher-order training signal that is composited from lower-level WM representations and error signals; the composition of the higher-order training signal is carried out by IPFC in the HER model, while the comparison of outcomes and predictions is undertaken by mPFC.

Naturally, since the proxy outcome depends on the lower-level error signal, the evolution of error units in mPFC in our simulations precedes the development of activity in IPFC training units (**Figure 5A**, top panel): the mPFC error signal ramps up rapidly at the onset of feedback, and decays quickly following feedback offset. In contrast, the IPFC training signal lags the mPFC error signal, and its activity is temporally blurred. Although the relative onset of the error and training signals is prefigured by the architecture of the HER model, the relative distribution (**Figure 5A**, bottom panel) of the signals emerges only due to the temporal dynamics introduced in these simulations. This emergent pattern qualitatively matches data recorded from monkey dACC and IPFC during performance of the reversal task (**Figure 5B**; Oemisch et al., 2019), and adds to the already considerable array of effects the HER model has been applied to (Alexander and Brown, 2015, 2018). Furthermore, the HER model provides a mechanistic explanation for both why prediction errors are observed in multiple brain regions during feedback processing, and for their relative time courses. Instead of multiple areas of the brain engaging in independent calculation of prediction errors, the HER model proposes that a single outcome prediction error calculation

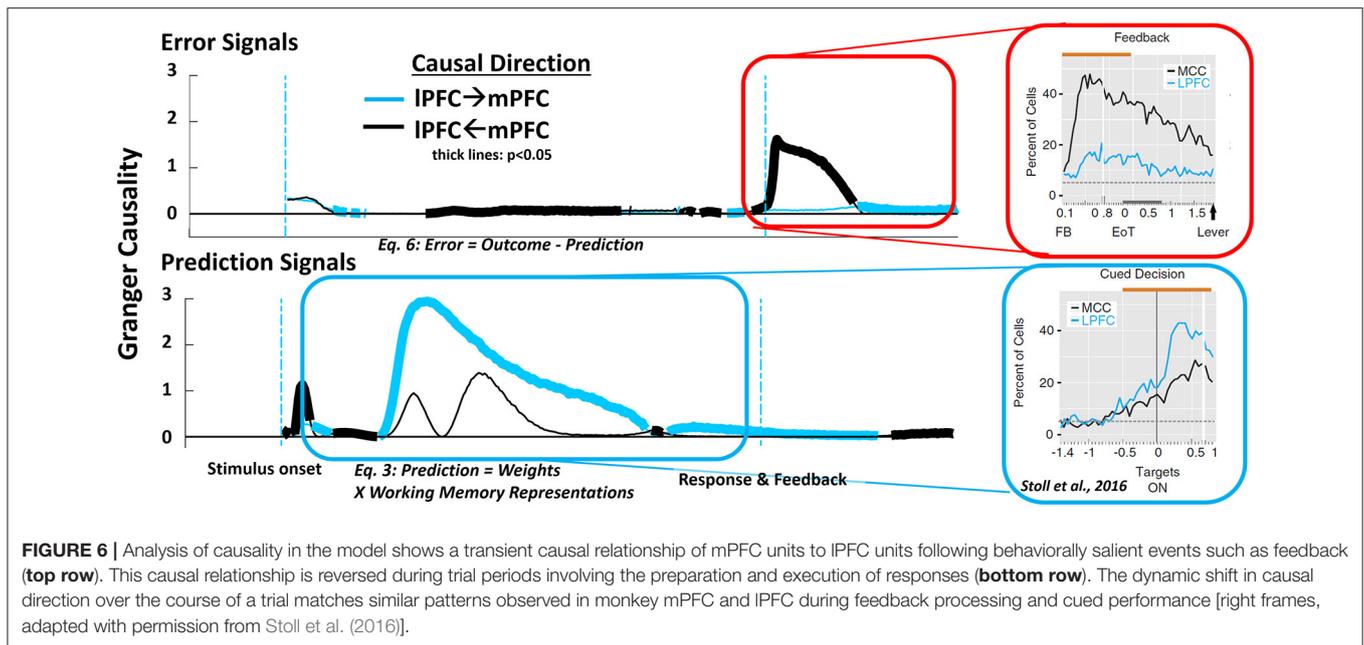
propagates through the cognitive control network in a bottom-up fashion.

## Information Flow

So far, our results are suggestive of how mPFC and IPFC interact during preparatory and feedback periods of a trial. To determine how these mPFC/IPFC interactions develop during a trial we turn to Granger causality (Luo et al., 2013) as a measure of how well one variable can be predicted by lagged values of another variable. Here, the variables are the unit activities in mPFC and IPFC and Granger causality indexes whether unit activity in one area is better predicted by the preceding unit activity from the other area than by its own past. Granger causality was computed for a time lag of 50 model cycles for both error-related and prediction units in the model, and trial timing was standardized—even if a response was indicated prior, the model was simulated for 300 cycles following the onset of a stimulus.

We first tested for causally significant error signals (Equation 6) and found they emerged prominently following performance-related feedback (**Figure 6**, top row). Immediately following the onset of feedback, Granger causality for the influence of mPFC on IPFC was significant at a level of  $p < 0.05$  (two sample  $t$ -test,  $df = 19$ , Bonferroni corrected), and remained so until feedback-related activity naturally decayed. IPFC also causally influenced mPFC during the feedback epoch of the task at a significance level of  $p < 0.05$  (two sample  $t$ -test,  $df = 19$ , Bonferroni corrected), but only after an outcome feedback had ended.

We next analyzed Granger causal effect of the prediction signal of the model (**Figure 6**, bottom row). We found



**FIGURE 6** | Analysis of causality in the model shows a transient causal relationship of mPFC units to IPFC units following behaviorally salient events such as feedback (**top row**). This causal relationship is reversed during trial periods involving the preparation and execution of responses (**bottom row**). The dynamic shift in causal direction over the course of a trial matches similar patterns observed in monkey mPFC and IPFC during feedback processing and cued performance [right frames, adapted with permission from Stoll et al. (2016)].

that prediction signals at the mPFC representation of the model immediately following stimulus onset Granger-cause activity in IPFC ( $p < 0.05$ ; two sample  $t$ -test,  $df = 19$ , Bonferroni corrected). While the HER model contains no mechanism by which predictions in mPFC influence processing activity in IPFC, Granger causality only indicates whether one signal can be predicted by previous values of another signal. In this case, predictive signals in mPFC and IPFC are correlated, but mPFC signals develop more rapidly than IPFC signals, producing a significant Granger causality effect. Following this transient effect, Granger causality for the influence of IPFC on mPFC becomes significant during the delay period prior to the generation of a response (**Figure 6**, bottom frame). This finding is consistent with the role of IPFC in maintaining information and implementing control demands. These simulation results are consistent with monkey neurophysiological studies demonstrating differences in the relative onset and peak activity of neurons in mPFC and IPFC that depend on the trial epoch (**Figure 6**, right panels; Stoll et al., 2016).

## DISCUSSION

In this manuscript, we have described additional simulations of the HER model in which real-time temporal dynamics were introduced to the model. The results of these simulations provide additional perspective on how the activity of mPFC and IPFC, as components of a hierarchical predictive coding framework (Alexander and Brown, 2018), might develop and interact following salient task events, and how the relative direction of this interaction evolves over the course of a trial. Beyond simply exploring the dynamics of model activity, our simulations demonstrate how the HER model can further account for

additional single-unit (Oemisch et al., 2019), behavioral (Wylie and Allport, 2000), and network effects (Stoll et al., 2016) previously reported in the literature.

At the level of neurons, the activity of single units in the HER model corresponding both to IPFC and mPFC, is observed to ramp-up following the onset of a task-relevant stimulus (**Figure 3**). Previous real-time models of mPFC (Alexander and Brown, 2011, 2014) have likewise shown units with ramping activity profiles, similar to those of reward- and error-predicting neurons in monkey mPFC (Amador et al., 2000; Amiez et al., 2006). The HER model, conceived as a temporally-coarse hierarchical extension to the PRO model (Alexander and Brown, 2015), was unable to replicate this pattern; by re-introducing real-time dynamics, and consequently reproducing effects from the PRO model that depended on temporal processes, our simulations underscore that the principal role and computational mechanisms attributed to mPFC by the PRO model remain intact in the HER model.

Furthermore, the simulations in this manuscript extend real-time processing to the registration of feedback and the development of error signals. The original PRO model (Alexander and Brown, 2011) was developed using temporal difference (TD) learning formulations (Sutton, 1988) in which the temporal profile of feedback signals (and subsequent error signals) was specified by the modeler (e.g., either as a punctate event or a box car profile). Here, the duration and magnitude of feedback signals is still modeler-defined, but the development of activity registering these signals is described by the same timing equations used to model all other unit activity in the model. By extending temporal dynamics to feedback processes, the simulations here are able to capture the temporal profile of error-signaling units recorded from monkey mPFC and IPFC, as well as the relative onset and decay of these signals (**Figure 5**).

Specifically, simulated error-related signals in mPFC peak earlier and decay more rapidly than signals observed in IPFC, consistent with recent reports (Shen et al., 2015; Oemisch et al., 2019). The HER model explains this through the role of mPFC in training error representations in IPFC: error signals generated directly by outcome feedback in the model are combined with active representations of task-stimuli to derive higher-order outcome and error signals (Equation 8), represented in IPFC. As the development of these higher-order signals is mechanically subsequent to direct error signals (Figure 1), the dynamics specified in Equation (11) dictate a later peak and lingering activity. Furthermore, the “proxy” outcome signals derived in IPFC are required for subsequent higher-order error calculations, suggested by the HER model to be carried out in hierarchically-superior regions of mPFC (Figure 4), and these error signals are subject to additional lag as lower-order error and training signal computations that support their calculation develop. The HER model thus provides a mechanistic explanation for the relative time course of error signals in caudal-to-rostral regions of mPFC (Polli et al., 2005).

Although computation of error (Equation 6) and prediction signals (Equation 3) is ongoing throughout our causality analysis over the entire course of a trial reveals that the net direction of information flow depends both on the type of information (prediction or error) computed, as well as the period within a trial. Following salient task events, such as stimulus onset or delivery of feedback, information in the model flows primarily from mPFC to IPFC, while during periods in which a salient event is expected but has yet to occur, information flows principally from IPFC to mPFC. This pattern maps well both to functional roles attributed to these regions, as well as the observed time course of interactions. Functionally, mPFC has long been associated with processing novel or behaviorally-relevant events, especially the occurrence of errors (Gehring et al., 1990) or otherwise surprising (Jessup et al., 2010; Ide et al., 2013) stimuli, while IPFC is implicated in slower processes involving information maintenance (Sawaguchi and Goldman-Rakic, 1991), representing task structure (Badre and D’Esposito, 2007), or implementing control in preparation for upcoming demands (Botvinick et al., 2001). This temporal dissociation, implied in the architecture of the HER model (Alexander and Brown, 2015) is made explicit in this manuscript, and the relative timing and flow of information in the model is consistent with human and monkey studies of PFC (Taren et al., 2011; Shen et al., 2015; Stoll et al., 2016; Oemisch et al., 2019).

The model distinguishes different levels of a computational hierarchy, which we associate with different brain areas to highlight the overarching functional contributions of a specific brain area. The association of the first (lowest) level of the hierarchy with mPFC and the next higher hierarchical level with IPFC should not imply that these areas are exclusively representing information pertaining to these hierarchical levels. Rather, task related information about the value of stimulus features and the type of prediction errors can be decoded from model units at all levels, albeit to different degrees.

This observation is consistent with a large body of literature demonstrating that prefrontal cortical neurons code task variables with mixed selectivity (Fusi et al., 2016), whereby individual neurons encode multiple task variables, i.e., they mix information about task rules, feature values, and outcome variables in their firing independent of the specific source of the information (Bernardi et al., 2020). According to these insights, mPFC and IPFC will host neurons tuned to similar task variables, but will show a “representational gradient” (Kyriazi et al., 2020) showing more prominent encoding of feedback information in mPFC and a more abstract stimulus value code in IPFC (Rigotti et al., 2013; Bernardi et al., 2020). Our model depiction is consistent with these insights and specifies the actual computations that might give rise to these representational gradients.

Finally, by introducing temporal dynamics, we were able to use the HER model to replicate sequential trial effects that are a staple of the cognitive control literature (Wylie and Allport, 2000; Fecteau and Munoz, 2003). Unsurprisingly, reaction times for the model are the most rapid for trials in which all features of the chosen stimulus are identical, and slowest for trials in which all features have changed relative to the previous trial. Of interest, however, are trials in which only one relevant feature (out of two in the current study) changes. In these cases, the identity of the changing stimulus can have differential effects on reaction time. The HER model solves structured tasks by decomposing stimulus dimensions hierarchically (Alexander and Brown, 2015): dimensions that are “concrete,” i.e., those that most directly inform the eventual response, are preferentially encoded at the lowest hierarchical level, while more abstract features are encoded at superior hierarchical levels. The simulations reported here suggest that changes in the concrete decision variable (in this case, the direction of the target response) may have a more profound influence on reaction time than changes in the more abstract variable. Recent work (Vassena et al., 2019) has begun to explore how interfering with the structure of a task through manipulations of presentation order might influence decision making and performance. The results of the present study suggest a complementary approach in which the differences in performance elicited through feature changes might be used to infer the representation of task structure.

In summary, the simulations of the extended HER model reported in this manuscript demonstrate that including temporal dynamics endows the model with additional explanatory power, and provides the basis for additional work investigating the function and interaction of regions within PFC, as well as how they contribute to behavior. More generally, the HER model, as an instance of predictive coding, suggests how additional regions in PFC may be organized (Alexander et al., 2017); specifically, the HER model is primarily concerned with how information is integrated in order to generate responses, but may have little to say about how information is acquired to begin with. It is possible that additional regions implicated in cognitive control may be integrated with the HER framework to describe information is actively selected and interpreted to assist in adaptive behavior.

## DATA AVAILABILITY STATEMENT

The original contributions presented in the study are included in the article/supplementary material, further inquiries can be directed to the corresponding author/s.

## AUTHOR CONTRIBUTIONS

WA and TW conceived of the study and wrote the manuscript. WA conducted computational simulations.

## REFERENCES

- Alexander, W., and Womelsdorf, T. (2018). Interactions of medial and lateral prefrontal cortex in hierarchical predictive coding. *bioRxiv*. doi: 10.1101/439927
- Alexander, W. H., and Brown, J. W. (2011). Medial prefrontal cortex as an action-outcome predictor. *Nat. Neurosci.* 14, 1338–1344. doi: 10.1038/nn.2921
- Alexander, W. H., and Brown, J. W. (2014). A general role for medial prefrontal cortex in event prediction. *Front. Comput. Neurosci.* 8:69. doi: 10.3389/fncom.2014.00069
- Alexander, W. H., and Brown, J. W. (2015). Hierarchical error representation: a computational model of anterior cingulate and dorsolateral prefrontal cortex. *Neural Comput.* 27, 2354–2410. doi: 10.1162/NECO\_a\_00779
- Alexander, W. H., and Brown, J. W. (2018). Frontal cortex function as derived from hierarchical predictive coding. *Sci. Rep.* 8, 3843. doi: 10.1038/s41598-018-21407-9
- Alexander, W. H., Vassena, E., Deraeve, J., and Langford, Z. D. (2017). Integrative modeling of prefrontal cortex. *J. Cogn. Neurosci.* 29, 1674–1683. doi: 10.1162/jocn\_a\_01138
- Amador, N., Schlag-Rey, M., and Schlag, J. (2000). Reward-predicting and reward-detecting neuronal activity in the primate supplementary eye field. *J. Neurophysiol.* 84, 2166–2170. doi: 10.1152/jn.2000.84.4.2166
- Amiez, C., Joseph, J. P., and Procyk, E. (2006). Reward encoding in the monkey anterior cingulate cortex. *Cereb. Cortex* 16, 1040–1055. doi: 10.1093/cercor/bhj046
- Badre, D., and D'Esposito, M. (2007). Functional magnetic resonance imaging evidence for a hierarchical organization of the prefrontal cortex. *J. Cogn. Neurosci.* 19, 2082–2099. doi: 10.1162/jocn.2007.19.12.2082
- Badre, D., and D'Esposito, M. (2009). Is the rostro-caudal axis of the frontal lobe hierarchical? *Nat. Rev. Neurosci.* 10, 659–669. doi: 10.1038/nrn2667
- Badre, D., and Nee, D. E. (2018). Frontal cortex and the hierarchical control of behavior. *Trends Cogn. Sci. (Regul. Ed.)* 22, 170–188. doi: 10.1016/j.tics.2017.11.005
- Barbas, H., and Pandya, D. N. (1989). Architecture and intrinsic connections of the prefrontal cortex in the rhesus monkey. *J. Comp. Neurol.* 286, 353–375. doi: 10.1002/cne.902860306
- Barbas, H., and Rempel-Clower, N. (1997). Cortical structure predicts the pattern of corticocortical connections. *Cereb. Cortex* 7, 635–646. doi: 10.1093/cercor/7.7.635
- Bernardi, S., Benna, M. K., Rigotti, M., Munuera, J., Fusi, S., and Salzman, C. D. (2020). The geometry of abstraction in the hippocampus and prefrontal cortex. *Cell* 183, 954–967.e21. doi: 10.1016/j.cell.2020.09.031
- Botvinick, M. M., Braver, T. S., Barch, D. M., Carter, C. S., and Cohen, J. C. (2001). Conflict monitoring and cognitive control. *Psychol. Rev.* 108, 624–652. doi: 10.1037/0033-295X.108.3.624
- Brown, J. W. (2009). Conflict effects without conflict in anterior cingulate cortex: multiple response effects and context specific representations. *Neuroimage* 47, 334–341. doi: 10.1016/j.neuroimage.2009.04.034
- Cole, M. W., and Schneider, W. (2007). The cognitive control network: integrated cortical regions with dissociable functions. *Neuroimage* 37, 343–360. doi: 10.1016/j.neuroimage.2007.03.071

All authors contributed to the article and approved the submitted version.

## FUNDING

The funding is supported by Florida Atlantic University.

## ACKNOWLEDGMENTS

This manuscript has been released as a pre-print at Alexander and Womelsdorf (2018).

- Fecteau, J. H., and Munoz, D. P. (2003). Exploring the consequences of the previous trial. *Nat. Rev. Neurosci.* 4, 435–443. doi: 10.1038/nrn1114
- Fusi, S., Miller, E. K., and Rigotti, M. (2016). Why neurons mix: high dimensionality for higher cognition. *Curr. Opin. Neurobiol.* 37, 66–74. doi: 10.1016/j.conb.2016.01.010
- Gehring, W. J., Coles, M. G. H., Meyer, D. E., and Donchin, E. (1990). The error-related negativity: an event-related potential accompanying errors. *Psychophysiology* 27, S34.
- Grossberg, S. (1988). Nonlinear neural networks: principles, mechanisms, and architectures. *Neural Netw.* 1, 17–61. doi: 10.1016/0893-6080(88)90021-4
- Hodgkin, A. (1969). Conduction of the nerve impulse. *Can. Med. Assoc. J.* 101, 359–361.
- Holroyd, C. B., and Yeung, N. (2012). Motivation of extended behaviors by anterior cingulate cortex. *Trends Cogn. Sci.* 16, 122–128. doi: 10.1016/j.tics.2011.12.008
- Ide, J. S., Shenoy, P., Yu, A. J., and Li, C. R. (2013). Bayesian prediction and evaluation in the anterior cingulate cortex. *J. Neurosci.* 33, 2039–2047. doi: 10.1523/JNEUROSCI.2201-12.2013
- Jessup, R. K., Busemeyer, J. R., and Brown, J. W. (2010). Error effects in anterior cingulate cortex reverse when error likelihood is high. *J. Neurosci.* 30, 3467–3472. doi: 10.1523/JNEUROSCI.4130-09.2010
- Kyriazi, P., Headley, D. B., and Paré D. (2020). Different multidimensional representations across the amygdalo-prefrontal network during an approach-avoidance task. *Neuron* 107, 717–730.e5. doi: 10.1016/j.neuron.2020.05.039
- Luo, Q., Lu, W., Cheng, W., Valdes-Sosa, P. A., Wen, X., Ding, M., et al. (2013). Spatio-temporal Granger causality: a new framework. *Neuroimage* 79, 241–263. doi: 10.1016/j.neuroimage.2013.04.091
- Nácher, V., Hassani, S. A., and Womelsdorf, T. (2019). Asymmetric effective connectivity between primate anterior cingulate and lateral prefrontal cortex revealed by electrical microstimulation. *Brain Struct. Funct.* 224, 779–793. doi: 10.1007/s00429-018-1806-y
- Niv, Y., and Langdon, A. (2016). Reinforcement learning with Marr. *Curr. Opin. Behav. Sci.* 11, 67–73. doi: 10.1016/j.cobeha.2016.04.005
- Oemisch, M., Westendorff, S., Azimi, M., Hassani, S. A., Ardid, S., Tiesinga, P., et al. (2019). Feature specific prediction errors and surprise across macaque fronto-striatal circuits during attention and learning. *Nat. Commun.* 10:176. doi: 10.1038/s41467-018-08184-9
- Polli, F. E., Barton, J. J. S., Cain, M. S., Thakkar, K. N., Rauch, S. L., and Manoach, D. S. (2005). Rostral and dorsal anterior cingulate cortex make dissociable contributions during antisaccade error commission. *Proc. Natl. Acad. Sci. U. S. A.* 102, 15700–15705. doi: 10.1073/pnas.0503657102
- Rigotti, M., Barak, O., Warden, M. R., Wang, X.-J., Daw, N. D., Miller, E. K., et al. (2013). The importance of mixed selectivity in complex cognitive tasks. *Nature* 497, 585–590. doi: 10.1038/nature12160
- Sawaguchi, T., and Goldman-Rakic, P. S. (1991). D1 dopamine receptors in prefrontal cortex: involvement in working memory. *Science* 251, 947–950. doi: 10.1126/science.1825731
- Shen, C., Ardid, S., Kaping, D., Westendorff, S., Everling, S., and Womelsdorf, T. (2015). Anterior cingulate cortex cells identify process-specific errors of attentional control prior to transient prefrontal-cingulate inhibition. *Cereb. Cortex* 25, 2213–2228. doi: 10.1093/cercor/bhu028

- Shenhav, A., Botvinick, M. M., and Cohen, J. D. (2013). The expected value of control: an integrative theory of anterior cingulate cortex function. *Neuron* 79, 217–240. doi: 10.1016/j.neuron.2013.07.007
- Stoll, F. M., Fontanier, V., and Procyk, E. (2016). Specific frontal neural dynamics contribute to decisions to check. *Nat. Commun.* 7, 11990. doi: 10.1038/ncomms11990
- Sutton, R. S. (1988). Learning to predict by the methods of temporal difference. *Mach. Learn.* 3, 9–44. doi: 10.1007/BF00115009
- Taren, A. A., Venkatraman, V., and Huettel, S. A. (2011). A parallel functional topography between medial and lateral prefrontal cortex: evidence and implications for cognitive control. *J. Neurosci.* 31, 5026–5031. doi: 10.1523/JNEUROSCI.5762-10.2011
- Vassena, E., Deraeve, J., and Alexander, W. H. (2019). Task-specific prioritization of reward and effort information: novel insights from behavior and computational modeling. *Cogn. Affect. Behav. Neurosci.* 19, 619–636. doi: 10.3758/s13415-018-00685-w
- Wilson, H. R., and Cowan, J. D. (1972). Excitatory and inhibitory interactions in localized populations of model neurons. *Biophys. J.* 12, 1–24. doi: 10.1016/S0006-3495(72)86068-5
- Wylie, G., and Allport, A. (2000). Task switching and the measurement of “switch costs.” *Psychol. Res.* 63, 212–233. doi: 10.1007/s004269900003
- Yeung, N., Botvinick, M. M., and Cohen, J. D. (2004). The neural basis of error detection : conflict monitoring and the error-related negativity. *Psychol. Rev.* 111, 931–959. doi: 10.1037/0033-295X.111.4.931

**Conflict of Interest:** The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

*Copyright © 2021 Alexander and Womelsdorf. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.*