



Published in final edited form as:

*J Cogn Neurosci.* 2011 November ; 23(11): 3598–3619. doi:10.1162/jocn\_a\_00047.

## From an Executive Network to Executive Control: A Computational Model of the *N*-back Task

Christopher H. Chatham<sup>1</sup>, Seth A. Herd<sup>1</sup>, Angela M. Brant<sup>1,2</sup>, Thomas E. Hazy<sup>1</sup>, Akira Miyake<sup>1</sup>, Randy O'Reilly<sup>1</sup>, and Naomi P. Friedman<sup>2</sup>

<sup>1</sup>Department of Psychology and Center for Neuroscience, University of Colorado, Boulder, Boulder CO, 80302, USA

<sup>2</sup>Institute for Behavioral Genetics, University of Colorado, Boulder, Boulder CO, 80302, USA

### Abstract

A paradigmatic test of executive control, the *n*-back task, is known to recruit a widely-distributed parietal, frontal and striatal “executive network,” and is thought to require an equally wide array of executive functions. The mapping of functions onto substrates in such a complex task presents a significant challenge to any theoretical framework for executive control. To address this challenge, we developed a biologically-constrained model of the *n*-back task that emergently develops the ability to appropriately gate, bind, and maintain information in working memory in the course of learning to perform the task. Furthermore, the model is sensitive to proactive interference in ways that match findings from neuroimaging, and shows a U-shaped performance curve after manipulation of prefrontal dopaminergic mechanisms similar to that observed in studies of genetic polymorphisms and pharmacological manipulations. Our model represents a formal computational link between anatomical, functional neuroimaging, genetic, behavioral, and theoretical levels of analysis in the study of executive control. In addition, the model specifies one way in which the prefrontal cortex, basal ganglia, parietal, and sensory cortices may learn to cooperate and give rise to executive control.

---

Goal-directed behaviors are enabled by executive functions that help stop prepotent responses, resolve interference, update working memory, shift mental sets, and coordinate multiple tasks (e.g., Friedman & Miyake, 2004; Logie, Cocchini, Della Sala, & Baddeley, 2004; Miyake et al., 2000; Salthouse, Atkinson, & Berish, 2003). Such broad categories of executive function can be fractionated into lower-level component processes. For example, working memory updating tasks require storing information, gating information into and out of working memory, tracking serial order, and selective attention. These processes may in turn be mapped to diverse parietal, frontal and striatal substrates (e.g., Wager & Smith, 2003), posing a many-to-many problem in mapping executive functions to their neural substrates. A paradigmatic example of this many-to-many mapping problem is the *n*-back task (e.g., Kirchner, 1958). The main purpose of this article is to elucidate the mechanistic basis of this complex task using a biologically constrained computational model.

### The *N*-back Task

In the *n*-back task, subjects identify over consecutive trials whether the current stimulus matches a stimulus presented *n* trials previously. At the cognitive level, this task is thought to involve numerous executive processes: active maintenance of the last *n* items; updating of new items so that they can be actively maintained; rapid binding of items to their serial order, so that responses are based on the match between the current item and the *n*-back item

---

and not between items matching at a non- $n$  lag; and resolution of any proactive interference arising from non- $n$  lag items. At the biological level, neuroimaging, pharmacological, and genetic polymorphism studies indicate that  $n$ -back performance is associated with a distributed network of parietal, frontal and striatal sites (Owen, McMillan, Laird & Bullmore, 2005; Tsuchida & Fellows, 2009; Olesen et al., 2003) and dopaminergic mechanisms (Aalto, Bruck, Laine, Nagren & Rinne, 2005; Apud & Weinberger, 2007; Goldberg et al., 2003; Egan et al., 2001; Mattay et al., 2003; Meyer-Lindenberg et al., 2005; Tan et al., 2007). This cognitive and neurobiological complexity makes the  $n$ -back task a useful test case for formal accounts of how executive functions arise from their neural substrates.

One feature of the  $n$ -back task makes it especially appropriate for this undertaking: It appears to require rapid binding of stimuli to representations of their serial order. Symbolic cognitive models (e.g., ACT-R) fulfill this requirement through the use of propositional representations and explicit variables, and have yielded a working  $n$ -back model (Juvina & Taatgen, 2007). However, the brain operates on the basis of distributed representations and slowly adapting synaptic connections. The difficulty in reconciling this distributed, slowly adapting neural substrate with the  $n$ -back's rapid binding requirements could explain the absence of more biologically constrained models of this task. Here we present a model that overcomes this challenge and is capable of learning the  $n$ -back without a mechanism specifically implemented for symbolic binding.

## Model Architecture

Our model is rooted in the biologically plausible prefrontal-basal ganglia working memory (PBWM) architecture (Hazy, Frank & O'Reilly, 2006, 2007, 2010; O'Reilly & Frank, 2006). PBWM's essential principle is that task-relevant information can be maintained in PFC and help guide successful task performance by a process of biased competition (Desimone & Duncan, 1995); the reward signals resulting from successful task performance, in the form of phasic dopamine, can then train the basal ganglia through reinforcement learning to send an "updating signal" for gating new information into the PFC. PBWM thus integrates numerous ideas from computational neuroscience, implementing reinforcement learning in terms of phasic dopamine via the Primary Value Learned Value (PVLV) mechanism (Hazy et al, 2010; O'Reilly et al, 2006), resolving the stability-flexibility dilemma (Goschke, 2000) with flexible gating mechanisms, and yielding biased competition via prefrontal representations stabilized through recurrent connectivity and tonic dopamine (e.g., Cohen, Braver & Brown, 2002). The architecture supporting these interactions is schematically illustrated in Figure 1.

Our implementation of PBWM, depicted in Figure 2, is most closely based on the PBWM model of the phonological loop (O'Reilly & Frank, 2006). As in that previous work, the model receives input from a layer in which one of 10 different units is activated on every trial of the task, each unit corresponding to a different stimulus. The network's response on every trial is indicated by the patterns of activation across two output layers: a "verbal output" layer with 10 units, corresponding to each of the input stimuli indicating the network's best guess as to the  $n$ -back stimulus, and a "manual output" layer with 2 units, corresponding to match and nonmatch responses, indicating the network's best guess as to whether the current stimulus matches the  $n$ -back stimulus. (Some  $n$ -back tasks require subjects only to indicate whether there is a match between the current and  $n$ -back stimulus, and not the actual identity of the  $n$ -back stimulus; we included both output requirements in our model because subjects are likely to keep the identity of the  $n$ -back stimulus identity in memory regardless of the precise variant of  $n$ -back they are performing.) Finally, a "posterior cortex" layer of 100 units is bidirectionally connected with each of these layers,

and provides a substrate for biased competition to take place. We refrain from identifying this layer with a particular neocortical area, as it contains no special mechanisms which might be thought to differentiate it from many areas of neocortex.

Superimposed on this structure are the core components of PBWM. These components include prefrontal layers organized into stripes, consistent with the functional macrocolumns observed in the monkey prefrontal cortex (e.g., Levitt, Lewis, Yoshioka & Lund, 1993; Pucak, Levitt, Lund & Lewis, 1996; Rao, Williams & Goldman-Rakic, 1999). The units constituting these stripes are unique relative to all other units in two ways: They are recurrently self-connected, and they contain an excitatory hysteresis current. When combined, these features enable persistent, self-sustaining patterns of activity. The resulting temporally-stable patterns of activity are gated on a stripe-specific basis by a set of corresponding stripes in a basal ganglia “matrix” layer, modeled after the medium spiny projection neurons of the striatal matrix. Each stripe in the matrix (stripes are represented as visible subgroups within the basal ganglia layers in Figure 2) contains “Go” and “NoGo” units that are spatially intermixed (as they are biologically). “Go” units correspond to the direct pathway of the striatum, and “NoGo” units correspond to the indirect pathway. As such, “Go” units have a disinhibitory effect on corticothalamic gating, thereby allowing working memory to be updated; “NoGo” units have an inhibitory effect on corticothalamic gating, thereby helping to keep the contents of working memory the same despite new incoming information.

## Learning Algorithms in the Model

Central to the PBWM architecture is the use of the PVLV algorithm, which can be seen as a biologically-plausible implementation of traditional temporal difference reinforcement learning (Hazy, Frank & O’Reilly, 2010; O’Reilly, Frank, Hazy & Watz, 2007). The PVLV algorithm is used specifically, and selectively, to train the “Go” and “NoGo” units of the striatum. Ultimately, PVLV trains Go units to fire in response to stimuli that predict reward (and which might therefore be updated into working memory) whereas NoGo units learn to fire when stimuli do not predict anything more rewarding than the information currently represented in working memory. In conjunction with the prefrontal layers, PBWM implements mechanisms that at a higher level of analysis can enable basic executive functions like active maintenance and gating (e.g., O’Reilly & Frank, 2006).<sup>2</sup>

The other components of the model are all trained with a standard Hebbian learning rule and an error-driven learning rule (O’Reilly & Munakata, 2000). The end result of this combination of learning rules and PVLV is that, by the end of training, networks learned to fire primarily “Go” units in certain stripes in the basal ganglia, such that the particular stripes activated depend on the activity patterns in other layers. This stripe-specific “Go” firing within the basal ganglia updates corresponding stripes in the PFC with information currently present in the input layer. Basal ganglia stripes that are not used for a given trial fire primarily “NoGo” units, resulting in the preserved maintenance of information from preceding trials. Finally, PFC activity representing this important maintained information biases the posterior layer, which in turn biases the verbal and manual output layers. These connection weights are incrementally refined via Hebbian and error-driven learning so that they are most likely to produce the correct verbal and manual outputs (see Appendices I & II for additional details).

---

<sup>2</sup>We note that the dense interconnectivity of the PBWM architecture is based on known neurobiology, and is therefore taken as a given in the current attempt to map from neurobiology to executive function. Previous work has identified that the intact striatal and prefrontal mechanisms of PBWM are necessary for good serial recall performance (O’Reilly & Frank, 2006), of which the *n*-back is a particularly demanding variant.

## Serial Order Representations of the Model

Interestingly, the model autonomously learns to take advantage of the stripe-specific gating possible within the PBWM architecture so as to solve the variable binding problem posed by the  $n$ -back task (Juvina & Taatgen, 2007). With training, the basal ganglia send an increasingly differentiated gating signal such that the PFC can learn to maintain items in different stripes conditional on their serial order. This increased specificity of gating enables distribution of the task's mnemonic demands across multiple stripes, and solves the rapid binding problem by allowing the model to autonomously bind representations of items to their serial order (e.g., O'Reilly, Busby & Soto, 2003).

One crucial addition to this standard PBWM architecture is the parietal layer, which represents the serial order of successive stimuli using a graded and compressive code (in which representations are distributed and increasingly similar to one another as the serial order of the current stimulus increases; depicted in Figure 2B for serial orders 1, 2 and 3, respectively). The localization of such a serial order representation to parietal cortex is consistent with previous models (Botvinick & Plaut, 2006; Botvinick & Watanabe, 2007), with electrophysiology and neuroimaging of serial order representation in the intraparietal sulcus (IPS; Nieder, Diester & Tudusciuc, 2006; Marshuetz, Reuter-Lorenz, Smith, Jonides & Noll, 2006; Marshuetz, Smith, Jonides, DeGutis & Chenevert, 2000), and with the IPS activity observed across an  $n$ -back meta-analysis (Owen, McMillan, Laird & Bullmore, 2005). Recent evidence suggests that working memory contents are encoded as a function of their ordinal position in the sequence of to-be-remembered items (van Dijck & Fias, in press), consistent with our use of a parietally-based serial order mechanism to satisfy the working memory updating demands of the  $n$ -back task. Thus, the serial order representations in our model are different from the representations expected to support processing of other attributes (e.g., color or shape), in that they are explicitly based on the known tuning curves of neurons coding for serial order in the IPS.

Importantly, we implemented serial order representations not as a continuous number line that stretches to the number of trials, but as a periodic repeat of item positions. For example, in the 2-back task, the serial order representations alternate between 1 and 2, whereas in the 3-back task, they repeatedly cycle through 1, 2, and 3. This periodicity of the serial order representations is imposed by fiat or "prescribed." Although we return to this issue in the discussion, it is a difficult and outstanding problem of how such serial order representations and their dynamics might be learned. We abstract over this difficulty here. This nonetheless leaves much to be solved: The model must still autonomously learn that these serial order representations are important, to bind them to the items presented on each trial, and to update and maintain this information appropriately.

## Organization of the Current Paper

The results of our simulations are outlined as follows. After describing the details of the model and the way in which the  $n$ -back task and its instructions were presented to the network, we demonstrate the capacity of the network to replicate hallmark findings from the  $n$ -back literature, spanning multiple levels of analysis (behavioral, hemodynamic, and genetic). We then quantitatively analyze the model's prefrontal and striatal functioning to support an expository description of the model's functioning. We next discuss how our simulations inform cognitive theorizing about how executive functions like active maintenance, gating, and the resolution of proactive interference may emerge from a highly interactive fronto-parieto-striatal circuit. Finally, we describe how our model provides not only a computationally explicit example of the prefrontal-parietal interactions commonly observed in the considerable neuroimaging literature on this and other executive control

tasks (Owen, McMillan, Laird & Bullmore, 2005; Tsuchida & Fellows, 2009; Olesen et al., 2003; de Frias et al., 2010; Egan et al., 2003), but also how it leads to new theoretical insights and untested predictions.

## METHODS

### Implementation

To illustrate the biologically-constrained nature of our model, here we briefly review the Leabra framework (O'Reilly, 2001). This framework simulates neural processing in terms of interconnected units, each of which has a membrane potential determined by separate excitatory, inhibitory and leak conductances. Fluctuations in the resulting membrane potential are thresholded and transformed to yield a rate-coded output that contributes to the excitatory conductance of all other units to which a particular unit is connected, in proportion to the connection weight. Connection and bias weights are initially randomized but are shaped over the course of training according to Hebbian, reward-driven, and biologically realistic error-driven learning rules (see below). Units are grouped into layers that undergo a k-winners-take-all function for simulating the influence of local inhibitory interneurons. These biologically-inspired mechanisms have been used in over 40 models to capture a variety of detailed phenomena (e.g., O'Reilly & Munakata, 2000), indicating that these simple biological mechanisms can yield human-like performance in a number of domains.

In addition to the PBWM implementation (see Appendix I, and depiction in Figure 2A), sequential order was represented via the scaled log-normal function (Botvinick & Watanabe, 2007):

$$R_{\rho}(r) = \exp\left(-\frac{(\ln r - \ln \rho)^2}{2\sigma^2}\right)$$

where  $R_{\rho}(r)$  is the activation level of the  $\rho^{\text{th}}$  unit in the layer on a trial with rank  $r$ , and  $\sigma$  is a parameter determining the relative specificity of each unit to its preferred rank. We elaborated on this scheme by convolving the activation levels specified by this scaled-log normal function with gaussian variance conforming to that empirically observed in the IPS (Nieder et al., 2006); this convolution is included here as an additional biological constraint on the parietal layer, and can be observed as “noise” in the activation dynamics depicted in the parietal layer of Figure 2B. In the model, this parietal layer is interconnected both with the prefrontal cortex and (strongly) with the basal ganglia, consistent with the known anatomy of humans and other primates (Fernandez-Miranda et al., 2008; Yeterian & Pandya, 1993) and with the functionally interconnected parieto-fronto-striatal network commonly observed in neuroimaging studies of the executive functions.

Each named layer of the model contains features which uniquely associate its layers with the identified brain regions. For example, prefrontal layers are unique due to recurrent connections and an excitatory hysteresis current, as well as the stripe organization connected with a parallel stripe organization in striatal layers; parietal layers are unique due to the graded and compressive activation dynamics imposed there; striatal layers are unique due to their dopamine-driven reinforcement learning. The posterior layer is distinct because it contains none of the unique features above, but only the more general mechanisms implemented by Leabra and thought to apply to neocortex in general. Moreover, the connectivity among these layers is based on known neurobiology (Hazy, Frank & O'Reilly, 2006, 2007, 2010; O'Reilly & Frank, 2006).

## Training & Testing

All models were run in batches of 25 networks, and each network was initialized with random patterns of connection weights. To compare performance on the 2- and 3-back tasks, we employed networks with 12 PFC stripes so that the same networks were capable of learning both tasks, as the 3-back task seemed to require more working memory “capacity” than the 2-back task. For all other analyses, we used a scaled-down model consisting of only 6 stripes, both to speed training time and make detailed analyses of network behavior more tractable.

Training on the 2- and 3-back tasks consisted of activating one of 10 possible input units and the corresponding distributed representation of serial order in the parietal layer (each trial corresponds to one of the three serial orders illustrated in Figure 2B). “Lure” trials, in which the current stimulus matched a previous stimulus at a non- $n$  lag, were allowed to occur. “Recent” lure trials are those where the current stimulus matches the  $n-1$  stimulus; “Non-recent” lure trials are those where the current stimulus matches a preceding stimulus with a lag larger than  $n$ .

Human subjects are instructed on the value of  $n$  for each  $n$ -back task they perform. In our simulations, the network was informed of the value of  $n$  by way of a small, probabilistic bias to replace stimuli occurring at values of  $n$ . This bias was implemented in the 2-back task by increasing the activity level of the Go units in the matrix layer on a random 10% of the trials in which they had not fired on the previous trial. Similarly, in the 3-back task, the activity level of those units was increased on a random 10% of the trials in which they had not fired on the previous two trials. This probabilistic bias yields a proportion of trials in which the prefrontal cortex is updated with a periodicity of  $n$ . For this bias to yield good performance, the network must not only perform correctly on the few trials where this probabilistic updating occurs, but must also generalize that behavior across all trials and stimuli.

For testing, the patterns of activity in the verbal and manual output layers were recorded after those activity patterns had stabilized or a maximum number of cycles had occurred (here we use the Leabra default of 60 cycles). The most active output unit was considered the network’s response, and this output was compared with the correct output for computing the error statistics described in Results. Networks were trained in epochs of 500 trials each until the network was tested to perform above 80% correct in terms of both its verbal and manual outputs for 7 consecutive epochs. This performance criterion allows networks to develop individual differences in the range of those observed in humans: Some networks will perform substantially better than 80% correct by the end of training, whereas others may have a shallower learning curve. For all analyses except those pertaining to learning across the entire course of training, network behavior is tested during the final 10% of training.

For individual differences analyses, three batches of 25 networks were run with variations in the gain of prefrontal units (a proxy for tonic prefrontal dopamine) but the same 25 random seeds were used to initialize weights across each batch to ensure comparability across model runs. Generalization was assessed in terms of the verbal responses in a distinct batch of 25 networks on a randomly selected set of 10 trial sequences; these 10 trial sequences had been entirely omitted from the training set. For example, the sequence A1X2B1 might have been excluded from the training set for the 2-back network, where the intervening “X” stimulus could have been any of the possible stimuli.

The activation dynamic resulting from training is schematically illustrated in Figure 3 for the 2-back task. The first trial is a nonmatch trial with input stimulus “A” and serial order “1” (in Results, this type of trial is represented with the phrase “A1”). A subset of basal ganglia

stripes fire (the leftmost three basal ganglia units in Figure 3), resulting in maintenance of stimulus “A” within a corresponding subset of PFC stripes (the left-most three PFC units in Figure 3). On the following trial a different subset of basal ganglia stripes fire, resulting in the maintenance of the next stimulus (“D”) within the corresponding new subset of PFC stripes. This two-part activation dynamic repeats across all subsequent trials, but is illustrated for several trials in Figure 3 for clarity, including a recent lure trial, a non-recent lure trial, and a match trial. Three-part activation dynamics emerge in networks trained to perform the 3-back task, such that PFC, parietal, and basal ganglia layers have three distinct activation states (as opposed to the two distinct states illustrated in Figure 3). The only remaining difference in 3-back is that the correct verbal and manual outputs correspond to matches between the item presented currently and that presented 3 trials previously in the 3-back.

## RESULTS AND DISCUSSION

### *The Model Captures Benchmark Findings in the N-back Literature*

Our model was capable of *de novo* learning of both the 2-back and 3-back tasks, without an underlying symbolic variable system for performing rapid binding. This learning was not rote, in that all networks generalized to untrained sequences at a rate significantly above chance  $t(1,24)=17.9, p<.001$  for 2-back, and  $t(1,24)=12.9, p<.005$  for 3-back. As described below, the model also captured numerous benchmark features of human performance in the *n*-back task.

One hallmark finding in the *n*-back literature is reduced accuracy as *n* increases from 2 to 3. The model also showed this pattern, such that 2-back accuracy was higher than 3-back accuracy,  $F(1,24)=10.54, p<.005$ , as shown in Figure 4A. This result arises from two features of the *n*-back: Relative to 2-back, 3-back requires an additional item be maintained by the prefrontal layers; also, 3-back involves a less reliable signal of a current item’s serial order, owing to the logarithmic compression of the parietal layer. These two constraints jointly produce lower performance on (and also slower learning of) the 3-back task, because they diminish the ability of the network to appropriately bind an item to its serial order, and to maintain this binding over subsequent trials.

A second benchmark finding in the *n*-back literature is that performance is sensitive to the presence of lures – items that match a preceding item but not at the critical *n* lag. The model also captures this phenomenon, such that accuracy was significantly lower for recent lures than non-recent lures (Figure 4B) in both the 2-back,  $F(1,24)=77.2, p<.001$ , and the 3-back,  $F(1,24)=15.8, p=.001$ . This effect reflects interference caused by items in the input which match items maintained in memory, albeit with a different temporal order, thereby yielding a tendency for the network to inappropriately detect a match on lure trials. Moreover, accuracy is particularly low on recent lure trials (*n*-1), reflecting proactive interference, because the prefrontal layers are more likely to represent items with lags less than *n* than items with lags greater than *n* (the latter are task-irrelevant); thus the network is more prone to erroneously detect matches in the former case.

One counterintuitive result from the *n*-back literature is that the effect of *n*-1 lures, relative to the effect of non-recent lures (i.e., lures at positions  $> n$ ), is reduced as *n* moves from 2 to 3-back (Oberauer, 2005). Although this effect is counterintuitive – one might expect that the cost of lure trials on accuracy would increase proportionally with overall difficulty – the model reproduced the observed result (Figure 4B;  $F(1,24)=18.14, p<.001$ ). Consistent with the model’s functioning, this effect reflects the fact that proactive interference arising from a match between the current item and maintained items is *diluted* when more items are being simultaneously maintained, as in the 3-back task.

Neuroimaging studies of this kind of proactive interference reveal a larger hemodynamic response in the lateral prefrontal cortex to recent relative to nonrecent lures (Jonides & Nee, 2006; Badre & Wagner, 2005; Jonides et al 1998). The hemodynamic response is thought to reflect metabolic demands; furthermore, 50-80% of the brain's energy consumption reflects the input and output activity of its neurons (Buzsáki, Kaila & Raichle, 2007). As an approximation of this metabolic demand, we calculated a proxy hemodynamic response by summing the net input to each unit in the PFC, with each unit's contribution to the sum weighted by its net output. Consistent with extant neuroimaging data on proactive interference, our simulated hemodynamic response was markedly increased in prefrontal layers during recent lures, relative to nonrecent lures or targets,  $t(24)=11.35$ ,  $p<.0005$  and  $t(24)=5.01$ ,  $p<.001$ , respectively; see Figure 5. This result was not due solely to simulated excitatory neurotransmission: The same pattern was observed in terms of net inhibitory input (see Appendix I for details about inhibitory currents in Leabra), consistent with theories of inhibitory contributions to the hemodynamic response (Buzsáki, Kaila & Raichle, 2007), and with the involvement of inhibition in resolving proactive interference (Jonides et al., 1998).

### The Model Captures Individual Differences in Human N-back Performance

In addition to capturing the above hallmark phenomena in the *n*-back task, we also tested whether the model captures individual differences. One source of individual differences is genetic variation related to dopaminergic functioning, such as the Val158Met polymorphism in the gene coding for catechol-O-methyl transferase (COMT), the principal enzyme that degrades dopamine in the prefrontal cortex (Boulton & Eisenhofer, 1998). The low efficiency variant (the met allele) yields a large net reduction in prefrontal dopamine metabolism relative to carriers of the higher efficiency val allele (Chen et al., 2004; Mannisto & Kaakkola, 1999). This differing efficacy results in a higher tonic level of prefrontal dopamine in met carriers (Bilder, Volavka, Lachman & Anthony, 2004).

Consistent with the hypothesized inverted U-shaped curve relating prefrontal dopamine levels to executive control (see Mattay et al, 2003), homozygotes for the val allele perform worse on the *n*-back than met carriers, either in terms of performance (e.g., Goldberg et al., 2003) or efficiency (i.e., neural activation required to achieve the same level of performance; e.g., Egan et al., 2001). Additionally, met carriers perform worse following pharmacological manipulations thought to increase prefrontal dopamine levels, such as administration of amphetamine (Mattay, et al, 2003). Some recent studies suggest that the effect of the Val158Met polymorphism on *n*-back performance is weak, if it exists at all (e.g. Barnett et al, 2008). However in practice any conclusion about the influence of COMT polymorphisms is complicated by other unmeasured and confounding genetic differences that may also distinguish val and met carriers (e.g., linkage disequilibrium or functional epistasis; Meyer-Lindenberg et al., 2006; Tan et al., 2007). Biologically-constrained computational modeling can offer clarity to this situation, as a way of testing the underlying hypothesis that extremes in prefrontal dopamine should be associated with worse performance when all other factors are held constant.

Higher extracellular dopamine levels are frequently thought to increase the gain in individual pyramidal cells' activation function so as to make strongly active cells more active – an excitatory effect – and weakly active cells less active – an inhibitory effect (Cohen, Braver & Brown, 2002). The net result is an increase in signal-to-noise ratio for PFC as a whole (Durstewitz, Seamans & Sejnowski, 2000; Stefanis, van Os, Avramopoulos, Smyrnis, Evdokimidis & Stefanis, 2005; Winterer et al., 2006). In this way, individual differences at the Val158Met locus of the COMT gene might be hypothesized to produce differences in the relative sharpness of active representations in prefrontal cortex. All things being equal, sharper, sparser representations will promote faster processing and more robust

maintenance in PFC areas (O'Reilly & Munakata, 2000). Thus, to mimic the putative effects of individual differences in COMT function, we trained models to perform the 2-back task under variations in tonic dopamine's aforementioned (and most widely-hypothesized) influence on prefrontal cortex: signal-to-noise ratio (Durstewitz, Seamans & Sejnowski, 2000; Stefanis, van Os, Avramopoulos, Smyrnis, Evdokimidis & Stefanis, 2005; Winterer et al, 2006; Cohen, Braver & Brown, 2002). Specifically, we increased the gain of the sigmoidal activation function on the units in the prefrontal layers from the default value (from 400 to 600). The gain was also decreased from the default value as a proxy for reduced levels of prefrontal dopamine (from 400 to 100).

Results of the simulations indicate that while none of these variations in prefrontal dopamine precluded learning of the 2-back task to criterion, the final levels of performance reached by these networks after training conformed to the expected U-shaped curve,  $F(1,24)=5.25$ ,  $p=.027$  and  $F(1,24)=9.09$ ,  $p=.006$  for target and lure trial accuracy, respectively; see Figure 6. In our model, the U-shaped curve arises by rebalancing the flexibility-stability tradeoff: With low gain/tonic dopamine, prefrontal representations are somewhat unstable, but with high gain/tonic dopamine, prefrontal representations become somewhat difficult to update. Note that the individual differences resulting from PFC gain are not unique to *n*-back or PBWM; other models positing similar dopamine effects in PFC may exhibit similar results (e.g., Tagamets & Horwitz, 2000; Chadderdon & Sporns, 2006; Deco, 2006), and as such, this result represents an important point of convergence across multiple formalisms.

Another source of individual differences in the *n*-back relates to the influence of control strategies and response bias on behavioral performance. Juvina and Taatgen (2007) showed that subjects encouraged to use a high-control strategy in this task – i.e., to rely on active maintenance as opposed mere familiarity – show a positive correlation between accuracy on recent lure trials and on target *n*-back trials. In contrast, subjects encouraged to use a low-control familiarity strategy demonstrate a negative correlation between these trial types. Because our model includes only the mechanisms thought to be involved in high-control strategies, the model should also show this positive correlation. Indeed, we observed a robust positive correlation,  $r(73)=-.44$ ,  $p<.0005$ , between lure and target accuracy in the 2-back task (Figure 7) performed by models that varied in their PFC gain parameters (as described in the previous paragraph). This positive correlation arises because networks differing in prefrontal gain consequently also differ in their ability to update and maintain information in working memory – abilities which support performance on both lure and target trials alike<sup>3</sup>.

### Inside The N-back: How Gating, Binding, and Resolution of Proactive Interference Occur

As described above, our model captures numerous empirical phenomena from the *n*-back task. Crucially, this good match to empirical data is enabled not by the explicit fitting of parameters, but rather by the types of representations that develop through learning in PBWM. These representations can be readily understood as instantiations of the very executive functions hypothesized to be crucial for *n*-back performance: the need to flexibly update working memory, to bind stimulus representations to representations of serial order, and to manage proactive interference. Below, we demonstrate how these functions are accomplished using quantitative analysis of the model's learning trajectories.

<sup>3</sup>We suggest this effect is intrinsic to the model's emergent behavior, and not merely epiphenomenal, for two reasons. First, a more likely result would have been a negative correlation between the accuracy on lure and target trials, owing to the fact that Leabra involves the learning of "bias weights" which might produce the widely-observed tradeoff between hit and false alarm rate in target detection tasks. Second, this correlation was specific to networks in the trained state; no significant correlation between lure and target accuracy was observed following the first epoch of training ( $r=-.10$ , ns).

**Gating**—One principal executive function important for the  $n$ -back task is working memory updating; our model reveals what form this updating may take as a result of the striatal reinforcement learning mechanisms implemented in our model. In particular, the striatal layers learn to maximize reinforcement by firing differentially in terms of the serial order of each stimulus (i.e., 1, 2 or 3) instead of stimulus identity (i.e., A, B, C, etc). This policy develops because it is supported by network connectivity (such that parietal layers project particularly strongly to striatal layers) but also because it maximizes reinforcement. Had striatal layers learn to fire differentially based on stimulus identity information (i.e., A, B, C, etc), then it would be up to prefrontal layers to learn to represent whether *each* of those stimuli had been seen 1, 2, or 3 trials ago or *not at all*, as would most commonly be the case. Because it is less efficient to use limited prefrontal resources to represent stimuli that have *not* been recently experienced than to specifically represent those stimuli that have been seen recently, the latter updating policy is what emerges naturally through reinforcement learning.

The order-based gating striatal policy can be seen in how the activity patterns of these layers become more discrete with respect to serial order as training progresses. Formally, this change can be quantified as a reduction in entropy (such that lower entropy reflects greater certainty in which striatal units will be activated by a particular serial order) over the course of learning, as illustrated in Figure 8A. Thus, increasingly distinguishable neural patterns in the basal ganglia occur for distinct serial orders as training progresses, thereby yielding an order-specific gating signal.

The importance of this reduction in entropy can be seen in its relationship to performance. Although all networks ultimately reached approximately the same level of updating ability (i.e., near-zero entropy by the end of training), differences in accuracy on the task at that final point could be predicted based on the *history* of striatal entropy. That is, networks that were less error-prone at the end of training showed no difference in striatal entropy at that late point, but rather lower striatal entropy only *early* in training (as illustrated in Figure 8B). This effect occurs because the separation of items occurring with different serial orders to different prefrontal stripes is essential for two subsequent developments: the differentiation of items by prefrontal stripes, and the active maintenance of this information to resolve proactive interference. Networks that achieve earlier reductions in striatal entropy have a “head start” in these subsequent and slow refinements, each discussed in turn below.

**Binding**—A stable, order-specific gating signal is a prerequisite for prefrontal units to learn to differentiate items occurring with different serial orders – i.e., binding, an important process for  $n$ -back performance (Oberauer, 2005; Badre & Wagner, 2006). The binding that occurs in our model is distinct from that typically occurring in connectionist models, which relies on coactivation of shared features. Our model binds items to serial order in a fundamentally different way, as described below.

First, the order-specific gating policy developed by striatal reinforcement learning mechanisms exposes particular PFC stripes to stimuli occurring with a particular serial order, and other PFC stripes to stimuli occurring with other serial orders. By itself, this order-specific gating policy does not suffice for binding; the network must also be able to differentiate between the stimuli of any given serial order (e.g., to differentiate an A of serial order 1 from a B of serial order 1). Because this stimulus identity information is not provided by firing in striatal layers (which is serial-order based), the prefrontal layers must discriminate stimulus identity on the basis of posterior representations. Ultimately, this discrimination is accomplished through a process of representational differentiation supported by Hebbian and error-driven learning<sup>4</sup>, such that each prefrontal stripe will learn to discriminate all items occurring with the serial order that that stripe is preferentially

exposed to (via the order-based striatal gating policy). This progressive differentiation both within and across stripes in the PFC reflects the network emergently learning to *bind* items to their serial order.

The nature of the resulting bound representations can be quantified in terms of the Euclidean distance between prefrontal activity patterns across the various items and order combinations. This high-dimensional analysis can be illustrated with a cluster plot in which the y-axis represents item by order combinations, horizontal lines represent the Euclidean distances between clusters, and cluster membership is indicated by vertical lines. We performed this type of hierarchical cluster analysis on the activity patterns, both before and after training, of one PFC stripe that learned to code for items appearing with one serial order (Figure 9A & 9B) and for a different PFC stripe that learned to code for items appearing with a different serial order (Figure 9C & 9D). The resulting figure reveals an initially haphazard pattern of representational similarity across items by order combinations (represented as a letter followed by a number; e.g., “D1”; left panels of Figure 9). After learning, this disorganization resolves into a highly structured representational scheme (right panels of Figure 9) in which all items occurring with a non-preferred serial order for a given PFC stripe are highly similar, as indicated by very short horizontal lines linking the items into large clusters. In contrast, items of a preferred serial order become much more differentiated, as indicated by the increasingly pairwise clusters.

Thus, reinforcement learning mechanisms drive an order-based gating policy, while Hebbian/error-driven mechanisms support representational differentiation within particular serial orders. These processes jointly give rise to the active maintenance of item-context bindings in our model, such that items are bound to their context in terms of which PFC stripe they are gated into. Our model further suggests that this binding may occur through a sensitivity of the prefrontal cortex to serial order; indeed, empirical evidence suggests the prefrontal cortex encodes information about serial order (Amiez & Petrides, 2007), and our model demonstrates how such sensitivity might emerge.

**Proactive Interference**—As the prefrontal-striatal circuit learns to gate and actively maintain information about items and their serial order, the network must also learn to resolve proactive interference from lure trials. In essence, the network must produce responses based on the match or mismatch between the current stimulus and information that was updated  $n$ -trials previously, while avoiding responses that would be based on any mapping between the current stimulus and the information updated and maintained from non- $n$  lag lure trials. This constraint is at its core a *selection* problem: The prefrontal stripe with stimulus identity information for the current trial’s serial order – and not stripes with information from other serial orders – must convey this information to the verbal output and posterior cortical layers so that corresponding match/non-match outputs can be activated.

The network learns to solve this selection problem through two mechanisms that emerge over learning. First, the stimulus identity information relevant to the current trial’s serial order biases the verbal output unit that corresponds to that stimulus’s identity, as learned in the weights connecting that prefrontal stripe and the verbal output layer. Second, the posterior cortex acts as a kind of comparator, such that error-driven and Hebbian learning mechanisms craft a set of weights in the posterior cortical layer to detect matches between the stimulus input and verbal output layers and activate the appropriate manual output (see Appendix II for more details). The source of lure errors is therefore multicausal: Some errors (approximately 25%) reflect inappropriate detection of input/output matches by the posterior cortical area (i.e., the verbal output is correct and does not match the current stimulus, but

---

<sup>4</sup>Appendix II describes how error-driven and Hebbian learning cooperate to support representational differentiation.

the manual output nevertheless indicates a match response). Other errors (approximately 75%) reflect item confusion in the prefrontal layers as a result of interference from current stimuli, ultimately leading to the biasing of the incorrect unit in the verbal output layer (Appendix II provides a detailed analysis of recent lure errors in the 2-back task, which provides no evidence for the interpretation that lure errors arise because of incorrect gating on previous trials).

Figure 10 illustrates that this selection problem is solved relatively slowly over the course of training, with more rapid reductions in the proportion of errors that occur on non-recent lures and match trials than on recent lure trials, as well as a lower asymptotic error rate on those trials. We thus observed a relatively protracted development of resistance to interference, which is consistent with new evidence on developmental trajectories in the  $n$ -back (Schleepe & Jonkman, 2010). Our model shows this protracted development due to an interdependency between the resolution of proactive interference and other executive functions: Gating, maintenance, and binding control processes, supported by reinforcement learning (in the case of gating) as well as Hebbian and error-driven learning, must first construct a relatively stable state before those representations can be incrementally refined to reduce proactive interference through additional Hebbian and error-driven learning.

## GENERAL DISCUSSION

Here we report a biologically based model of the parieto-fronto-striatal system that learns to perform the  $n$ -back task, emergently producing representations that support executive functions like gating, binding and the resolution of proactive interference. The model's acquisition of these functions enables its close match to empirical data, including behavioral, genetic, and neuroimaging findings, without the need for fine-grained tuning of the model's underlying parameters. Specifically, the updating of working memory is accomplished as the basal ganglia learn to provide a gating signal that is increasingly differentiated by an item's serial order, quantified above in terms of entropy. Active maintenance occurs as the prefrontal cortex learns to bind items and their serial order, quantified above via hierarchical cluster analysis. Finally, proactive interference resulting from recent lure trials affects these prefrontal representations via increases in net input, unit activation, and inhibitory neurotransmission, consistent with the increased BOLD response observed in the PFC during proactive interference (Jonides et al., 1998; Badre & Wagner, 2005). The model also captures the effect of individual differences in prefrontal tonic dopamine, individual differences observed when humans use the same type of high-control strategy implemented by our model, as well as decreased overall accuracy and an increase in the relative accuracy of recent lures as  $n$  moves from 2 to 3.

Our model integrates previous work to use an order representation in PBWM (as in the phonological loop model developed by O'Reilly & Frank, 2006) by more firmly rooting it in the biology of the IPS (as in Botvinick & Watanabe, 2007) and thereby differentiating the serial order signal from representations that might be used for dimensions like color or form. This work extends the phonological loop model in two important ways. First, the  $n$ -back task differs from the serial recall performed by these models, in that serial order representations must now do "double-duty" – simultaneously supporting the recall of old information as well as the storage of new information. Second, our model addresses the lack of an explicit external frame of reference for the representation of serial order by positing an internally generated one, such that there is a *periodicity* relation rather than a strictly serial relation between successive stimuli. While the importance of a self-imposed periodicity function for our model actually leads to testable predictions, how these representations might autonomously develop nonetheless remains an important and unsolved problem.

## Insights

Our model also leads to a number of theoretical insights. First, previous literature suggests that executive functions show unity and diversity at both behavioral and genetic levels (e.g., Friedman et al., 2008), but our model may challenge modular interpretations of this unity and diversity. Executive functions emerge here from an integrated parieto-fronto-striatal circuit instead of discrete mechanisms involved in only some executive tasks (c.f. Cooper & Davelaar, 2010). Instead, a more emergent view of unity and diversity may enable a better match to neural mechanisms. Our model also indicates this emergent view will need to include frontal, striatal, and parietal areas, at the minimum. Executive functions are often discussed in terms of frontal, parieto-frontal (e.g., Corbetta & Shulman, 2008), or fronto-striatal (O'Reilly & Frank, 2006) substrates, but theoretical accounts of frontal, parieto-frontal or fronto-striatal interactions may be substantially incomplete without considering all three parts to the larger, integrated network. For example, though executive control might be considered relatively distinct from abilities like serial order processing, our model indicates the neural mechanisms supporting behavior across these domains may be related; this relationship should be explicitly considered in determining the role of parietal cortex in updating tasks.

However, it is also likely that other areas of parietal cortex contribute to performance in ways that do not selectively relate to time or order representations (Collette et al., 2005), and our model does not encapsulate the only form of parieto-frontal interaction. These anatomically and functionally diverse regions (e.g., Rushworth, Behrens, & Johansen-Berg, 2006) are likely to support multiple forms of processing. Thus, other computational and theoretical models of fronto-parietal function (Edin, et al., 2009; Corbetta & Shulman, 2002; Corbetta, Patel & Shulman, 2008) may describe aspects of the parietal cortex not captured by our model, and vice versa.

The model also provides insight into the ability to resolve proactive interference. Our simulations suggest that the increased hemodynamic response observed during recent lures may *reflect* the presence of interference, and not a separate control process recruited to *resolve* interference. Instead, proactive interference resolution unfolds as an emergent consequence of the network's learning in general. That is, proactive interference resolution depends on striatal gating becoming increasingly specific to serial orders (which reduces interference across different serial orders that involved presentation of the same item), and increased representational differentiation among items of the same serial order within PFC (which reduces interference across different items presented with the same serial order). However, proxy-hemodynamic increases to recent lures also reflect that a similar preceding item is being strongly maintained and that the current stimulus is being fully processed. These multiple facets of the hemodynamic response to proactive interference may explain the seemingly paradoxical findings that activation of the lateral prefrontal cortex positively correlates with fluid intelligence (Gray, Chabris & Braver, 2003), which presumably relies on strong maintenance and full processing of stimuli, but also positively correlates with behavioral indices of proactive interference (Nee, Jonides & Berman, 2007). Our model thus offers one explanation for these apparent contradictions in the current empirical literature.

## Predictions and Extensions

Our *n*-back model also leads to new testable predictions. First, because the models rely on a periodic serial order representation, 2-back accuracy should be differentially disrupted if subjects must simultaneously complete a task that requires a different periodicity of serial order representations (e.g., a 3-movement spatial tapping task relative to a 2-movement one). Indeed, serial order may be important for precisely this type of motor control (Salinas, 2009).

Second, the neural substrates of self-imposed periodicity should be identifiable with fMRI, using regressors whose onsets correspond to a periodicity of  $n$ . Striatal activation should show the same parametric variation with  $n$  as has been previously observed in the cortex: Our model predicts these areas form a highly-interconnected circuit modulated by memory demands. Moreover, representational similarity analysis or other multi-voxel pattern analysis methods might reveal the same striatal hemodynamics reported here in terms of representational differentiation patterns.

Third, to the extent that humans are capable of good performance on  $n > 3$  -back tasks, they may recruit additional mechanisms, such as the hippocampal complex (HPC), to compensate for the logarithmically compressed nature of serial order representations in the IPS. Indeed, the HPC has only been inconsistently observed during performance of 2- and 3-back tasks (de Frias et al., 2010; Egan et al., 2003), and other accounts of  $n$ -back might predict HPC involvement only insofar as subjects adopt a low-control (i.e., familiarity-based) strategy (Juvina & Taatgen, 2007). In contrast, the current model predicts high-control strategies will involve additional mechanisms not modeled here, when  $n > 3$ .

This third set of predictions suggests several possible of the model to capture different strategies and training effects. The parietal layer is an important constraint on the ability of networks to perform adequately on  $n > 3$ -back tasks, because its tuning curves becomes increasingly compressed at higher serial orders. Parietal serial order representations simply become too compressed at high levels of  $n$  to support discrete representations. Nonetheless, because humans are apparently capable of learning  $n > 3$  back tasks with training, one extension to our model would be a top-down projection to this area from the prefrontal cortex. Over training, the network might learn to support increasingly discrete serial order representations using a top-down biasing signal (e.g., Edin et al., 2009). Our model might also be extended to capture the HPC mechanisms possibly used by subjects adopting a familiarity-based strategy. Conceptually similar mechanisms are used in a symbolic model of the  $n$ -back task (Juvina & Taatgen, 2007), such that a “time tagging” system is integrated with a familiarity system that relies on declarative memory. Different control strategies are then simulated in terms of whether the time tags are actively-maintained (as in our current model), or retrieved only when familiarity is detected. With the appropriate biological extensions, our model might capture these and more  $n$ -back phenomena.

Our model may be relevant to the burgeoning field of executive functions training, in which the  $n$ -back is playing a prominent role. For example,  $n$ -back performance improves following training on the letter memory task (Dahlin et al., 2008). Our model is also capable of performing the letter memory task, and the types of executive functions that emerge in our model from its training on letter memory are extremely similar to those reported here. However, in the current report our models were trained only on the  $n$ -back task; clearly, human performance in any task relies on a longer and more varied history of experience than the training we provided to our model. Future work will pretraining models on a larger variety of more elemental cognitive tasks and test transfer effects.

## Conclusions and Future Directions

Follow-up work is ongoing, including the more complete modeling of these and other neural structures with a role in executive functioning and a more elaborate mapping of this type of model to behavior in other executive function tasks. Indeed, the PBWM framework used here can model a number of other tasks, and the overlap among these models may reveal the computational origin of the unity and diversity of executive functions (Friedman et al., 2008; Miyake et al., 2000). The current work represents a first step in that direction, by specifying a formal computational link between anatomical connectivity studies demonstrating a highly interconnected parieto-fronto-striatal network with studies of genetic polymorphisms, with

individual differences at the behavioral level, and with theoretical accounts of the executive functions important for working memory updating tasks.

## Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

## Acknowledgments

Jeremy Reynolds and National Institute of Health (MH063207 and MH079485).

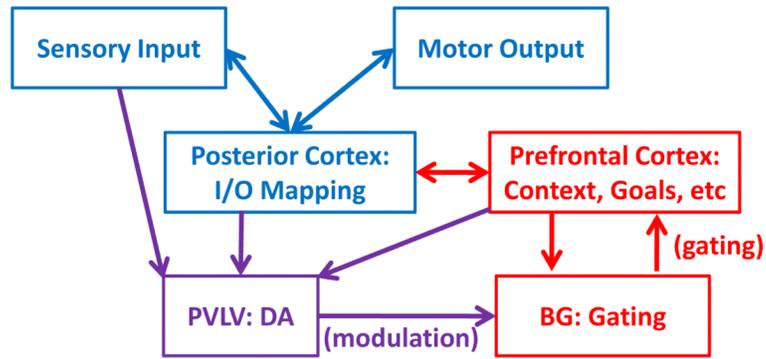
## REFERENCES

- Aalto S, Brück A, Laine M, Någren K, Rinne JO. Frontal and temporal dopamine release during working memory and attention tasks in healthy humans: A positron emission tomography study using the high-affinity dopamine D2 receptor ligand [11C]FLB 457. *Journal of Neuroscience*. 2005; 25:2471–2477. [PubMed: 15758155]
- Amiez C, Petrides M. Selective involvement of the mid-dorsolateral prefrontal cortex in the coding of the serial order of visual stimuli in working memory. *Proceedings of the National Academy of Sciences*. 2007; 104:13786–13791.
- Apud JA, Weinberger DR. Treatment of cognitive deficits associated with schizophrenia: Potential role of catechol-O-methyltransferase inhibitors. *CNS Drugs*. 2007; 21:535–557. [PubMed: 17579498]
- Badre D, Wagner AD. Computational and neurobiological mechanisms underlying cognitive flexibility. *Proceedings of the National Academy of Sciences U. S. A.* 2006; 103:7186–7191.
- Bilder RM, Volavka J, Lachman HM, Grace AA. The catechol-O-methyltransferase polymorphism: Relations to the tonic-phasic dopamine hypothesis and neuropsychiatric phenotypes. *Neuropsychopharmacology*. 2004; 29:1943–1961. [PubMed: 15305167]
- Botvinick M, Plaut DC. Short-term memory for serial order: A recurrent neural network model. *Psychological Review*. 2006; 113:201–233. [PubMed: 16637760]
- Botvinick M, Watanabe T. From numerosity to ordinal rank: A gain-field model of serial order representation in cortical working memory. *Journal of Neuroscience*. 2007; 27:8636–8642. [PubMed: 17687041]
- Boulton AA, Eisenhofer G. Catecholamine metabolism: From molecular understanding to clinical diagnosis and treatment. *Advances in Pharmacology*. 1998; 42:273–292. [PubMed: 9327897]
- Buzsáki G, Kaila K, Raichle M. Inhibition and brain work. *Neuron*. 2007; 56:771–783. [PubMed: 18054855]
- Chadderdon GL, Sporns O. A large-scale neurocomputational model of task-oriented behavior selection and working memory in prefrontal cortex. *Journal of Cognitive Neuroscience*. 2006; 18:242–257. [PubMed: 16494684]
- Chen J, Lipska BK, Halim N, Ma QD, Matsumoto M, Melhem, et al. Functional analysis of genetic variation in catechol-O-methyltransferase (COMT): Effects on mRNA, protein, and enzyme activity in postmortem human brain. *American Journal of Human Genetics*. 2004; 75:807–821. [PubMed: 15457404]
- Cohen JD, Braver TS, Brown JW. Computational perspectives on dopamine function in prefrontal cortex. *Current Opinion in Neurobiology*. 2002; 12:223–229. [PubMed: 12015241]
- Collette F, Van Der LM, Laureys S, Delfiore G, Degueldre C, Luxen A, Salmon E. Exploring the unity and diversity of the neural substrates of executive functioning. *Human Brain Mapping*. 2005; 25:409–423. [PubMed: 15852470]
- Cooper, RP.; Davelaar, EJ. Modelling the correlation between two putative inhibition tasks: A simulation approach. In: Salvucci, DD.; Gunzelmann, G., editors. *Proceedings of the 10th International Conference on Cognitive Modeling*; Philadelphia, PA: Drexel University; 2010. p. 31-36.

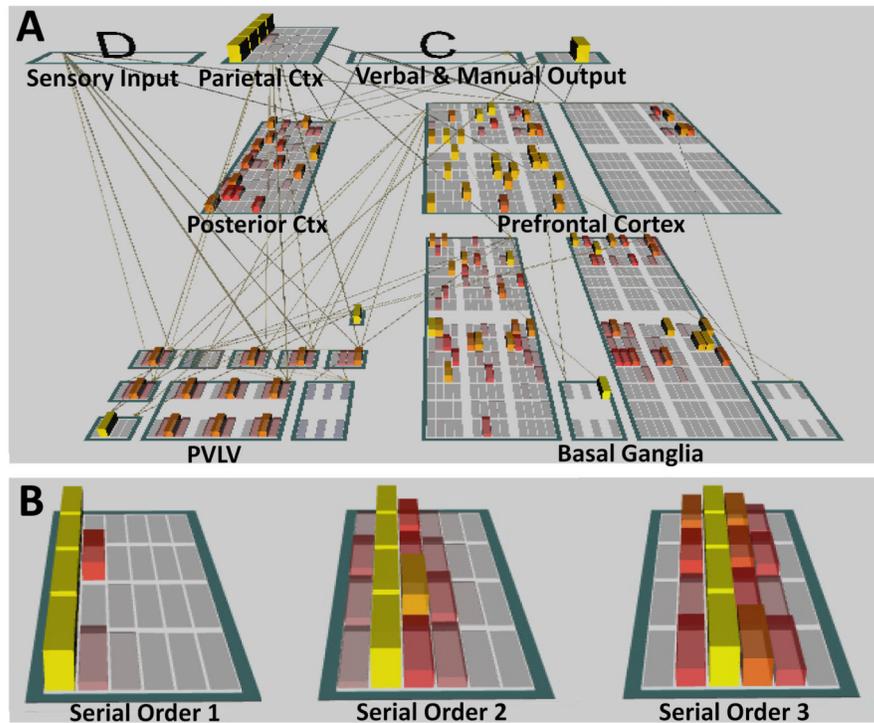
- Corbetta M, Shulman GL. Control of goal-directed and stimulus-driven attention in the brain. *Nature Reviews Neuroscience*. 2002; 3:215–229.
- Corbetta M, Patel GH, Shulman GL. The reorienting system of the human brain: From environment to theory of mind. *Neuron*. 2008; 58:306–324. [PubMed: 18466742]
- Dahlin E, Neely AS, Larsson A, Bäckman L, Nyberg L. Transfer of learning after updating training mediated by the striatum. *Science*. 2008; 320:1510–1512. [PubMed: 18556560]
- de Frias CM, Marklund P, Eriksson E, Larsson A, Oman L, Annerbrink K, et al. Influence of COMT gene polymorphism on fMRI-assessed sustained and transient activity during a working memory task. *Journal of Cognitive Neuroscience*. 2010; 22:1614–1622. [PubMed: 19642882]
- Deco G. A dynamical model of event-related FMRI signals in prefrontal cortex: Predictions for schizophrenia. *Pharmacopsychiatry*. 2006; 39:65–67.
- Desimone R, Duncan J. Neural mechanisms of visual selective attention. *Annual Review of Neuroscience*. 1995; 18:192–222.
- Durstewitz D, Seamans JK, Sejnowski TJ. Neurocomputational models of working memory. *Nature Neuroscience*. 2000; 3:1184–1191.
- Edin F, Klingberg T, Johansson P, McNab F, Tegnér J, Compte A. Mechanism for top-down control of working memory capacity. *Proceedings of the National Academy of Sciences U. S. A.* 2009; 106:6802–6807.
- Egan MF, Goldberg TE, Kolachana BS, Callicott JH, Mazzanti CM, Straub RE, et al. Effect of COMT Val108/158 Met genotype on frontal lobe function and risk for schizophrenia. *Proceedings of the National Academy of Sciences U. S. A.* 2001; 98:6917–6922.
- Egan MF, Kojima M, Callicott JH, Goldberg TE, Kolachana BS, Bertolino A, et al. The BDNF val66met polymorphism affects activity-dependent secretion of BDNF and human memory and hippocampal function. *Cell*. 2003; 112:257–269. [PubMed: 12553913]
- Fernández-Miranda JC, Rhoton AL, Alvarez-Linera J, Kakizawa Y, Choi C, de Oliveira EP. Three-dimensional microsurgical and tractographic anatomy of the white matter of the human brain. *Neurosurgery*. 2008; 62:989–1026. [PubMed: 18695585]
- Friedman NP, Miyake A. The relations among inhibition and interference control functions: A latent variable analysis. *Journal of Experimental Psychology: General*. 2004; 133:101–135. [PubMed: 14979754]
- Friedman NP, Miyake A, Young SE, Defries JC, Corley RP, Hewitt JK. Individual differences in executive functions are almost entirely genetic in origin. *Journal of Experimental Psychology: General*. 2008; 137:201–225. [PubMed: 18473654]
- Goldberg TE, Egan MF, Gscheidle T, Coppola R, Weickert T, Kolachana BS, et al. Executive subprocesses in working memory: Relationship to catechol-O-methyltransferase Val158Met genotype and schizophrenia. *Archives of General Psychiatry*. 2003; 60:889–896. [PubMed: 12963670]
- Goschke, T. Involuntary persistence and intentional reconfiguration in task-set switching. In: Monsell, S.; Driver, J., editors. *Attention and performance XVIII: control of Cognitive Processes*. MIT Press; Cambridge, MA: 2000. p. 331-355.
- Gray JR, Chabris CF, Braver TS. Neural mechanisms of general fluid intelligence. *Nature Neuroscience*. 2003; 6:316–322.
- Hazy TE, Frank MJ, O'Reilly RC. Banishing the homunculus: Making working memory work. *Nature Neuroscience*. 2006; 139:105–118. [PubMed: 16343792]
- Hazy TE, Frank MJ, O'Reilly RC. Towards an executive without a homunculus: Computational models of the prefrontal cortex/basal ganglia system. *Philosophical Transactions of the Royal Society B*. 2007; 362:1601–1613.
- Hazy TE, Frank MJ, O'Reilly RC. Neural mechanisms of acquired phasic dopamine responses in learning. *Neuroscience & Biobehavioral Reviews*. 2010; 34:701–720. [PubMed: 19944716]
- Jonides J, Nee DE. Brain mechanisms of proactive interference in working memory. *Neuroscience*. 2006; 139:181–193. [PubMed: 16337090]
- Jonides J, Smith EE, Marshuetz C, Koeppel RA, Reuter-Lorenz PA. Inhibition in verbal working memory revealed by brain activation. *Proceedings of the National Academy of Sciences U. S. A.* 1998; 95:8410–8413.

- Juvina, I.; Taatgen, NA. Modeling control strategies in the n-back task. Proceedings of the 8th International Conference on Cognitive Modeling; Michigan, USA: Ann Arbor; 2007.
- Kirchner WK. Age differences in short-term retention of rapidly changing information. *Journal of Experimental Psychology*. 1958; 55:352–358. [PubMed: 13539317]
- Levitt JB, Lewis DA, Yoshioka T, Lund JS. Topography of pyramidal *Neuron* intrinsic connections in macaque monkey prefrontal cortex (Areas 9 & 46). *The Journal of Comparative Neurology*. 1993; 338:360–376. [PubMed: 8113445]
- Logie RH, Cocchini G, Delia Sala S, Baddeley AD. Is there a specific executive capacity for dual task coordination? Evidence from Alzheimer's disease. *Neuropsychology*. 2004; 18:504–513. [PubMed: 15291728]
- Männistö PT, Kaakkola S. Catechol-O-methyltransferase (COMT): Biochemistry, molecular biology, pharmacology, and clinical efficacy of the new selective COMT inhibitors. *Pharmacology Review*. 1999; 51:593–628.
- Marshuetz C, Reuter-Lorenz PA, Smith EE, Jonides J, Noll DC. Working memory for order and the parietal cortex: An event-related fMRI study. *Neuroscience*. 2006; 139:311–316. [PubMed: 16417974]
- Marshuetz C, Smith EE, Jonides J, Degutis J, Chenevert TL. Order information in working memory: fMRI evidence for parietal and prefrontal mechanisms. *Journal of Cognitive Neuroscience*. 2000; 12:130–144. [PubMed: 11506653]
- Mattay VS, Goldberg TE, Fera F, Hariri AR, Tessitore A, Egan MF, et al. Catechol O-methyltransferase val158-met genotype and individual variation in the brain response to amphetamine. *Proceedings of the National Academy of Sciences U S A*. 2003; 100:6186–6191.
- Meyer-Lindenberg A, Nichols T, Callicott JH, Ding J, Kolachana B, Buckholtz J, et al. Impact of complex genetic variation in COMT on human brain function. *Molecular Psychiatry*. 2006; 11:867–877. [PubMed: 16786032]
- Miyake A, Friedman NP, Emerson MJ, Witzki AH, Howerter A, Wager TD. The unity and diversity of executive functions and their contributions to complex “frontal lobe” tasks: A latent variable analysis. *Cognitive Psychology*. 2000; 41:49–100. [PubMed: 10945922]
- Nee DE, Jonides J, Berman MG. Neural mechanisms of proactive interference-resolution. *Neuroimage*. 2007; 38:740–751. [PubMed: 17904389]
- Nieder A, Diester I, Tudusciuc O. Temporal and spatial enumeration processes in the primate parietal cortex. *Science*. 2006; 313:1431–1435. [PubMed: 16960005]
- Oberauer K. Binding and inhibition in working memory: Individual and age differences in short-term recognition. *Journal of Experimental Psychology: General*. 2005; 134:368–387. [PubMed: 16131269]
- Olesen PJ, Westerberg H, Klingberg T. Increased prefrontal and parietal activity after training of working memory. *Nature Neuroscience*. 2003; 7:75–79.
- O'Reilly RC. Generalization in interactive networks: The benefits of inhibitory competition and hebbian learning. *Neural Computation*. 2001; 13:1199–1242. [PubMed: 11387044]
- O'Reilly, RC.; Busby, RS.; Soto, R. Three forms of binding and their neural substrates: Alternatives to temporal synchrony. In: Cleeremans, A., editor. *The Unity of Consciousness: Binding, Integration, and Dissociation*. Oxford University Press; Oxford: 2003. p. 168-192.
- O'Reilly RC, Frank MJ. Making working memory work: A computational model of learning in the frontal cortex and basal ganglia. *Neural Computation*. 2006; 18:283–328. [PubMed: 16378516]
- O'Reilly RC, Frank MJ, Hazy TE, Watz B. PVLV: The Primary Value and Learned Value pavlovian learning algorithm. *Behavioral Neuroscience*. 2007; 121:31–49. [PubMed: 17324049]
- O'Reilly, RC.; Munakata, Y. *Computational explorations in cognitive neuroscience: Understanding the mind by simulating the brain*. MIT Press; Cambridge: 2000.
- Owen AM, McMillan KM, Laird AR, Bullmore E. N-back working memory paradigm: A meta-analysis of normative functional neuroimaging studies. *Human Brain Mapping*. 2005; 25:46–59. [PubMed: 15846822]
- Pucak ML, Levitt JB, Lund JS, Lewis DA. Patterns of intrinsic and associational circuitry in monkey prefrontal cortex. *The Journal of Comparative Neurology*. 1996; 376:614–630. [PubMed: 8978474]

- Rao SG, Williams GV, Goldman-Rakic PS. Isodirectional tuning of adjacent interneurons and pyramidal cells during working memory: Evidence for microcolumnar organization in PFC. *Journal of Neurophysiology*. 1999; 81:1903–1916. [PubMed: 10200225]
- Rushworth MF, Behrens TE, Johansen-Berg H. Connection patterns distinguish 3 regions of human parietal cortex. *Cerebral Cortex*. 2006; 16:1418–1430. [PubMed: 16306320]
- Salinas E. Rank-order-selective neurons form a temporal basis set for the generation of motor sequences. *Journal of Neuroscience*. 2009; 29:4369–4380. [PubMed: 19357265]
- Salthouse TA, Atkinson TM, Berish DE. Executive functioning as a potential mediator of age-related cognitive decline in normal adults. *Journal of Experimental Psychology: General*. 2003; 132:566–594. [PubMed: 14640849]
- Schleepen TJ, Jonkman LM. The development of non-spatial working memory capacity during childhood and adolescence and the role of interference control: An n-back task study. *Developmental Neuropsychology*. 2010; 35:37–56. [PubMed: 20390591]
- Stefanis NC, van Os J, Avramopoulos D, Smyrnis N, Evdokimidis I, Stefanis CN. Effect of COMT Val158Met polymorphism on the Continuous Performance Test, Identical Pairs Version: tuning rather than improving performance. *American Journal of Psychiatry*. 2005; 162:1752–1754. [PubMed: 16135641]
- Tagamets MA, Horwitz B. A model of working memory: Bridging the gap between electrophysiology and human brain imaging. *Neural Networks*. 2000; 13:941–952. [PubMed: 11156203]
- Tan HY, Chen Q, Goldberg TE, Mattay VS, Meyer-Lindenberg A, Weinberger DR, Callicott JH. Catechol-O-methyltransferase Val158Met modulation of prefrontal-parietal-striatal brain systems during arithmetic and temporal transformations in working memory. *Journal of Neuroscience*. 2007; 49:13393–13401. [PubMed: 18057197]
- Tsuchida A, Fellows LK. Lesion evidence that two distinct regions within prefrontal cortex are critical for *n-back* performance in humans. *Journal of Cognitive Neuroscience*. 2009; 21:2263–2275. [PubMed: 19199405]
- Van Dijck J, Fias W. A working memory account for spatial-numerical associations. *Cognition*. (in press).
- Wager TD, Smith EE. Neuroimaging studies of working memory: A meta-analysis. *Cognitive, Affective and Behavioral Neuroscience*. 2003; 3:255–274.
- Winterer G, Musso F, Vucurevic G, Stoeter P, Konrad A, Seker B, et al. COMT genotype predicts BOLD signal and noise characteristics in prefrontal circuits. *Neuroimage*. 2006; 32:1722–1732. [PubMed: 16884927]
- Yeterian EH, Pandya DN. Striatal connections of the parietal association cortices in rhesus monkeys. *Journal of Comparative Neurology*. 1993; 332:175–197. [PubMed: 8331211]



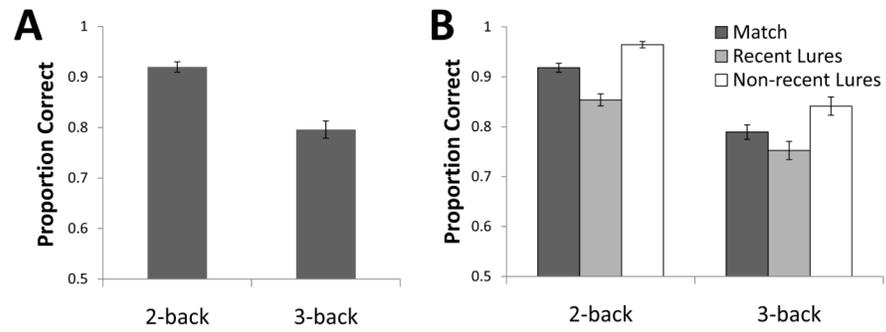
**Figure 1.** Schematic illustration of core PBWM architecture, in which prefrontal context representations of relevant prior information and current goals bias the sensory-motor mappings that are learned by posterior cortical “hidden” layers. The prefrontal context representations are updated via dynamic gating by the basal ganglia. These gating functions are learned by the basal ganglia on the basis of input from the PVLV system, which provides modulatory dopaminergic input depending on the reward value of the actions performed by the basal ganglia.



**Figure 2.**

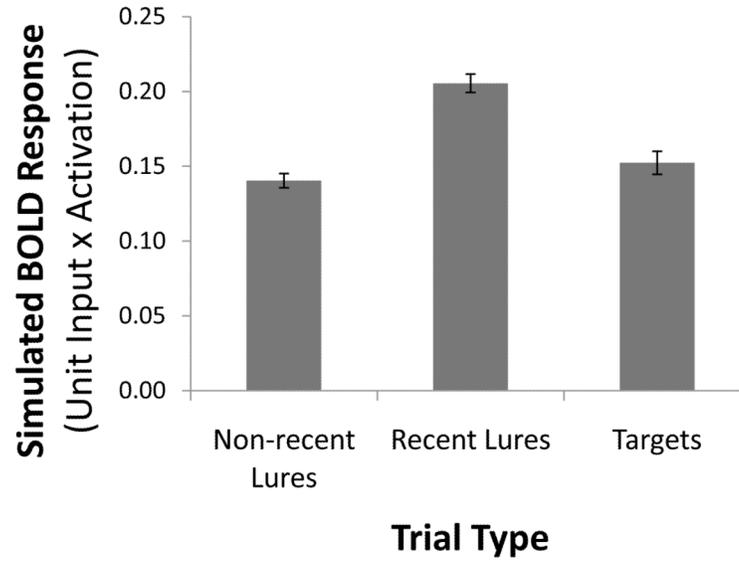
**A.** The PBWM architecture includes units based on the prefrontal cortex and basal ganglia, including ventral and dorsal striatum, grouped into “stripes” (the visible subgroups within prefrontal and striatal layers). Input is provided to the model about the identity of the current stimulus and its serial order; the model is required to produce a manual output about whether the current stimulus matches that presented  $n$  trials previously, and a verbal output corresponding to the identity of the stimulus presented  $n$  trials previously. **B.** The parietal layers represent the serial order of successive trials in terms of  $n$ , using a graded and compressive code based on the mean and variance observed in the tuning curves of rank order sensitive neurons in the horizontal segment of the intraparietal sulcus.



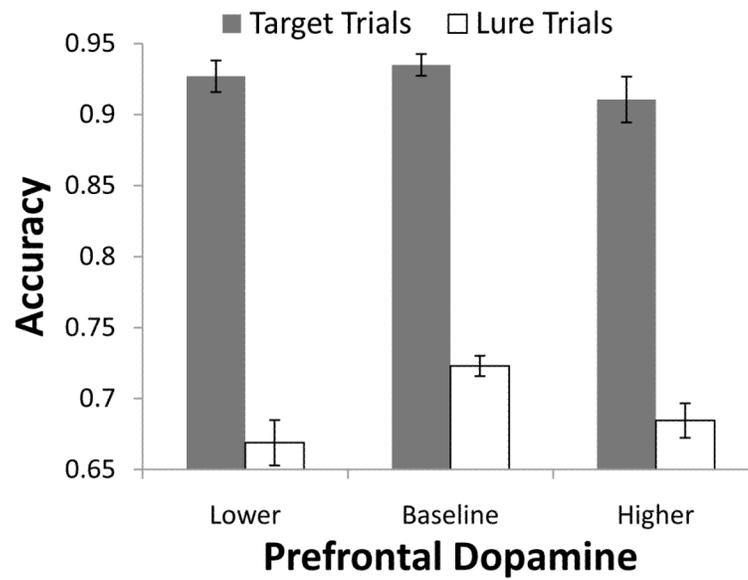


**Figure 4.**

**A.** The model reproduces the benchmark result of lower accuracy on 3-back than 2-back. **B.** The model shows reduced accuracy on recent ( $n-1$ ) lures, relative to both non-recent lures ( $>n$ ) and match trials. In addition, the relative difference of these trial types is smaller in the 3-back task than the 2-back task, consistent with human data.

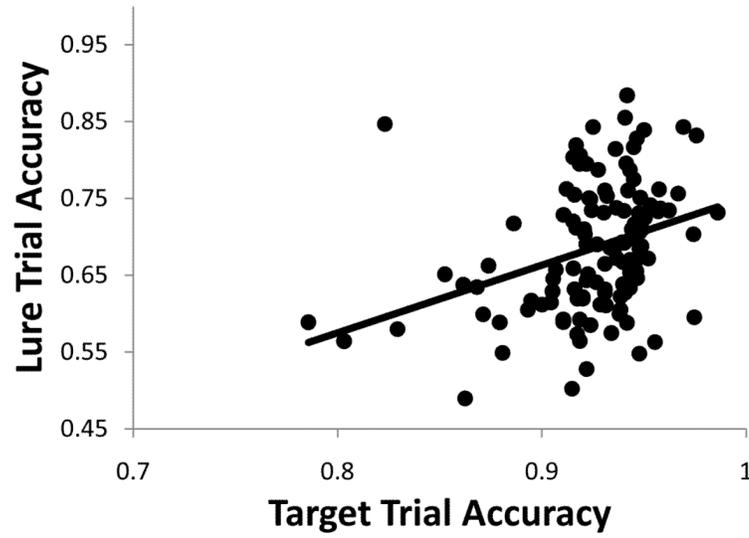


**Figure 5.** Recent lures were associated with a greater simulated hemodynamic response than non-recent lures and targets, where the hemodynamic response is simulated as the weighted average of unit inputs and unit activations in the prefrontal cortex layers.

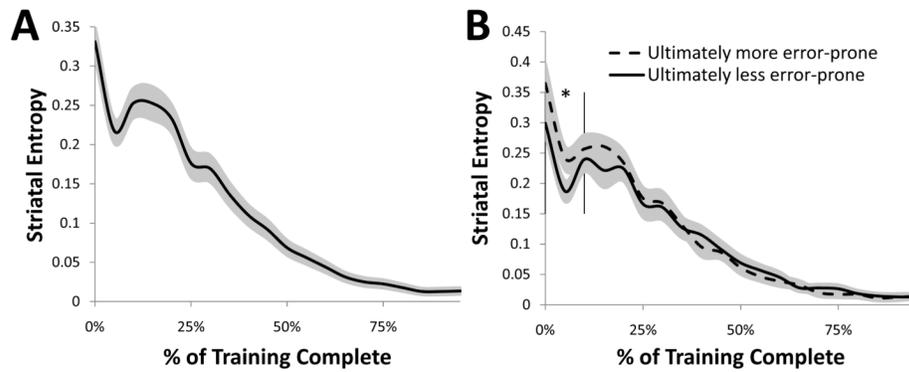


**Figure 6.**

As a proxy for the effects of the polymorphisms in the COMT gene, we manipulated the effects of dopamine in the prefrontal layers of the model. This manipulation revealed an inverted U-shaped curve relating dopamine levels to performance, consistent with the hypothesized effects of varying dopamine levels in prefrontal cortex.

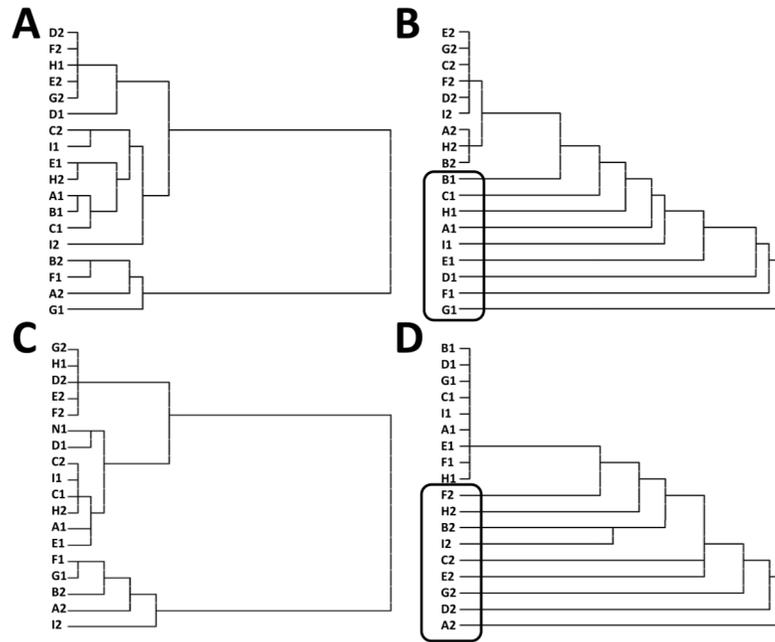


**Figure 7.** The model captures the individual differences in the relationship of lure and target trial accuracy observed empirically when subjects are encouraged to adopt the same strategy as adopted by our model.



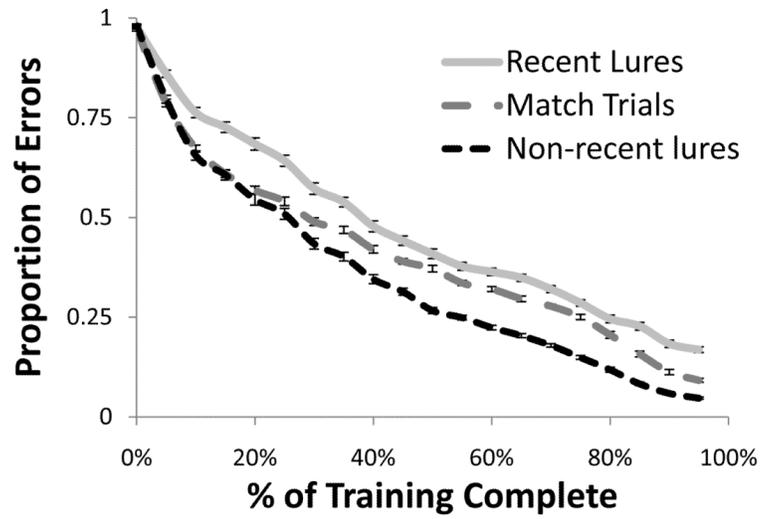
**Figure 8.**

**A.** The model learns to appropriate gate information into working memory by developing increasingly discrete firing patterns in the striatum over the course of training, here visualized in terms of reductions in entropy. **B.** Individual differences in the ultimate post-training performance of models across runs can be predicted based on the reduction in striatal entropy much earlier in training: networks that ultimately commit less errors following training (solid vs. dotted lines) show significantly more (\*  $p < .05$ ) discrete patterns of firing between 0 and 10% of the total training time (vertical bars). Shaded regions represent standard error of the mean for each time point.



**Figure 9.**

Cluster plots reflect the Euclidean distance (indicated by the length of horizontal lines) between every item (indicated by letters along the y axis) and serial order (indicated by numbers along the y-axis); thus, if the path from one item to another requires a large amount of horizontal travel, then the representations of those items are relatively distinct. **A.** One prefrontal stripe shows an initially haphazard pattern of representational similarity across items, as indicated by the lack of systematic clustering between items and their order. **B.** After training, the same prefrontal stripe illustrated in **A** develops a highly structured representation, by collapsing across all items of serial order 2 (upper half of cluster plot) but differentiating among all items of serial order 1 (as indicated by the large horizontal lines separating each item; lower half, enclosed by rounded rectangle). This stripe is preferentially tuned to code items of serial order 1. **C.** A different PFC stripe also shows initially haphazard representational similarity. **D.** After training this stripe shows a different pattern than that illustrated in **B**, in that it collapses equally across all items of serial order 1 (upper half of cluster plot) but increasingly differentiates every item occurring with the other serial order (lower half, enclosed by rounded rectangle).



**Figure 10.** Performance on recent lures trials undergoes a shallower learning curve than performance on all other trial types, reflecting a more rapid reduction in error rate on trials that do not require the resolution of proactive interference (match and non-recent lure trials as compared to recent lures).