Published in final edited form as:

Cognition. 2009 December ; 113(3): 281–292. doi:10.1016/j.cognition.2009.05.015.

Developing PFC representations using reinforcement learning

Jeremy R. Reynolds* and

Department of Psychology, University of Denver

Randall C. O'Reilly

Department of Psychology, University of Colorado

Abstract

From both functional and biological considerations, it is widely believed that action production, planning, and goal-oriented behaviors supported by the frontal cortex are organized hierarchically (Fuster, 1990, Koechlin, Ody, & Kouneiher, 2003, & Miller, Galanter, & Pribram, 1960) However, the nature of the different levels of the hierarchy remains unclear, and little attention has been paid to the origins of such a hierarchy. We address these issues through biologically-inspired computational models that develop representations through reinforcement learning. We explore several different factors in these models that might plausibly give rise to a hierarchical organization of representations within the PFC, including an initial connectivity hierarchy within PFC, a hierarchical set of connections between PFC and subcortical structures controlling it, and differential synaptic plasticity schedules. Simulation results indicate that architectural constraints contribute to the segregation of different types of representations, and that this segregation facilitates learning. These findings are consistent with the idea that there is a functional hierarchy in PFC, as captured in our earlier computational models of PFC function and a growing body of empirical data.

Keywords

PFC; representation; reinforcement learning; functional organization

The prefrontal cortex (PFC) plays a critical role in the execution of controlled behavior (Miller & Cohen, 2001). Many theories exist regarding the function of PFC (for reviews, see Stuss & Knight, 2002; Wood & Grafman, 2003), and this plethora of theories in part reflects our lack of understanding concerning the functional organization of the multiple anatomical areas that compose PFC (Duncan & Owen, 2000; Miller, 2000). This lack of understanding is not for a lack of trying; there have been a number of approaches to investigating this critical question. One approach has been to focus on various stimulus dimensions that have produced reliable dissociations in posterior areas. For example, researchers have hypothesized that the dorsal and ventral visual processing streams project into PFC, and cause a spatial vs. object dissociation along a dorsal and ventral gradient (Goldman-Rakic, 1987; Haxby, Petit, Ungerleider, & Courtney, 2000; Wilson, Scalaidhe, & Goldman-Rakic, 1993). Although there appears to be some evidence suggesting that areas of PFC may be dissociated along these stimulus dimensions (Johnson, Raye, Mitchell, Greene, & Anderson, 2003), there are also

Correspondence concerning this article should be addressed to Jeremy Reynolds, Department of Psychology, University of Denver, 2155 S. Race St, Denver, CO 80208. E-mail: jeremy.reynolds@psy.du.edu.

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

studies that have not been able to identify such distinctions (i.e. Nystrom, Braver, Sabb, Delgado, Noll, & Cohen, 2000).

A second approach has focused on investigating whether different processes determine the functional organization of the PFC. In particular, a common distinction in the literature has suggested that ventrolateral PFC (BA 44/45) is responsible for maintaining and rehearsing information in working memory (WM) whereas dorsolateral PFC (BA 9/46) is responsible for the manipulation and monitoring of such information (D'Esposito, Postle, Ballard, & Lease, 1999; Haxby et al., 2000; Petrides, 2000). Although there have been studies supporting this claim, there have been numerous studies demonstrating areas of dorsolateral PFC active under situations that do not require manipulation *prima facie* (Braver & Bongiolatti, 2002; Braver, Cohen, & Barch, 2002; Raye, Johnson, Mitchell, Reeder, & Greene, 2002). In addition to the hypothesis regarding dorsolateral and ventrolateral PFC, studies have also suggested that there is an additional gradient dissociating anterior PFC (BA 10) from more posterior areas of lateral PFC based on process, although the nature of that process is unclear. For example, data suggest that anterior PFC may process internally generated information (Christoff & Gabrieli, 2000; Christoff, Ream, Geddes, & Gabrieli, 2003), or that it may be associated with multi-tasking aspects of maintaining one goal in an active state while executing or scheduling a different one (Braver & Bongiolatti, 2002; Braver, Reynolds, & Donaldson, 2003; Koechlin, Basso, Pietrini, Panzer, & Grafman, 1999; Reynolds, West, & Braver, 2008). Finally, recent evidence has suggested that PFC may be hierarchically organized, such that more posterior, dorsal regions are intimately tied to determining the appropriate response, and more anterior regions process more abstract aspects of the task that are summed to inform the response (Koechlin, Ody, & Kouneiher, 2003; Koechlin & Summerfield, 2007). Future research will be needed to determine whether these distinctions are isomorphic, or whether they capture something fundamentally different about cognition and PFC function.

Despite the large number of theories and studies investigating the functional organization of PFC, there is not a satisfactory set of theories that is strongly supported by the available data. As mentioned above, one problem, particularly in the approach analyzing organization according to different processes, is that there do not appear to be general, grounded definitions of each proposed process. For example, What constitutes a manipulation in the context of the maintenance/manipulation distinction? Does simply updating the contents of working memory constitute a manipulation? Likewise, what does it mean to internally generate information? Any transformation upon externally presented information could be considered internally generated; where does one draw the line? Which specific computations are being performed that facilitate multi-task and sub-goal performance? Finally, what are the specific computations/dimensions that determine the aspect of the task that is being represented; (e.g. in the Koechlin, Ody, and Kouneiher (2003) tasks, what causes something to represent task context vs. episode context)?

One approach that can help address the ambiguity over the functional definition of various processes is computational modeling. To this end, the current simulations investigated the extent to which a biologically-based computational modeling framework can provide insight into the development of specialized neural representations in different PFC areas, with a focus on the hierarchical structuring of such representations (Miller, Galanter, & Pribram, 1960; Fuster, 1991; Koechlin et al., 2003). In particular, there are substantial gains to be had by hierarchically organized systems; systems organized in such a way frequently learn faster than systems that contain no such organization (Botvinick, 2008), and display substantially greater generalization (Lashley, 1951; Miller et al., 1960; Botvinick, 2008). Effectively, a hierarchically organized system allows individuals to form abstractions (e.g. a plan to traverse a room to the opposite door) that can be applied in a variety of different contexts. Once discovered, these abstractions, or chunks, can then be applied to novel situations.

We explore these issues in the context of a relatively sophisticated task that has a hierarchical structure: the 1-2 AX- continuous performance task (Frank, Loughry, & O'Reilly, 2001). In this task, letters and numbers are presented sequentially over time, and participants must detect specific target sequences. The appropriate response to a particular sequence (such as A-X) is dependent on the most recently viewed number; thus, the cues (A's) and probes (X's) are nested hierarchically within an outer-loop of number information (see Figure 1). For example, if a 1 was last seen, the target is A-X, but if a 2 was last seen, the target is B-Y (see Table 1). The computational model we employ here can successfully learn to perform this task (O'Reilly & Frank, 2006) using a reinforcement learning algorithm called PVLV (O'Reilly, Frank, Hazy, & Watz, 2007), coupled to a biologically-based model of the Basal Ganglia (BG) and PFC circuitry, all of which are detailed in the online supplementary materials.

Specifically, we build on previous work suggesting that the BG can dynamically and selectively gate the contents of PFC, and thus permit task-relevant information to be maintained in PFC while preventing extraneous information from interfering with performance (Frank et al., 2001; O'Reilly & Frank, 2006). A key property of this approach is that it allows for selective updating via the inclusion of parallel loops of connectivity between the BG and PFC, such that a particular set of neurons within the BG can control the updating of a particular set of neurons within the PFC (which we refer to as a stripe). This ability to selectively update some contents of WM while leaving other content intact is a process that is fundamental to hierarchical behavior, because the nature of the task representations have, by definition, different temporal dynamics. Larger goals are relevant over longer periods of time, and thus should not be updated once one sub-goal is completed and another is begun. The current framework provides an ideal avenue for exploring hierarchical behavior, because this ability to perform asynchronous updating has been established across several tasks, including the one explored in the current manuscript (O'Reilly & Frank, 2006).

While previous explorations have interrogated the ability of this type of network to learn and perform these tasks (O'Reilly & Frank, 2006), there has been little attention to the types of representations that are developed. Within the context of the task we investigated, outer-loop information (number) is available at the time at which the network processes inner-loop information (letter cues and probes). Thus, outer-loop information has the opportunity to guide and shape the developing inner-loop representations. We investigated the extent to which various manipulations of the network connectivity and other parameters would influence the degree to which the representation of these two types of information (inner-loop vs. outer-loop) were segregated or interdigitated.

The first simulation investigated whether a natural connectivity hierarchy would be sufficient to constrain the development of representations within PFC. This type of hierarchy has been associated with neuroanatomical data (Fuster, 1991, 2004; Petrides & Pandya, 2007), and has been shown to encourage the representation of stable context representations in previous computational models (Botvinick, 2007).

Additionally, we investigated the idea that a hierarchical structure allows for more generalized, flexible behavior by having higher-level areas dynamically contextualize and modulate the input/output mappings represented in lower-level systems (Lashley, 1951; Miller et al., 1960; Botvinick, 2008).

The third simulation is motivated by developmental data suggesting that a) cognitive development seems to proceed in a hierarchical fashion, where lower levels of cognitive processing are established prior to the development of more complex processing (Bunge & Zelazo, 2006; Halford, Wilson, & Phillips, 1999), and b) that this behavioral timecourse could be mediated by developments in corresponding areas of PFC (Bunge & Zelazo, 2006; Gogtay,

Giedd, Lusk, Hayashi, Greenstein, Vaituzis, Nugent, Herman, Clasen, Toga, Rapoport, & Thompson, 2004; O'Donnell, Noseworth, Levine, & Dennis, 2005). In order to model this type of developmental trajectory, we included additional tasks in the training regimen. These tasks (called the 3-4 tasks) had exactly the same structure as the 1-2 AX task, but the stimulus contents of each inner loop were unique to each task (e.g. in the 3 task, R-P might be a target sequence, and those stimuli would never occur in any other outer-loop context, see Table 1). As such, the cue-probe pair uniquely determines the response, and there is no need to maintain or process the outer-loop information. In the context of these tasks, a developmental trajectory would suggest that one learns the easier, less complex tasks first, the representational content of which occupies resources. As one begins to learn more complicated tasks, the resources that were previously available are now occupied, and additional resources need to be utilized in order to solve the problem. The implementation of each approach is detailed in the Methods below.

To foreshadow the results, we found that a.) anatomical constraints provide a strong constraint in the learning of such representations, and resulting hierarchical segregated representations facilitate learning, b) Outer-loop information is represented both in the form of stable representations across trials as well as conjunctive representations that dynamically update each trial, and c.) differential plasticity schedules have little influence on the development of segregated or hierarchical representations within the context of the various models implemented here.

Methods

Tasks

The model was trained to perform four different versions of the AX-CPT. In each of the versions of the AX-CPT modeled, the model was asked to make a target response to a particular probe (e.g. "X") that follows a particular cue (e.g. "A"), and to make a non-target response after all other cue-probe combinations. Two of the versions followed this exact format; the only difference was the set of stimuli used for each task (see Table 1, columns 1-5). In the other two versions of the task, the exact same stimuli were used across the tasks, but the appropriate response mappings changed as a function of the task-level (henceforth called outer-loop) context (the number; see Table 1, columns 6-9). Specifically, if the most recent task cue was a "1," then the appropriate target sequence would be A-X, but if the most recent task cue was a "2," then the appropriate target sequence would be B-Y (see Figure 1). The probability of the cue-probe pairs was adjusted from the typical expectancy versions of the AX-CPT (Braver, Barch, Keys, Carter, Cohen, Kaye, Janowsky, Taylor, Yesavage, Mummenthaler, Jagust, & Reed, 2001) by decreasing the relative probability of the target sequences and increasing the relative probability of non-target trials that have no stimulus overlap with the target sequence (e.g. BY trials when AX is the target). The primary reason behind this change was so that the network would not be able to identify the current task by keeping track of the relative frequency of the AX and BY sequences. In this way, the stimulus frequencies provide no diagnostic information with respect to what task the model is supposed to be doing. Each task was equally likely to appear, and once the model encountered a task cue, the model would be asked to perform between 1 and 5 trials associated with that task prior to seeing a new task cue.

Each input and output was represented by its own localist representation. All other layers had distributed representations that were learned over the course of training (including the PFC, see below).

Modeling and training

The implemented models use the Leabra framework described in the appendix (O'Reilly, 1998, 2001; O'Reilly & Munakata, 2000), with the additional specialized prefrontal-cortex

basal ganglia working memory (PBWM) mechanisms described in O'Reilly and Frank (2006). As a brief summary, the Leabra framework uses point neurons with excitatory, inhibitory, and leak conductances contributing to an integrated membrane potential, which is then thresholded and transformed via an $x/(x + 1)$ sigmoidal function to produce a rate code output. Each layer uses a k-winners-take all (kWTA) function that computes an inhibitory conductance that keeps roughly the k most active units above firing threshold and keeps the rest below threshold. Units learn according to a combination of Hebbian, error-driven, and reinforcement learning, with the error-driven component computed using the generalized recirculation algorithm (GeneRec; O'Reilly, 1996), which computes backpropagation derivatives using two phases of activation settling. The cortical layers in the model use standard Leabra parameters and functionality, while the PBWM systems require some additional mechanisms to implement the DA modulation of Go/NoGo units, and toggling of PFC maintenance currents, as detailed in the supplementary online materials.

The base network architecture followed the organization depicted in Figure 2, and it was largely consistent across simulation. The input, hidden, and output layers consisted of 16, 49, and 2 units, respectively. The PFC layer consisted of 4 stripes of 36 units each, while the Matrix layer consisted of 4 stripes of 28 units each (14 Go and 14 NoGo units). Implementation of the primary value- learned value (PVLV) system is detailed in the supplementary materials. If two layers were connected, all units in one layer were connected with all units in the other, except for those connected to the PFC. The set of projections to the PFC was a target of manipulation within the Simulations, and the details regarding its various connectivity patterns are enumerated below.

Training procedure

The networks were trained to a similar criterion as in previous models using this task (O'Reilly & Frank, 2006; O'Reilly et al., 2007). Specifically, the networks had to complete 80 sequences of 1-5 trials with 0 errors (i.e. one epoch). Each epoch was constrained to include at least 4 trials of each type defined in Table 1. Although the stopping criterion is defined differently from previous published work with this task, the current criterion corresponds to approximately the same number of correct trials per outer loop context as previously used criteria (current = 60 per outer-loop context, previous = mean of 62.5 per outer-loop context). Each network was trained 100 different times, with each training run initialized with different random seeds.

In order to collect behavioral and PFC activity data, learning was turned off after training, and each network was run through the tasks for an additional 240 sequences.

Data and Analysis

Behavior / Performance—The behavior and performance of the networks were tested in two different ways. First, the rate of successful training and the time to criterion (number of epochs) was recorded for each replication, so that any manipulations could be tested to determine whether they influence learning. Second, we analyzed performance over the testing session in order to confirm that the networks were using active maintenance rather than weight-based memory. If weight-based memory were being utilized, then performance would greatly degrade during testing, as the only mechanism for memory is through active maintenance. Performance is reported for each type of task [i.e. low complexity (3-4) vs. high complexity (1-2)].

Representations—Due to the nature of our manipulations, we collapsed the four PFC stripes into two groups of two stripes. Within these two groups of stripes, a neurophysiology approach was taken to probe the types of representations the network was using to perform the task. This procedure involved using the activity patterns at the time of each probe in order to identify

units demonstrating a significant response relative to baseline in any one of the four tasks ($p < 0.001$), and then probing further to determine whether that response was selective to the task. This selectivity was determined by t -tests comparing the response to that particular task to each of the other tasks; if the response was different across all comparisons, it was said to be selective.

In order to interrogate differences between the two groups of units, an index was used to compare the proportion of each type of selective unit, controlling for the overall proportion of selective units:

$$\frac{p(\text{Exp}) - p(\text{Control})}{p(\text{Exp}) + p(\text{Control})} \quad (1)$$

where $p(\text{Control})$ is the proportion of selective units in the set of stripes with normal connectivity from the hidden layer and BG, and $p(\text{Exp})$ is the proportion of selective units in the set of stripes with the experimental manipulation (see Figure 2B-D). As reflected in Figure 2B-D, the experimental group was defined for each simulation by either a) replacing the projections between it and the hidden layer with projections to and from the other PFC group, b) adding an additional projection to the BG stripes responsible for updating the control group of PFC units, or c) delaying their learning. The experimental group always corresponds to the group of units that would be predicted to have more frequent outer-loop representations. The index places all measures on a -1 to 1 scale in which -1 means all selective units occur in the control group, 1 means all selective units occur in the experimental group, and 0 is an equal distribution of selective units across the groups. These ratios were then analyzed across multiple training runs to determine whether there were systematic differences in the type of representations that develop.

Simulations

The first two simulations investigate whether two similar, but distinct, computational constraints could potentially lead to the development of segregated representations: the ability of one set of units to directly influence the content of other units via direct biasing and the ability of one set of units to directly influence the timing of the updating signal associated with other units. The first simulation investigated whether a natural connectivity hierarchy would be sufficient to constrain the development of representations within PFC. This type of hierarchy has been associated with neuroanatomical data (Fuster, 1991, 2004; Petrides & Pandya, 2007), and has been shown to encourage the development of stable context representations in previous computational models (Botvinick, 2007; Paine & Tani, 2005). This hypothesis was implemented as shown in Figure 2C, and is characterized by having bidirectional connections between the hidden layer and a group of stripes in the PFC layer, and bidirectional connectivity between these PFC stripes and the remaining PFC stripes. This is an explicit hierarchical structure similar to that used by Botvinick (2007) and Paine and Tani (2005), in which the stripes of PFC connected directly to the hidden layer are considered proximal the processing pathway, and as such, should come to represent information that is closer to the output pathway. Conversely, the PFC stripes connected only via the intermediate processing of PFC should reflect the apex of the processing pathway, and is predicted to represent the most abstract/temporally extended information. This network is labeled the PFC-Hidden network. The representations developed by this network can then be compared to the representations developed in the context of a standard network in which all of PFC has access to information from the hidden layer, and does not direct connect with other areas of PFC (see Figure 2B).

The second simulation investigated whether architectural changes allowing for higher-level areas to dynamically contextualize and modulate the input/output mappings represented in lower-level systems would encourage the development of segregated representations. This hypothesis was implemented as displayed in Figure 2D; the experimental group of PFC units (see Figure 2D), in addition to projecting to their corresponding BG stripes, also projected to the BG stripes associated with the group of control PFC units (labeled the PFC-BG network). This structure of projections was chosen for one primary reason. This projection formalizes the idea that anterior PFC has a strong direct influence over the timing of the updating process within lateral PFC. By influencing when lateral PFC is updated, the alternate area of PFC could likely shape the type of representations used by lateral PFC.

The third set of simulations interrogated whether different plasticity schedules could influence the development of segregated representations. This set of simulations was based on the observation that children learn to perform more simple tasks prior to learning more complex tasks (Halford, 1984, 1993; Robin & Holyoak, 1995), and further, that different areas of PFC tend to have different schedules of development (Shaw, Kabani, Lerch, Eckstrand, Lenroot, Gogtay, Greenstein, Clasen, Evans, Rapoport, Giedd, & Wise, 2008; Sowell, Thompson, Holmes, Jernigan, & Toga, 1999). The hypothesis is that if particular sets of PFC units are learning early, while the model is learning simple tasks, then those units would be pre-disposed to learn about the simple relationships in such tasks (e.g., between inner-loop cue-probe pairs). Once these tasks are learned, then if there are additional, non-committed units/stripes, then these non-committed units could be allocated to learning about additional sequential structure, namely about outer-loop task representations. In order to introduce this dynamic, the connections to and from the experimental group of PFC units did not learn until the less complex tasks achieved perfect performance for an epoch. Critically, the environment was constant over time, such that the network was always attempting to learn all 4 tasks, but early in training, only half of its PFC resources were available to learn (a subsequent investigation investigated the role that a dynamic environment may play, but did not influence the results). While this procedure is not particularly realistic, it serves as a proxy for the natural development of the brain and cognition, in which it appears that young children learn to solve relatively simple problems prior to solving more complex problems (Halford, 1984, 1993), and that this cognitive development somewhat tracks the development of the cerebral cortex (Bunge & Zelazo, 2006).

Simulation results

Training—The only network that influenced the training success or training time was the PFC-Hidden architecture (see Table 2). However, this difference in time to criterion appeared to be due largely to the number of stripes with access to hidden information, as a nonhierarchical network with only 2 stripes demonstrated similar time to criterion (325 epochs, $p = 0.9$).

Behavior—Across all networks, error rates during the testing phase were low, with the error rates in the 3 and 4 tasks being significantly lower than the error rates in the 1 and 2 tasks, $F(1, 377) = 877, p < 0.001; \eta_p^2 = 0.7$. This behavior did not fluctuate much as a function of the network architecture or plasticity schedule, as mean error rates in the 3-4 tasks were between 0.4% and 2.4%, and error rates in the 1-2 tasks were between 9.4% and 11.3%.

Selective Representations—As expected, the selectivity index in the baseline network was not significantly different from 0 for either the complex 1-2 tasks or the less complex 3-4 tasks (collapsed: $t(99) = 0.02$), indicating that there was no bias to develop representational specialization.

Constrained PFC-Hidden layer projections—In contrast to baseline, the PFC-Hidden network produced more segregated representations, such that the experimental group of units tended to code for outer-loop context more frequently than the control unit group, specifically when coding for information about the more complex, outer-loop tasks $t(81)=2.56, p=0.01$, see Figure 3. When compared directly to the control network, this network demonstrated a dissociation, such that the experimental group of units had a greater likelihood of representing outer context in the more complex tasks $F(1, 162)=4.84, p=0.03$.

Additionally, regression analyses indicated that there was a negative relationship between training time and the degree of representational segregation, specifically for the tasks requiring outer context $t(79)=-2.6, p=0.01, \Delta R^2=0.08$. This relationship was different than that observed in the baseline network, where the degree of representational segregation was not related to training time (significant network $\times$ index interaction $t(177)=3.6, p<0.001$).

Hierarchical PFC-BG projections—Similar to the constrained PFC model, there was a relatively high degree of representational segregation within the model with hierarchical projections between the PFC and the BG, particularly for the complex tasks $t(98)=3.0, p=0.003$, see Figure 3. When compared directly to the control network, this network demonstrated a dissociation in developed responses, with the areas of PFC capable of directly influencing the updating of other areas of PFC having a greater likelihood of representing outer context $F(1, 197)=4.5, p=0.03$. Similar to Simulation I, there was a negative relationship between time to train and the degree of representational segregation in the complex tasks, $t(96)=-2.8, p=0.006, R^2=0.07$, see Figure 4. Visual inspection and diagnostic measures revealed heteroskedasticity in the model (see Figure 4; significant Breusch-Pagan test: $\chi^2(1)=38, p<0.001$), but the use of a heteroskedasticity consistent error term (Long & Ervin, 2000) revealed the relationship was still reliable, even after taking this violation into consideration ($t(96)=-2.5, p=0.01$). No quadratic trends were significant. Similar to Simulation I, this negative relationship was significantly different from the relationship identified in the control network (condition $\times$ selectivity index interaction: $t(194)=3.8, p<0.001, \Delta R^2=0.07$). In this particular example, there was no overall increase in training time relative to the control network (see Table 2), so this is a pure benefit associated with segregated representation.

Plasticity manipulation—Contrary to the earlier simulations, the simulation interrogating the role of plasticity did not produce a difference in representational structure ($p=0.12$), nor was the representational structure measure associated with training time ($p=0.8$).

Additional simulations—Each of the 3 experimental manipulations could be considered orthogonal to one another; one can perform the plasticity manipulation on the network with constrained PFC connections or the network with hierarchical projections to the BG. When all of the possible models were run and included in a 2 (PFC-Hidden connectivity: either full or constrained) $\times$ 2 (PFC-BG connectivity: stripe specific vs. hierarchical) $\times$ 2 (plasticity: all learning vs. 1 delayed group) $\times$ 2 (complexity: low vs. high) ANOVA, the effects were additive and consistent with the simulations above (e.g. there was a complexity $\times$ PFC-Hidden connectivity interaction, and there was a complexity $\times$ PFC-BG connectivity interaction, but no three-way interaction, see Figure 3). The additive nature of the two significant effects indicates that both types of constraints may produce a similar tendency to segregate the nature of the representations learned, albeit by different underlying mechanisms. However, we only tested the additive nature of these effects when the organizational structure was completely overlapping. That is, we only interrogated the situation in which the most removed groups of PFC units are also the exact same units that can govern the updating of the intermediately located units. Future anatomical, empirical, and computational study will be needed to discern whether these effects occur when the overlap may not be quite so clear-cut.

Stable Representations—In addition to coding selective outer-loop representations, we also differentiated between selective, stable representations and selective, dynamic representations. Representations were considered stable if they were selective for a particular outer-loop context, and also did *NOT* vary as a function of the inner-loop cues and probes that occurred within that outer-loop context. Representations were considered dynamic if they did fluctuate as a function of the inner-loop information.

The first thing to note regarding this distinction is that it highlights two distinct encoding strategies the network could adopt. In one case, it could simply learn to encode direct mappings of the stimuli, and to maintain some representation of that stimulus information in a constant state while updating other, independent pieces of information. This strategy would be characterized by stable units. On the other hand, the network could adopt a more dynamic updating policy in which it encodes and updates each cue it encounters based on the outer loop information that is currently in memory. In this way, it could store a conjunction of the information regarding the cue and probe, and as such, would not necessarily need to maintain the outer-loop information in an isolated format. All networks predominantly utilized the latter pattern. For example, given a particular outer-level context within the baseline network, 85-95% of active units were task selective, depending on the task, which is consistent with values across other networks. However, less than 1% of these task-selective units were also stable. The percentage of stable units went up to approximately 3% for the hierarchical network if we specifically probe the tasks requiring outer loop contexts, but despite this increase, the predominant pattern of task encoding was that of a dynamic attractor, where outer-loop information was maintained across multiple patterns of activity that also reflected the task cue.

Discussion

The simulations presented here provide us with novel data that illustrate several points regarding PFC function and its role in the performance of sequential working memory tasks. Consistent with neuroanatomical data (Petrides & Pandya, 2007) and recent computational modeling work (Botvinick, 2007; Paine & Tani, 2005), anatomical constraints that impose a hierarchical structure within PFC resulted in segregated representational structure, with the most distal areas of PFC having higher likelihoods of representing outer-loop information. Additionally, anatomical constraints on the interactions between PFC and BG also produced such representational structure. While we also investigated the hypothesis regarding the utility of differential learning schedules, these manipulations resulted in minimal impact on the representational structure of the network. For those networks that promoted segregated representations, there was a unique benefit to training time, as those networks that had the strongest representational structure were also those that trained the fastest. Additionally, we probed the nature of the representations within these tasks, and demonstrated that all networks predominantly represented the task demands in a dynamic attractor in which each state was coded in terms of conjunctions between the outer loop context and most recent inner-loop cue.

Behavior

Each of the networks learned to perform the relevant tasks, and to perform such tasks, they relied upon the use of active maintenance processes subserved by interactions between the PFC and BG. Behavior during test was well above chance. However, there were clear differences in the ability of the networks to perform the two different types of tasks, as the networks were better at performing the tasks not requiring additional outer-loop information. This is not surprising given behavioral data from other paradigms, as there is substantial evidence that performance degrades as the number of dimensions an individual has to integrate increases (Christoff, Prabhakaran, Dorfman, Zhao, Kroger, Holyoak, & Gabrieli, 2001; Halford et al., 1999; Kroger, Sabb, Fales, Bookheimer, Cohen, & Holyoak, 2002).

Architectural constraints—Imposing an architectural constraint on the connectivity of PFC clearly had an effect on the types of representations that were developed and the degree to which the different types of representations were segregated across different tasks. The experimental networks produced stable representations of the outer loop information more frequently (albeit, still a very small overall proportion), and tended to place such representations in the more remote areas of PFC. Further, it appeared as though such segregation was associated with faster learning within the context of the manipulated networks. This finding is similar to findings in the hierarchical reinforcement learning domain, particularly the options framework discussed below (Botvinick, Niv, & Barto, 2008; Sutton, Precup, & Singh, 1999).

It should be pointed out that the relationship between training time and selectivity index was not symmetric around 0 for either of the networks that demonstrated such a relationship. If the key to fast training was representation segregation (meaning that outer context representations tended to be represented in one particular group while not being represented in the other), then the relationship between this score and time to criterion should be symmetric around 0. However, it is very clear that this is not the case (adding an additional squared-term to the linear model accounted for no additional variance, $F(1, 95) < 1$). In these networks, it clearly matters what representation is placed in what set of stripes. It is useful for the network to place outer-loop context in the remote areas of PFC because the extra step of processing insulates it from updating signals originating from the input, and thus makes it less likely to update (and see this in that these networks have a higher chance of producing stable representations). It is useful for the network to place the outer context information in the areas of PFC that provide direct input to the updating mechanism of the other stripes, because it can then provide more direct input on the updating of the other stripes, and in such a way, mirror the demands of the task. Alternately, if outer context is placed in the areas that are being modulated by other areas, then the gain from the ability to scope the updating signal is somewhat lost, as these are relatively stable representations (as far as the task demands are concerned). Likewise, if outer-loop information is placed in the proximal areas of PFC, then the remote areas of PFC are unable to be updated with the relevant, dynamically changing inner-loop information, because they only have access to the static outer-loop information.

Plasticity manipulation—This plasticity manipulation did not result in segregated, hierarchical representations. However, there are a number of potential reasons for this null effect. Based on our previous control simulations, it is not entirely surprising that there was no bias to segregate representations, as within that particular architecture, there is no computational pressure to form such types of representations. For the architecture used, it provides no additional benefit in terms of learning or performance, particularly when compared to the alternate networks. That being said, the additional models run were not influenced by the developmental training procedure, suggesting that what computational pressure may exist in the other model architectures did not provide the developmental procedure a way to bootstrap stronger segregated representations.

Relationship to Hierarchical Reinforcement Learning approaches and other Computational Models of Hierarchy

Other researchers have utilized reinforcement learning to approach the question of hierarchy (for reviews, see Botvinick, 2007 and Botvinick et al., 2008). We focus on the relationship between our approach and an alternate approach in which temporally abstracted actions, referred to as *options*, facilitate learning by providing subgoals that can be attained prior to the achievement of some more distant reward (Sutton et al., 1999). At first glance, the approaches are quite similar in that they both use reinforcement learning techniques in order to accomplish temporally distant goals. As such, the current set of simulations, and the PBWM framework

more generally, could be recast in the options framework, such that maintained outer-loop context could be considered an identifier of a particular *option* in that it has a relatively long-term outcome, and is being used to make more myopic decisions on a trial by trial basis. The maintenance of such information and its influence on ongoing processing would then be analogous to the utilization of such an option's policy (Botvinick et al., 2008).

However, a very clear difference between the two approaches arises when investigating the nature and discovery of component actions (also called subgoals). One problem that the *options* approach encounters is that the subgoals need to be identified and compiled, typically prior to task learning. A number of mechanisms for doing such identification have been suggested, ranging from analyzing paths through problem space for relevant statistical structure (Pickett & Barto, 2002), to using intrinsic motivation as a possible mechanism (Singh, Barto, & Chentanez, 2005). This is a problem that, in many ways, the PBWM system has been able to solve (at least in the restricted environments that it has been placed). Within the PBWM framework, the appropriate options and subgoal states are learned via the mechanism that governs the updating system of PFC (O'Reilly & Frank, 2006); to the extent that the network discovers and segregates the types of information relevant for long-term performance, then it could be said that appropriate subgoals and options are acquired. Without segregation of the long-term information, it is analogous to a model attempting to solve a complex, temporally extended problem with RL, without the benefit of the options framework. This interpretation is consistent with the identified relationship between training time and degree of representational segregation.

Previous models exploring how hierarchical representations may develop (Botvinick, 2007; Paine & Tani, 2005) have produced consistent results with those produced here. Namely, architectural constraints tend to be sufficient to bias the network into developing representations that segregated as a function of content and temporal abstraction. This consistency is in spite of huge differences in the actual tasks and algorithms used to investigate the question. Whereas previous approaches have varied between the purely computational approach of using genetic algorithms (Paine & Tani, 2005) and recurrent backpropagation (Botvinick, 2007), the current approach is an advance in that it uses a biologically inspired algorithm to ask similar questions, and demonstrates that hierarchical representations can be selected using such an algorithm. In addition to these models, other models have been used to interrogate hierarchy without regards to the development or origin of the representations. Specifically, previous models have focused on understanding and distinguishing the potential computational roles of various areas of PFC (Koechlin et al., 2003; Koechlin & Hyafil, 2007; Koechlin & Summerfield, 2007; Reynolds & Mozer, 2009). The approach utilized by Koechlin argues that different areas of PFC are governed by different types of information conveyed by stimuli, and that such information may be understood in the context of a hierarchy, whereas the approach taken by Dayan (2008) and Reynolds and Mozer (2009) argues that no explicit hierarchy is needed to elicit hierarchical behavior. Although these approaches do not focus on how such representations may develop, they provide constraints for future models – the ability to develop representations consistent with one or another framework could prove to be a powerful tool in discriminating between the various models.

Stable Representations and conjunctive representations

The number of units that coded for conjunctions of various stimuli (e.g. A in the context of a 1) was quite large, and it was interesting that there was such a strong bias for the network to develop such representations. Although somewhat surprising, it is consistent with a range of neurophysiology data in which conjunctive codes have been identified in a number of paradigms, including working memory (Barone & Joseph, 1990; Rao, Rainer, & Miller, 1997), and task-switching paradigms (Wallis, Anderson, & Miller, 2001). It is possible that

the development of more stable, abstract representations requires more extensive training of a particular sort (Rougier, Noelle, Braver, Cohen, & O'Reilly, 2005), and that such training allows for task-independent representations that can be utilized when learning about more complex tasks. It is quite possible that the tasks used in the current set of simulations is not ideal for generating or creating hierarchical representations; specifically, there is not notion of a particular sequence of behaviours that can be learned, “chunked,” and then applied in a novel situation (see Botvinick & Plaut, 2004 and Reynolds, Zacks, & Braver, 2007 for alternate paradigms with exactly such a structure). As such, it is somewhat surprising that these tasks and constraints produced a measurable difference in the representational structure of the network at all. Further investigation will have to be performed in order to determine whether such true abstractions (either in the form of sequences or in the form of dimensional extractions) produce more stable or more specialized representations.

Despite this endeavor being in its infancy, it makes a strong prediction regarding the role of different areas of PFC in the performance of this sequential working memory task and novel variants. First, it suggests that this task may differ substantially from other empirical tasks used to prove the hierarchical organization of behavior (Badre & D'Esposito, 2008; Koechlin et al., 2003). The relatively small number of units that were statically coding for outer-loop context, combined with the large number of units that were multiplexing outer-loop information and inner-loop information suggests that one may find much stronger dynamics on trial-by-trial effects, rather than in sustained maintenance across trials (Reynolds, 2005). Second, this bias towards dynamic attractors begs the question of what could increase the probability of forming a more static attractor? Although there are many possibilities, it seems like one potentially important and relevant variable is the timing between the inner-loop cue and probe information. Counter-intuitively, if you reduce this delay, there may not be enough time to re-encode the inner-loop cue as a conjunction with the outer-loop context. As such, manipulating this timing parameter may be sufficient to increase the maintenance of the outer-loop stimulus, and may account for differences in empirical investigations of this task (Reynolds, 2005) and other tasks where the outer-most piece of information is the only item presented prior to a multidimensional imperative stimulus (Badre & D'Esposito, 2008; Koechlin et al., 2003). Ongoing simulations and experiments are currently investigating this question.

Although the developmental plasticity hypothesis did not play out as predicted, there is still growing evidence that different regions of PFC mature at different rates (Brown, Lugar, Coalson, Miezin, Petersen, & Schlaggar, 2005; Shaw et al., 2008; Sowell et al., 1999), and there is likely some functional consequence of this (Bunge & Zelazo, 2006). While previous experimental approaches have suggested that there is a posterior-anterior gradient in terms of either rule complexity or rule abstraction (Badre & D'Esposito, 2008), recent data have suggested that the PFC may not develop in a strict posterior-to- anterior gradient. Specifically, it appears that posterior and anterior areas of PFC are the first to develop, with the areas in between developing later (Gogtay et al., 2004; Shaw et al., 2008). Additional investigation will be needed to understand how the differences in maturation schedule influence the nature of representation, and how these developmental trajectories may relate to current theories of the organization of PFC.

Despite our best intentions to concretely define what is meant by the maintenance of outer-loop information, there are several different variables that may influence whether information is considered outer-loop by the cognitive and neural system. For example, there are at least two differences between outer-loop information and inner-loop information in the current model, both of which are consistent with the notion of anterior PFC playing a high-level role within a hierarchy of goals. The first is that outer-loop information must be maintained while something else is updated. The second is that outer-loop information must be maintained for longer periods of time than inner-loop information. Finally, it is very possible that the cognitive

system considers information to be outer- loop only when both of these constraints are met. The nature of the distinction between outer- and inner-loop information is a key question that is being investigated currently. Despite the potential concerns and limitations, the current set of simulations has allowed for a concrete, implemented definition regarding one particular dimension across which PFC may be organized, and provided a biologically inspired mechanism by which such an organization could develop.

The current set of simulations did not interrogate all of the potential dimensions of PFC organization. Rather, the current set of simulations should be viewed as an early step at using computational modeling techniques in order to probe the nature of such representations and differences between them. The current models could be extended in numerous ways to capture other potential dimensions. For example, there is growing evidence suggesting that left inferior PFC can be subdivided in to regions that are differentially sensitive to semantic or phonological properties of stimuli (with more anterior areas of left inferior PFC being more closely associated with semantics; Poldrack, Wagner, Prull, Desmond, Glover, & Gabrieli, 1999). This could be viewed along a similar axis of abstraction as that investigated here, such that the fine grained muscle movements necessary to create utterances are required before verbal communication can develop. These muscle representations are likely used in the perception and comprehension of language (e.g., phonology; Rizzolatti & Arbib, 1998; Rizzolatti, Fadiga, Gallese, & Fogassi, 1996). Once such a basis of movements is learned and used to, novel combinations of such movements can be generated and associated with an abstraction that represents the semantics of such a motion in a many-to-many mapping.

Conclusions

The current set of simulations provides convergent evidence that different types of representations within PFC can be developed by having specific kinds of architectural constraints. The segregation of these representations leads to faster learning, and as such, may provide some evidence for the computational pressures that govern the organization of PFC.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

The authors would like to thank Nicole Speer, Thomas Hazy, Seth Herd, and the rest of the Computational Cognitive Neuroscience laboratory for helpful comments and suggestions. This research was supported in part by an NRSA post-doctoral training grant (1 F32 MH075300-01A2).

References

- Badre D, D'Esposito M. Functional magnetic resonance imaging evidence for a hierarchical organization of the prefrontal cortex. *Journal of cognitive neuroscience* 2008;19
- Barone P, Joseph JP. Prefrontal cortex and spatial sequencing in macaque monkey. *Experimental brain research* 1990;78:447–464.
- Botvinick M, Niv Y, Barto AC. Hierarchically organized behavior and its neural foundations: A reinforcement learning perspective. *Cognition*. 2008
- Botvinick M, Plaut DC. Doing without schema hierarchies: a recurrent connectionist approach to normal and impaired routine sequential action. *Psychological review* 2004;111:395–429. [PubMed: 15065915]
- Botvinick MM. Multilevel structure in behaviour and in the brain: a model of fuster's hierarchy. *Philos Trans R Soc Lond B Biol Sci* 2007;362(1485):1615–1626. [PubMed: 17428777]
- Botvinick MM. Hierarchical models of behavior and prefrontal function. *Trends in cognitive sciences* 2008;12(11)

- Braver TS, Barch DM, Keys BA, Carter CS, Cohen JD, Kaye JA, Janowsky JS, Taylor SF, Yesavage JA, Mumuthaler MS, Jagust WJ, Reed BR. Context processing in older adults: evidence for a theory relating cognitive control to neurobiology in healthy aging. *Journal of experimental psychology* 2001;130:746–763. [PubMed: 11757878]
- Braver TS, Bongiolatti SR. The role of frontopolar cortex in subgoal processing during working memory. *NeuroImage* 2002;15:523–536. [PubMed: 11848695]
- Braver, TS.; Cohen, JD.; Barch, DM. The role of the prefrontal cortex in normal and disordered cognitive control: A cognitive neuroscience perspective. In: Stuss, DT.; Knight, RT., editors. *Principles of frontal lobe function*. Oxford; Oxford University Press; 2002. p. 428–448.
- Braver TS, Reynolds JR, Donaldson DI. Neural mechanisms of transient and sustained cognitive control during task switching. *Neuron* 2003;39:713–726. [PubMed: 12925284]
- Brown TT, Lugar HM, Coalson RS, Miezin FM, Petersen SE, Schlaggar BL. Developmental changes in human cerebral functional organization for word generation. *Cerebral cortex* (New York, N 2005;15:275–290.
- Bunge SA, Zelazo PD. A brain-based account of the development of rule use in childhood. *Current Directions in Psychological Science* 2006;14(3):118–121.
- Christoff K, Gabrieli JDE. The frontopolar cortex and human cognition: Evidence for a rostrocaudal hierarchical organization within the human prefrontal cortex. *Psychobiology* 2000;28:168–186.
- Christoff K, Prabhakaran V, Dorfman J, Zhao Z, Kroger JK, Holyoak KJ, Gabrieli JD. Rostrolateral prefrontal cortex involvement in relational integration during reasoning. *NeuroImage* 2001;14(5): 1136–4119. [PubMed: 11697945]
- Christoff K, Ream JM, Geddes LPT, Gabrieli JDE. Evaluating self-generated information: anterior prefrontal contributions to human cognition. *Behavioral neuroscience* 2003;117(6):1161–1168. [PubMed: 14674837]
- Dayan P. Bilinearity, rules, and prefrontal cortex. *Frontiers in computational neuroscience* 2008;1(1)
- D'Esposito M, Postle BR, Ballard D, Lease J. Maintenance versus manipulation of information held in working memory: an event-related fmri study. *Brain and cognition* 1999;41:66–86. [PubMed: 10536086]
- Duncan J, Owen AM. Common regions of the human frontal lobe recruited by diverse cognitive demands. *Trends in neurosciences* 2000;23:475–482. [PubMed: 11006464]
- Frank MJ, Loughry B, O'Reilly RC. Interactions between the frontal cortex and basal ganglia in working memory: A computational model. *Cognitive, Affective, and Behavioral Neuroscience* 2001;1:137–160.
- Fuster, JM. Prefrontal cortex and the bridging of temporal gaps in the perception-action cycle. In: Diamond, A., editor. *The development and neural bases of higher cognitive functions*. Vol. 608. New York: New York Academy of Science Press; 1991. p. 318–336.
- Fuster JM. Upper processing stages of the perception-action cycle. *Trends in cognitive sciences* 2004;8 (4):143–145. [PubMed: 15551481]
- Gogtay N, Giedd JN, Lusk L, Hayashi KM, Greenstein D, Vaituzis AC, Nugent TF, Herman DH, Clasen LS, Toga AW, Rapoport JL, Thompson PM. Dynamic mapping of human cortical development during childhood through early adulthood. *Proceedings of the National Academy of Sciences of the United States of America* 2004;101(21):8174–8179. [PubMed: 15148381]
- Goldman-Rakic PS. Circuitry of primate prefrontal cortex and regulation of behavior by representational memory. *Handbook of Physiology — The Nervous System* 1987;5:373–417.
- Halford GS. Can young children integrate premises in transitivity and serial order tasks? *Cognitive Psychology* 1984;16:65–93.
- Halford, GS. *Children's understanding: The development of mental models*. Hillsdale, NJ: Erlbaum; 1993.
- Halford GS, Wilson WH, Phillips S. Processing capacity defined by relational complexity: implications for comparative, developmental, and cognitive psychology. *The Behavioral and brain sciences* 1999;21:803. [PubMed: 10191879]
- Haxby JV, Petit L, Ungerleider LG, Courtney SM. Distinguishing the functional roles of multiple regions in distributed neural systems for visual working memory. *NeuroImage* 2000;11:380–391. [PubMed: 10806025]

- Johnson MK, Raye CL, Mitchell KJ, Greene EJ, Anderson AW. Fmri evidence for an organization of prefrontal cortex by both type of process and type of information. *Cerebral cortex* (New York, N 2003;13:265–273.
- Koechlin E, Basso G, Pietrini P, Panzer S, Grafman J. The role of the anterior prefrontal cortex in human cognition. *Nature* 1999;399:148–151. [PubMed: 10335843]
- Koechlin E, Hyafil A. Anterior prefrontal function and the limits of human decision-making. *Science* (New York, N 2007;318(5850):594–598.
- Koechlin E, Ody C, Kouneiher F. Neuroscience: The architecture of cognitive control in the human prefrontal cortex. *Science* 2003;424:1181–1184. [PubMed: 14615530]
- Koechlin E, Summerfield C. An information theoretical approach to prefrontal executive function. *Trends in cognitive sciences* 2007;11(6):229–235. [PubMed: 17475536]
- Kroger JK, Sabb FW, Fales CL, Bookheimer SY, Cohen MS, Holyoak KJ. Recruitment of anterior dorsolateral prefrontal cortex in human reasoning: a parametric study of relational complexity. *Cerebral cortex* (New York, N 2002;12:477–485.
- Lashley, KS. The problem of serial order in behavior. In: Jeffress, LA., editor. *Cerebral mechanisms in behavior: the hixon symposium*. New York: Wiley; 1951. p. 112-136.
- Long JS, Ervin LH. Using heteroscedasticity consistent standard errors in the linear regression model. *The American Statistician* 2000;54:217–224.
- Miller EK. The prefrontal cortex: No simple matter. *Neuroimage* 2000;11:447–450. [PubMed: 10806030]
- Miller EK, Cohen JD. An integrative theory of prefrontal cortex function. *Annual review of neuroscience* 2001;24:167–202.
- Miller, GA.; Galanter, E.; Pribram, KH. *Plans and the structure of behavior*. New York: Holt; 1960.
- Nystrom LE, Braver TS, Sabb FW, Delgado MR, Noll DC, Cohen JD. Working memory for letters, shapes, and locations: fmri evidence against stimulus-based regional organization in human prefrontal cortex. *NeuroImage* 2000;11:424–446. [PubMed: 10806029]
- O'Donnell S, Noseworth M, Levine B, Dennis M. Cortical thickness of the frontopolar area in typically developing children and adolescents. *NeuroImage* 2005;24(4):948–954. [PubMed: 15670671]
- O'Reilly RC. Biologically plausible error-driven learning using local activation differences: The generalized recirculation algorithm. *Neural Computation* 1996;8(5):895–938.
- O'Reilly RC. Six principles for biologically-based computational models of cortical cognition. *Trends in Cognitive Sciences* 1998;2(11):455–462.
- 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, Frank MJ. Making working memory work: A computational model of learning in the prefrontal 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*. Cambridge, MA: The MIT Press; 2000.
- Paine RW, Tani J. How hierarchical control self-organizes in artificial adaptive systems. *Adaptive Behavior* 2005;13(3):211–225.10.1177/105971230501300303
- Petrides M. The role of the mid-dorsolateral prefrontal cortex in working memory. *Experimental brain research* 2000;133:44.
- Petrides M, Pandya DN. Efferent association pathways from the rostral prefrontal cortex in the macaque monkey. *The Journal of neuroscience* 2007;27(43):11573–11586. [PubMed: 17959800]
- Pickett, M.; Barto, A. Policyblocks: An algorithm for creating useful macro-actions in reinforcement learning. *ICML '02: Proceedings of the Nineteenth International Conference on Machine Learning*; San Francisco, CA, USA: Morgan Kaufmann Publishers Inc; 2002. p. 506-513.
- Poldrack RA, Wagner AD, Prull MW, Desmond JE, Glover GH, Gabrieli JD. Functional specialization for semantic and phonological processing in the left inferior prefrontal cortex. *NeuroImage* 1999;10:15–35. [PubMed: 10385578]
- Rao SC, Rainer G, Miller EK. Integration of what and where in the primate prefrontal cortex. *Science* 1997;276(5313):821–824. [PubMed: 9115211]

- Raye CL, Johnson MK, Mitchell KJ, Reeder JA, Greene EJ. Neuroimaging a single thought: dorsolateral pfc activity associated with refreshing just-activated information. *NeuroImage* 2002;15:447–453. [PubMed: 11798278]
- Reynolds J, West R, Braver T. Distinct neural circuits support transient and sustained processes in prospective memory and working memory. *Cerebral cortex* (New York, N. 2008)10.1093/cercor/bhn164
- Reynolds J, Zacks J, Braver T. A computational model of event segmentation from perceptual prediction. *Cognitive Science* 2007;31:613–634.
- Reynolds, JR. PhD thesis. Washington University in St. Louis; 2005. On the roles of duration and computational complexity in the recruitment of frontopolar prefrontal cortex.
- Reynolds JR, Mozer MC. Temporal dynamics of cognitive control. *Advances in neural information processing systems*. 2009
- Rizzolatti G, Arbib MA. Language within our grasp. *Trends in Neurosciences* 1998;21:188–194. [PubMed: 9610880]
- Rizzolatti G, Fadiga L, Gallese V, Fogassi L. Premotor cortex and the recognition of motor actions. *Brain Res Cogn Brain Research* 1996;3:131–141. [PubMed: 8713554]
- Robin, N.; Holyoak, KJ. Relational complexity and the functions of prefrontal cortex. In: Gazzaniga, MS., editor. *The cognitive neurosciences*. Vol. 1. Cambridge, MA: MIT Press; 1995. p. 987–997.
- Rougier NP, Noelle D, Braver TS, Cohen JD, O'Reilly RC. Prefrontal cortex and the flexibility of cognitive control: Rules without symbols. *Proceedings of the National Academy of Sciences* 2005;102(20):7338–7343.
- Shaw P, Kabani N, Lerch J, Eckstrand K, Lenroot R, Gogtay N, Greenstein D, Clasen L, Evans A, Rapoport J, Giedd J, Wise S. Neurodevelopmental trajectories of the human cerebral cortex. *J* 2008;28:3586–3594.
- Singh, S.; Barto, A.; Chentanez, N. Intrinsically motivated reinforcement learning. In: Saul, LK.; W, Y.; B, L., editors. *Advances in neural information processing systems 17: Proceedings of the 2004 conference*. Cambridge: MIT Press; 2005. p. 1281–1288.
- Sowell ER, Thompson PM, Holmes CJ, Jernigan TL, Toga AW. In vivo evidence for post-adolescent brain maturation in frontal and striatal regions. *Nature neuroscience* 1999;2(10):859–861.
- Stuss, DT.; Knight, RT., editors. *Principles of frontal lobe function*. New York, New York: Oxford University Press; 2002.
- Sutton R, Precup D, Singh S. Between mdps and semi-mdps: A framework for temporal abstraction in reinforcement learning. *Artificial Intelligence* 1999;112(12):181–211.
- Wallis JD, Anderson KC, Miller EK. Single neurons in prefrontal cortex encode abstract rules. *Nature* 2001;411:953–956. [PubMed: 11418860]
- Wilson FA, Scalaidhe SP, Goldman-Rakic PS. Dissociation of object and spatial processing domains in primate prefrontal cortex. *Science* (New York, N 1993;260:1955–1957.
- Wood JN, Grafman J. Human prefrontal cortex: processing and representational perspectives. *Nature reviews* 2003;4:139–147.

Figure 1.

The 1-2 AX-CPT task. Stimuli are presented one at a time in sequence. The participant responds by pressing the right key (R) to the target sequence; otherwise, a left key (L) is pressed. If the participant last saw a 1, then the target sequence is an A followed by an X. If a 2 was last seen, then the target is a B followed by a Y. The maintenance of the task stimuli (1 or 2) constitutes a temporal outer loop around multiple inner-loop memory updates required to detect the target sequence.

Figure 2.

Network architectures for the simulations explored. Panel A reflects the overall structure of the network and its connections. The dashed box is expanded in panels B-D. Panel B reflects the structure of a typical network in this framework, with each PFC stripe governing the the updating of itself. Panel C reflects the implementation of the first hypothesis. In this architecture, there is a different connectivity pattern associated with PFC; a subset of PFC units receive direct input from and provide direct output to hidden and output layers, while the rest of PFC only receives information from other areas of PFC. Panel D reflects the second hypothesis, in which the updating of one set of PFC units is governed both by its own activity as well as the activity of the other PFC units. "Control" and "Exp." in panels B-D refer to the groups associated with the values in Equation 1.

Figure 3.

Representational Segregation for each type of task (low complexity = 3-4 tasks, high complexity = 1-2 tasks). Each set of bars corresponds to the difference between the control stripes and the experimental stripes for a different architecture or training regime. Dark bars correspond to the differences in type of representation for the less complex 3-4 tasks, whereas light bars correspond to differences in representation for the more complex 1-2 tasks. The first 3 groups correspond to the three different architectures presented in Figure 2. The fourth set of bars correspond to the results when the Control Network is placed under a condition with developmental plasticity, such that the experimental stripes of PFC have a later onset of plasticity, and the fifth set of bars correspond to the results when both architectural constraints are included in the same model.

Figure 4. Relationship between representational specialization and training time for the PFC-BG network. An increased selectivity index was associated with faster training times. A similar pattern existed for the PFC-Hidden network, but did not exist for the baseline or development networks.

Table 1

The task paradigm performed by the model. The first row indicates each of the 4 outer-loop contexts (i.e. tasks). The Second row indicates the inner-loop context cue nested within that task. For example, A & B in the last 2 columns appear in the context of the “2” outer-loop context, whereas R & S in the first two columns appear in the context of the “3” outer-loop context. The first column refers to the probes that may appear within each task. Within the table, empty fields mean those stimulus pairings do not appear, whereas T means that a target response is required on that combination of outer-loop, inner-loop, and probe and NT means that a non-target response is required. Numbers within parentheses refer to the relative percentage of the sequences within that particular outer-loop context (1, 2, 3 or 4). As can be seen from the table, the outer context is unnecessary in the 3 and 4 contexts, as an appropriate response can be determined purely from the inner cue and the probe.

		3		4		1		2	
		R	S	M	N	A	B	A	B
O		T (30%)	NT (20%)						
P		NT (20%)	NT (30%)						
Q				NT (30%)	NT (20%)				
T				NT (20%)	T (30%)				
X						T (30%)	NT (20%)	NT (30%)	NT(20%)
Y						NT (20%)	NT (30%)	NT(20%)	T (30%)

Table 2

Training Information. The only network that produced different training statistics from the others was the network with constrained projections between the PFC and Hidden Layers(C). This network converged less frequently, $\chi^2(1)=18$, $p < 0.001$, and took longer to train when it successfully reached criterion, $W=1818$, $p < 0.001$. Letters in parentheses correspond to their component structure in Figure 2.

proportion trained	median time to criterion
Base(B)100	66
PFC-Hidden(C)82***	296***
PFC-BG(D)99	92
Plasticity100	77.5