

Dissociable Frontal–Striatal and Frontal–Parietal Networks Involved in Updating Hierarchical Contexts in Working Memory

Derek Evan Nee^{1,2} and Joshua W. Brown¹

¹Department of Psychological and Brain Sciences, Indiana University, Bloomington, IN 47405, USA ²Department of Psychology, University of Michigan, Ann Arbor, MI 48109, USA

Address correspondence to Derek Evan Nee, 1101 E Tenth Street, Bloomington, IN 47405, USA. Email: denec@indiana.edu

Recent theories propose that the prefrontal cortex (PFC) is organized in a hierarchical fashion with more abstract, higher level information represented in anterior regions and more concrete, lower level information represented in posterior regions. This hierarchical organization affords flexible adjustments of action plans based on the context. Computational models suggest that such hierarchical organization in the PFC is achieved through interactions with the basal ganglia (BG) wherein the BG gate relevant contexts into the PFC. Here, we tested this proposal using functional magnetic resonance imaging (fMRI). Participants were scanned while updating working memory (WM) with 2 levels of hierarchical contexts. Consistent with PFC abstraction proposals, higher level context updates involved anterior portions of the PFC (BA 46), whereas lower level context updates involved posterior portions of the PFC (BA 6). Computational models were only partially supported as the BG were sensitive to higher, but not lower level context updates. The posterior parietal cortex (PPC) showed the opposite pattern. Analyses examining changes in functional connectivity confirmed dissociable roles of the anterior PFC–BG during higher level context updates and posterior PFC–PPC during lower level context updates. These results suggest that hierarchical contexts are organized by distinct frontal–striatal and frontal–parietal networks.

Keywords: basal ganglia, executive control, fMRI, prefrontal cortex, striatum

Introduction

A hallmark of intelligent behavior is the ability to flexibly adjust action plans based upon the context. For example, whereas one might typically turn right at an intersection to drive home, one might instead turn left if intending to first pick up groceries. Such rule-based behavior depends on a balance between stable maintenance of contexts and flexible updating. Excessive stability can lead to perseveration, while excessive flexibility can lead to distractibility. Mounting evidence suggests that the balance of stability and flexibility involves interactions between the prefrontal cortex (PFC) and basal ganglia (BG). Whereas the PFC is thought to help maintain representations of contexts and protect them from distraction, the BG has been proposed to act as a gate that affords flexible updating of PFC representations (Braver and Cohen 2000; Frank et al. 2001; Rougier et al. 2005; O'Reilly and Frank 2006; Reynolds and O'Reilly 2009). Hence, PFC–BG interactions may be critical for the representation of appropriate contexts online in working memory (WM) to guide flexible behavior.

A growing number of theories propose that the representations and/or processes instantiated by the PFC are

hierarchically organized along a rostral–caudal axis (Fuster 1990; Badre 2008; Botvinick 2008; Badre and D'Esposito 2009; O'Reilly 2010). According to such frameworks, more anterior portions of the PFC maintain more abstract, higher level contexts, whereas more posterior portions of the PFC maintain more concrete, lower level contexts. A number of recent functional magnetic resonance imaging (fMRI) studies have demonstrated such abstraction gradients in the PFC (Koechlin et al. 2003; Koechlin and Jubault 2006; Badre and D'Esposito 2007; Kounieher et al. 2009; Badre et al. 2010), although proposals differ somewhat with regard to what is abstracted (see Badre 2008; Botvinick 2008 for reviews).

While PFC–BG interactions may be critical for updating contexts, it is unclear whether this holds true across different levels of abstraction. Some have suggested that updating hierarchically dependent representations within the PFC depends on distinct “stripes” in the PFC and BG (Frank et al. 2001; O'Reilly and Frank 2006; Reynolds and O'Reilly 2009). In this framework, separate “stripes” maintain and update distinct items or contexts in WM. Through this organization, an item/context in WM can be updated without disrupting other information stored in WM. For example, this would allow the maintenance and updating of nested levels of context. These authors have not made explicit whether such “stripes” reflect different PFC/BG sub-regions or different cell populations within a given particular sub-region. Given the evidence of hierarchical structure in the PFC along distinct rostral–caudal sub-regions, it may be that the PFC “stripes” indeed reflect different sub-regions of the PFC. Moreover, different sub-regions of the PFC are anatomically connected to distinct sub-regions of the BG (Alexander et al. 1986) and recent proposals suggest that there may be abstraction gradients in the BG similar to that of the PFC (Cools 2006). Hence, it is possible that updating of hierarchically structured contexts depends on separable PFC–BG loops.

Here, we examined the neural correlates of updating hierarchically structured contexts in WM. We used a variant of the AX continuous performance task (AX-CPT) that has served as a model for hierarchical updating in WM (Frank et al. 2001; O'Reilly and Frank 2006). The task required responding to stimuli based upon a hierarchical series of cues stored in WM. We used fMRI to examine neural responses, while subjects updated WM with higher and lower level context cues. Based on previous research, we anticipated that the PFC would be sensitive to the level of abstraction of the information updated, but it was less clear how the BG would respond. We hypothesized that the BG could either 1) vary by level of abstraction such that distinct BG sub-regions were sensitive to updating different levels of representation, 2) respond to different levels of abstraction without regional distinction, or

3) respond to only particular levels of abstraction. Hence, these data served to empirically test computational models of WM (Frank et al. 2001; O'Reilly and Frank 2006), and illuminate the mechanisms of flexible behavior more generally.

Materials and Methods

Participants

Twenty-one (11 females) right-handed native English speakers with normal or corrected-to-normal vision participated in the experiment (mean age 23.7 years; range 21–32). Informed consent was obtained in accordance with the Institutional Review Board at Indiana University. Subjects were compensated at a rate of \$20 per hour for participation plus a performance based bonus (mean \$2.43; range \$1.24–3.76).

Procedure

The task is depicted in Figure 1. Subjects performed a variant of the AX-CPT (Servan-Schreiber et al. 1996; MacDonald 2008; Barch et al. 2009) referred to as the 1–2-AX-CPT (Frank et al. 2001; O'Reilly and Frank 2006). The 1–2-AX-CPT requires subjects to hold 2 levels of contexts in mind in order to make responses. These contexts are hierarchical, forming higher and lower level action rules (sometimes referred to as the outer and inner loops). In this task, subjects observed a series of visually presented digits and letters and made responses to the letters “X” and “Y”. Responses to these letters were based on a hierarchical digit–letter sequence. Under the “1” context, subjects made a target response to the letter “X” if it was preceded by the letter “A” and made a non-target response otherwise. Under the “2” context, subjects made a target response to the letter “Y” if it was preceded by the letter “B” and made a non-target response otherwise. Hence, subjects had to keep in WM both a higher level context (“1” or “2”) and a lower level context (“A” or “B”) to determine how to respond to the letters “X” and “Y”.

For the sake of exposition, a sequence of “A/B” (lower level context) followed by “X/Y” (response) was considered a trial. On one-third of the trials, the digit “1” or “2” was presented prior to the lower level context, thereby requiring an update of the higher level context. The digits alternated so that each digit presentation required an update (i.e. the same higher level context never repeated twice in a row). Stimuli within a trial were ordered. Higher level context cues, when they occurred, always followed a response and preceded lower level context cues. Lower level context cues were always separated by

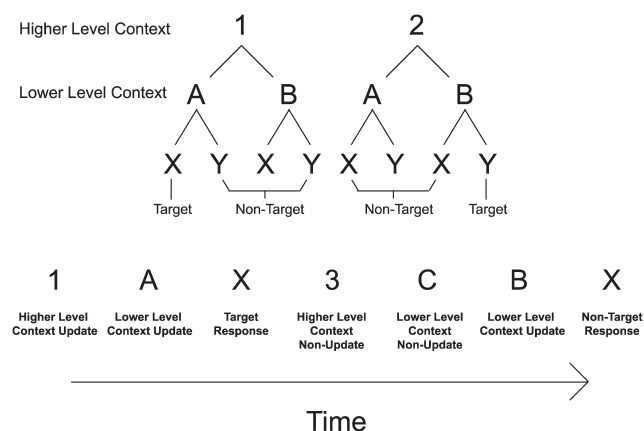


Figure 1. Hierarchical dependencies and task design. Subjects made responses to the letters “X” and “Y” based upon a hierarchical digit–letter context. In the “1” context, subjects made a target response to the letter “X” if it was preceded by the letter “A”, and made a non-target response otherwise. In the “2” context, subjects made a target response to the letter “Y” if it was preceded by the letter “B”, and made a non-target response otherwise. Hence, responses depended on both a higher level (“1” or “2”) and lower level (“A” or “B”) context. Irrelevant stimuli (e.g. “3, C”) were inter-mixed to provide controls for higher and lower level context updating. Subjects were instructed to withhold these items from WM.

a response. Consecutive responses never occurred without an intervening context cue.

In order to isolate activation related to updating higher and lower level contexts, irrelevant stimuli were presented that subjects were instructed to ignore (i.e. withhold from WM). On one-third of the trials, the digit “3” was presented during the interval that “1” or “2” normally appeared. On one-third of the trials, the letter “C” was presented during the interval that “A” or “B” normally appeared. On one-third of the trials, the letter “Z” was presented during the interval that “X” or “Y” normally appeared. Hence, contrasts of neural responses to “1/2” versus “3” isolated updating of the higher level context. Contrasts of neural responses to “A/B” versus “C” isolated updating of the lower level context. Contrasts of neural responses to “X/Y” versus “Z” isolated the motor response. No more than one irrelevant stimulus was presented in-between relevant stimuli on a given trial (e.g. consecutive presentations of “C” never occurred).

Responses were made with the index finger of either hand with the target hand counter-balanced between subjects. All relevant digits (“1/2”) and letters (“A/B, X/Y”) appeared in equal proportions throughout the experiment. Each stimulus was presented for 1 s. Jittered 4–6 s intervals (increments of 1 s) separated successive letter stimuli and each digit stimulus was preceded and followed by a 10-s interval. These intervals afforded isolation of the hemodynamic response corresponding to each stimulus. The lengthened intervals surrounding digit stimuli were designed for multi-variate analyses not relevant for present purposes. Subjects completed 6 runs of 18 trials each while being scanned. The session consisted of 36 higher level update cues, 36 higher level non-update cues, 108 lower level update cues, 36 lower level non-update cues, 108 response events, and 36 no-go events. Within a week prior to scanning, subjects completed a full session outside the scanner in order to minimize learning effects during scanning. These data were not analyzed and were used merely for practice purposes and to limit potentially confounding effects of learning during scanning.

Imaging Acquisition and Preprocessing

Images were acquired on a 3T Siemens Trio. Stimuli were presented to the subject via a projector at the rear of the scanner, reflected off a mirror mounted to the headcoil. Experimental tasks were presented using E-Prime software version 1.2 (Psychology Software Tools, Inc., Pittsburgh, PA).

Functional T_2^* -weighted images were acquired using an EPI sequence with 35 contiguous slices and $3.44 \times 3.44 \times 3.75 \text{ mm}^3$ voxels (repetition time (TR) = 2000 ms; echo time = 25 ms; flip angle = 70; field of view = 220). Phase and magnitude images were collected to estimate the magnetic inhomogeneity. T_1 -weighted MPRAGE images were collected for spatial normalization ($256 \times 256 \times 192$ matrix of $1 \times 1 \times 1 \text{ mm}^3$ voxels; TR = 1800 ms; echo time = 2.67 ms; flip angle = 9).

Functional data were spike-corrected to reduce the impact of artifacts using AFNI’s 3dDespike (<http://afni.nimh.nih.gov/afni>). Subsequent processing and analyses were done using SPM5 (<http://www.fil.ion.ucl.ac.uk/spm/>). Functional images were corrected for differences in slice timing using sinc-interpolation (Oppenheim et al. 1999) and head movement using least-squares approach and a 6 parameter rigid body spatial transformation. Structural data were coregistered to the functional data and segmented into gray and white-matter probability maps (Ashburner and Friston 1997). These segmented images were used to calculate spatial normalization parameters to the MNI template, which were subsequently applied to the functional data. As part of spatial normalization, the data were resampled to $2 \times 2 \times 2 \text{ mm}^3$. An 8-mm full-width at half-maximum isotropic Gaussian smoothing was applied to all functional images prior to analysis using SPM5. All analyses included a temporal high-pass filter (128 s), correction for temporal autocorrelation using an autoregressive AR(1) model, and each image was scaled to have a global mean intensity of 100.

Imaging Analysis

Univariate analyses were conducted using the general linear model implemented in SPM5. Separate regressors were included for the following events: higher level context update (“1” or “2”), higher level

context non-update (“3”), lower level context update (“A” or “B”), lower level context non-update (“C”), response (“X” or “Y”) and non-response (“Z”). For these regressors, each event was treated as an impulse (i.e. delta function) and convolved with the canonical hemodynamic response function. Additional regressors of non-interest were included for delay intervals in-between events to capture activation related to WM maintenance processes. We also included separate regressors for error trials, as well as intercept terms for each run. Error trials were excluded from the analyses described below. For subjects demonstrating >3 mm/degrees of motion across a session or >0.5 mm/degrees of motion between TRs, 24 motion regressors were included to capture linear, quadratic, differential, and squared differential residual motion variance (Lund et al. 2005). This procedure was only necessary for a single subject and all other subjects demonstrated <2 mm/degrees of motion throughout the scanning session.

Parameter estimates were calculated separately for each subject and used to conduct second-level (group) random effects analysis. This analysis consisted of a 2×2 ANOVA with factors of context level (higher or lower) and update (update or non-update). Within this ANOVA, simple effect contrasts were performed examining higher level context update – higher level context non-update (hereafter, higher level context update) and lower level context update – lower level context non-update (hereafter, lower level context update). A valid conjunction analysis (Nichols et al. 2005) was performed across these simple effect contrasts to identify regions common to context updating across levels of abstraction. Additionally, interaction contrasts were performed investigating regions significantly more involved in higher level context updating than lower level context updating and vice versa. These interaction contrasts were restricted to voxels showing a significant simple effect in order to avoid interactions being driven by deactivations of the control condition. All whole-brain univariate analyses were thresholded at $P < 0.001$ at the voxel-level with a 72 voxel ($2 \times 2 \times 2$ mm³) extent criterion providing a corrected P -value of $P < 0.05$ according to simulations using AlphaSim. Simulations were performed using smoothness estimations provided by AFNI’s 3dFWHMx and estimated smoothness (FWHM: 7.86 7.90 7.90) was similar to the smoothing applied to the data (FWHM: 8 8 8).

Whole-brain analyses were complemented by region-of-interest (ROI) analyses. These analyses sought to examine whether abstraction gradients were evident in critical PFC ROIs (see below). In order to create unbiased ROIs for analysis, data used to create and test ROIs were separated to ensure independence (Kriegeskorte et al. 2009). Specifically, we used a leave-one-subject-out procedure to define PFC ROIs. For each subject, ROIs were created using the other 20 subjects in the sample. Data for the held out subject were then tested within the independent group defined ROIs. This process was repeated for each subject. Parameter estimates for each condition of interest were averaged across each independent ROI and subjected to ANOVAs. ROIs were drawn from the left anterior PFC (aPFC) and left posterior PFC (pPFC). ROIs were also created in the bilateral BG and left PPC so that unbiased parameter estimates could be visualized.

Functional connectivity analyses were conducted to examine condition-related changes in functional connectivity between ROIs revealed in univariate analyses. These analyses were performed using the beta series correlation method (Rissman et al. 2004). Under this method, separate parameters (betas) are estimated for each event of

each trial. For a given event type, this creates a series of beta images (i.e. one per trial). Correlations between a given seed region and all other voxels of the brain across the beta series represent a measure of functional connectivity. Comparisons of these connectivity measures between conditions provide an assessment of task-related changes in functional connectivity. Accordingly, the model described above was re-estimated using separate predictors for each trial for each event. Eighty-one-voxel seeds (i.e. 5 mm spheres) were placed around the peak of left aPFC region active for higher level context updates and the left pPFC region active for lower level context updates. The size of the seeds was based upon ROI sizes used elsewhere (Nee et al. 2011) that followed from our own investigations which suggested that 5 mm spheres provide a good balance between functional specificity and noise reduction. Seeds were created so as to ensure independence between the seed and test data (Kriegeskorte et al. 2009). Specifically, seeds for each subject were based upon group activation peaks with that subject held out (i.e. leave-one-subject-out procedure). Although it is not clear that changes in functional connectivity should be related to univariate changes, the use of independent seeds ensures unbiased tests. The left aPFC seed was centered around the maximally active voxel with a y -coordinate >45 (mean peak -32.38 48.10 29.24). The left pPFC seed was centered around the maximally active voxel with $-30 < x < -20$ and $-10 < y < 0$ (mean peak -22.38 -2.19 63.81). Separate correlation maps were then computed for each subject between each seed region and each voxel in the brain for the following events: higher level context update, higher level context non-update, lower level context update, and lower level context non-update. These correlation maps were transformed with an arc-hyperbolic tangent function in order to approximate a normal distribution.

The normalized correlation maps were then submitted to second-level (group) 2×2 ANOVAs with factors of context level (higher or lower), and update (update or non-update). Separate ANOVAs were performed for each seed region (aPFC, pPFC). Planned contrasts examined updating-related changes in functional connectivity (i.e. update $>$ non-update) separately by level (higher, lower) and the interaction between levels (i.e. higher $>$ lower, lower $>$ higher). As we were primarily interested in frontal-striatal and frontal-parietal interactions, connectivity analyses were restricted to masks defined by univariate analyses described above. These masks included the left and right BG and the left PPC. All results were thresholded at $P < 0.05$ cluster corrected according to AlphaSim. Separate simulations were performed for analyses using the BG mask ($P < 0.05$, 78 voxel extent) and left PPC mask ($P < 0.05$, 129 voxel extent). Exploratory whole-brain analyses are reported in Supplementary Table 1 and were thresholded at the whole-brain levels used for univariate analysis described above.

Results

Behavioral Results

Behavioral data were analyzed to confirm expected signatures of WM performance based upon previous research (Fig. 2). Due to the symmetry of the data (effect of higher level context

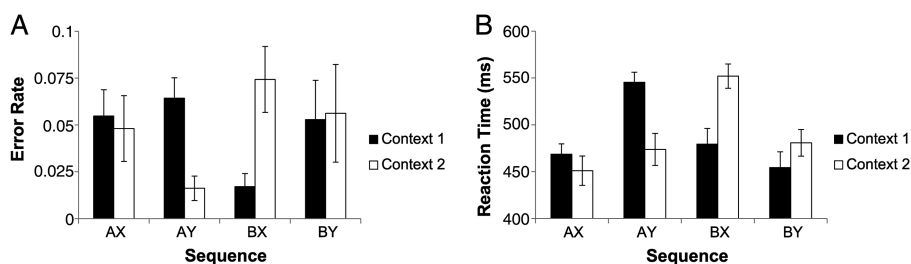


Figure 2. Behavioral data. Error rate (A) and reaction time (B) data were symmetrical for each higher level context (“1” or “2”) indicating appropriate use of higher level context cues to guide performance. Reduced performance on 1AY and 2BX sequences suggest that subjects used higher and lower level context cues to form an expectation of the target sequence and were slowed and less accurate when this expectation was violated.

$F_{1,20} < 1$ in both accuracy and reaction time), we collapsed across the higher level context for simplicity and to facilitate comparison with prior literature on the AX-CPT that has used only lower level contexts. Using the traditional parlance, we refer to “AX” as a target response (explicitly: the sequence “A-X” for context “1” and “B-Y” for context “2”). Similarly, “AY” refers to both “1-A-Y” and “2-B-X”, “BY” refers to both “1-B-Y” and “2-A-X”, and “BX” refers to both “1-B-X” and “2-A-Y”.

Accuracy was near ceiling (mean accuracy = 95.1%; d' context = 5.07) indicating that subjects adequately understood instructions and appropriately maintained contextual information in WM. Prior research with the AX-CPT has demonstrated that responses depend on context with better performance on non-target (“Y”) trials following “B” cues relative to “A” cues (Servan-Schreiber et al. 1996; MacDonald 2008; Barch et al. 2009). It is thought that following “A” cues, subjects expect the target sequence and performance is reduced when this expectation is violated. In prior research, such expectations are increased through frequent presentation of the target sequence (e.g. 70% “AX”, 10% “AY”). Although the present task used balanced frequencies (“AX” and “AY” were equally likely), these data were nevertheless consistent with prior research. Responses to “AY” trials were slowed relative to “BY” trials (mean difference: 73.97 ms; $t(20) = 9.89$, $P < 0.001$) and accuracy was non-significantly lower on “AY” trials relative to “BY” trials (mean difference: 1.88%; $t(20) = 1.85$, $P < 0.1$). Similarly, performance on “AY” trials was reduced relative to “BX” trials in both RT (mean difference: 50.15 ms; $t(20) = 5.08$, $P < 0.001$) and accuracy (mean difference: 5.26%; $t(20) = 4.97$, $P < 0.001$). Hence, these data suggest that subjects used contextual information to guide expectation.

Previous research has also demonstrated reduced performance on “BX” trials relative to “BY” trials. This decrement is thought to be due to loss of contextual information causing uncertainty with regard to how to respond to “X” stimuli. In the present task, loss of contextual information can operate at 2 levels: higher level context loss would lead to uncertainty regarding how to respond to “BY” trials whereas lower level context loss would lead to uncertainty regarding how to respond to “BX” trials. Here, responses to “BX” trials were slowed relative to “BY” trials (mean difference: 23.82 ms; $t(20) = 4.15$, $P < 0.001$), although this effect reversed in the accuracy data (mean difference: -2.11% , $t(20) = 2.50$, $P < 0.05$). These effects were not correlated ($r = -0.04$, $P > 0.85$) suggesting that they did not reflect a simple speed-accuracy tradeoff. These data indicate impacts of context loss on behavior. However, due to hierarchical dependencies, it is difficult to unambiguously relate either the reaction time or accuracy effects to a particular level of context since responses are always jointly determined by higher and lower level contexts.

Univariate fMRI Results

We began by separately examining neural activations relating to updates of higher and lower level contexts in WM. Based on prior literature regarding hierarchical organization in the PFC, we expected that higher level context updates would activate anterior regions of the PFC whereas lower level context updates would activate posterior regions of the PFC. Of particular interest was how the BG would respond to different levels of context updating. Univariate results are summarized in Tables 1 and 2.

Table 1
Neural correlates of context updating

	x	y	z	Extent	Z	BA	Region
Higher level context update							
Frontal/cingular	-4	10	34	2278	5.48	24	Anterior cingulate cortex
	2	2	40		5.35	24	Anterior cingulate cortex
	6	-26	40		4.97	23	Posterior cingulate cortex
	-32	48	28	266	4.70	46	Left anterior middle frontal gyrus
	32	48	30	153	3.75	46	Right anterior middle frontal gyrus
	38	48	20		3.50	46	Right anterior middle frontal gyrus
Sub-cortical	-18	2	-8	820	5.63		Left putamen/pallidum/ventral striatum
	8	-20	8		4.78		Right thalamus—medial dorsal nucleus
	-14	0	0		3.80		Left pallidum
	22	4	-8	477	5.05		Right putamen/pallidum/ventral striatum
	14	-2	-14		4.24		Right ventral striatum
	14	-10	-12		3.60		Right midbrain—subthalamic nucleus
Cerebellar	34	-50	-28	261	5.03		Right cerebellum—culmen
	26	-62	-24		3.76		Right cerebellum—culmen
	-32	-48	-28	180	4.67		Left cerebellum—culmen
	-24	-62	-24		4.31		Left cerebellum—culmen
Lower level context update							
Frontal/cingular	-26	-6	60	483	4.63	6	Left posterior superior frontal sulcus
	-10	-2	66		3.58	6	Left superior frontal gyrus
	-22	8	58			8	Left superior frontal sulcus
	-4	18	38	540	4.54	32/24	Anterior cingulate cortex
	-4	10	42		4.39	32/24	Anterior cingulate cortex
	0	26	30		4.37	32/24	Anterior cingulate cortex
	32	-6	56	99	3.68	6	Right posterior superior frontal sulcus/middle frontal gyrus
Parietal	-48	-34	48	1339	5.67	40	Left intraparietal sulcus/inferior parietal lobule
	-46	-42	50		5.34	40	Left intraparietal sulcus/inferior parietal lobule
	-38	-56	48		5.14	40/7/39	Left intraparietal sulcus
	8	-60	48	561	4.37	7	Precuneus
	6	-68	40		3.66	7	Precuneus
Cerebellar	36	-56	-34	447	5.57		Right cerebellum—culmen
	32	-50	-30		5.56		Right cerebellum—culmen
	26	-66	-28		4.27		Right cerebellum—culmen
	-34	-56	-32	328	5.04		Left cerebellum—Culmen

Table 2 Conjunctions and interactions							
	x	y	z	Extent	Z	BA	Region
Higher level context update and lower level context update							
Frontal/cingular	-4	6	42	198	4.02	32/24	Anterior cingulate cortex
	-4	-2	48		4.01	32/6	Anterior cingulate cortex/pre-supplementary motor area
	-2	20	32		3.74	24	Anterior cingulate cortex
Cerebellar	34	-50	-28	178	5.03		Right cerebellum—culmen
	28	-64	-26		3.54		Right cerebellum—culmen
	-34	-50	-30	111	4.27		Left cerebellum—culmen
	-28	-62	-26		3.25		Left cerebellum—culmen
Higher level context update > lower level context update							
Frontal/cingular	-4	6	32	110	3.84	24	Anterior cingulate cortex
	6	10	30		3.79	24	Anterior cingulate cortex
	-6	-2	36		3.46	24	Mid-cingulate cortex
	6	-26	40	416	4.52	23	Posterior cingulate cortex
	-6	-32	34		4.39	23	Posterior cingulate cortex
	-6	-36	42		4.25	31	Posterior cingulate cortex
Sub-Cortical	-16	2	-10	317	5.76		Left putamen/pallidum/ventral striatum
	-8	-4	-2		3.60		Left thalamus
	22	4	-8		5.66		Right putamen/pallidum/ventral striatum
	24	-4	-8	425	4.33		Right pallidum
	12	-12	-10		3.98		Right midbrain
Lower level context update > higher level context update							
Parietal	-46	-46	54	260	4.09	40	Left inferior parietal lobule
	-42	-40	50		3.89	40	Left inferior parietal lobule/intraparietal sulcus
	-42	-52	48		3.45	40	Left inferior parietal lobule




Figure 3. Univariate contrasts. Top: neural activations for the contrast of higher level context update – higher level context non-update. Activations included the aPFC, anterior cingulate, and bilateral BG. Bottom: neural activations for the contrast of lower level context update – lower level context non-update. Activations included ppFC in the caudal superior frontal sulcus, anterior cingulate, and PPC. BG activation was absent in this contrast. Results were thresholded at $P < 0.001$ at the voxel-level, corrected using cluster extent ($P < 0.05$ corrected).

Higher Level Context Update

We began by identifying regions sensitive to updating higher level contexts in WM (i.e. higher level context update – higher level context non-update; Fig. 3). Regions demonstrating activation increases when the higher level context was updated included bilateral aPFC (BA 46) and bilateral BG (striatum and pallidum). The anterior cingulate cortex (ACC), medial dorsal nucleus of the thalamus, and cerebellum were also sensitive to higher level context updates.

Lower Level Context Update

Regions demonstrating activation increases when the lower level context was updated included bilateral ppFC in the caudal superior frontal sulcus (BA 6; Fig. 3). The ACC, left posterior parietal cortex (PPC; intraparietal sulcus and inferior parietal lobule), and preCuneus also increased in activation when the lower level context was updated. Notably, the BG did not appear to be sensitive to lower level context updates, even at a lenient threshold ($P < 0.01$ uncorrected, 20 voxel extent).

Level × Update Interactions

In order to determine whether any region was selective to updating particular levels of representations in WM, we examined level × update interactions. These analyses revealed that the bilateral BG were selectively responsive to higher level context updates (higher level update > lower level update; Fig. 4B). By contrast, the left PPC was selectively responsive to lower level context updates (lower level update > higher level update; Fig. 4B). Level × update interactions were also found in the posterior cingulate and ACC, both of which demonstrated greater activation for higher than lower level context updates.

We also examined whether predicted abstraction gradients were present in the PFC using independent ROIs (see Materials and Methods). Averaged contrast estimates were drawn from the aPFC and pPFC separately for higher and lower level context updates. These averaged contrast estimates were then submitted to a 2 × 2 ANOVA with factors of region (aPFC, pPFC) and level (higher, lower). As anticipated, there was a significant region × level interaction (Fig. 4A; $F_{1,20} = 9.31$, $P < 0.01$). Although lower level context updates did activate the aPFC (mean contrast estimate (SEM): 0.52 (0.21); $t(20) = 2.43$, $P < 0.05$), higher level context updates activated the aPFC to a significantly greater degree (mean contrast estimate (SEM): 1.02 (0.26); $t(20) = 3.87$, $P < 0.001$; difference: $t(20) = 2.10$, $P < 0.05$). The reverse was true in the pPFC: while both forms of context updating activated the pPFC, this activation was numerically greater for lower (mean contrast estimate (SEM): 0.63 (0.11); $t(20) = 5.52$, $P < 0.0001$) than higher level context updates (mean contrast estimate (SEM): 0.38 (0.13); $t(20) = 2.85$, $P < 0.01$), although this difference did not achieve significance ($t(20) = 1.40$, $P > 0.15$). Hence, while neither PFC region was *selective* to one level of context updating, there was evidence of an abstraction gradient with anterior regions significantly more responsive to higher level context updating and posterior regions numerically, but not significantly, more responsive to lower level context updating.

Conjunction

Finally, we examined whether any regions were commonly active across both levels of updating in WM. Regions common to both higher and lower level context updates were located in the ACC (BA 32). This region was located along the cingulate sulcus and was dorsal to the ACC region that demonstrated a level × update interaction, which was located adjacent to the corpus callosum. The bilateral cerebellum was also common

to both forms of context updating. Critically, the BG were not conjointly active for both levels of context updating.

Univariate Summary

In summary, activations in the PFC demonstrated the anticipated sensitivity to level of abstraction with anterior regions of the PFC showing greater sensitivity to more abstract, higher level context updates and posterior regions of the PFC showing an opposite trend—greater sensitivity to less abstract, lower level context updates. Although the BG demonstrated activation during context updating, this activation was selective for higher level context updates. By contrast, lower level context updates demonstrated selective activation in the left PPC.

Before proceeding further, it is important to examine whether lower level context updating activation can be solely attributable to response preparation processes. In particular, when subjects received the lower level context cue “B” in the “1” context or “A” in the “2” context, a non-target response could have been proactively prepared (response certain). Analysis of motor cortex revealed the presence of such advanced response preparation (see Supplementary material, Results: Advanced Response Preparation). Interestingly, motor activation revealed advanced preparation of a target response following “1 + A” and “2 + B” combinations (response uncertain) even though these contexts provided no information regarding the forthcoming response. These data are consistent with the reduced performance on “AY” trials, as a prepared target response would need to be overcome on these trials. Given the patterns in motor cortex, it is possible that activations in pPFC and PPC regions may also have been driven by advanced response preparation. To explore this possibility, we calculated indices of advanced response preparation (response certainty × response hand interaction; see Supplementary material, Results, for full details). These analyses took into account response hand to target/non-target mappings that were counter-balanced between subjects (response hand).

Activations in the left pPFC were neither sensitive to response certainty ($F_{1,19} = 0.08$, $P > 0.75$) nor response preparation ($F_{1,19} = 1.27$, $P > 0.25$). Follow-up analyses re-evaluated prior contrasts while factoring out potential motor contributions. Since motor contributions affect the contralateral hemisphere, we considered activations ipsilateral to the prepared motor response. In other words, when evaluating left pPFC activations, we considered response certain contexts (e.g. non-target preparation) for subjects who made




Figure 4. Level × update interactions in ROIs. All depicted results were drawn from unbiased, test-independent ROIs (see Methods). (A) Contrast estimates for higher (gray) and lower (white) level updates in the left anterior PFC (L_aPFC) and left posterior PFC (L_pPFC). A significant level × update interaction indicated an anticipated PFC abstraction gradient. aPFC was activated more strongly by higher level updates, whereas pPFC was numerically more strongly activated by lower level updates. (B) The bilateral BG showed updating-related increases only for higher level context updates, as well as demonstrating significant update × level interactions. The left posterior parietal cortex (L_PPC) demonstrated the opposite pattern.

non-target responses with the left hand and we considered response uncertain contexts (e.g. target preparation) for subjects who made target responses with their left hand. Hence, these data have the least contribution of putative response preparation. We refer to these data as response non-prep contexts. When considering only response non-prep contexts, the left pPFC demonstrated a strong effect of lower level context updating ($t(20)=5.10, P<0.0001$). Furthermore, we repeated the region (aPFC, pPFC) \times level (higher, lower) ANOVA reported above considering only response non-prep contexts. The region \times level interaction was still significant ($F_{1,20}=4.79, P<0.05$). These results demonstrate that the left pPFC activations could not be explained by contributions of response preparation. Finally, the left PPC was also insensitive to response preparation effects ($F_{1,19}=0.12, P>0.7$). However, the left PPC did show an effect of response certainty ($F_{1,19}=9.16, P<0.01$) with greater activation when the response was certain relative to when it was not. Updating-related activation in the left PPC was still significant when considering only response uncertain contexts ($t(20)=3.99, P<0.001$) and the level \times update interaction remained significant, as well ($t(20)=1.92, P<0.05$, one-tailed). As a result, our results cannot be fully explained by advanced response preparation (see also Supplementary material, Results: Contrasting Attention, Response Preparation, and Responding in posterior PFC). These data confirm that activations in the left pPFC and PPC are driven by WM updating independently of response preparation.

Functional Connectivity fMRI Results

Models of WM suggest that updating the contents of WM depend not only on activation of the PFC and BG, but also upon the *interaction* of the PFC and BG. In order to explore these interactions, we examined changes in functional

connectivity during WM updating. Specifically, we examined whether the aPFC and BG demonstrated enhanced functional coupling during higher level context updates relative to lower level context updates. Given that the left PPC was selective for lower level context updates, we were also interested in whether the pPFC and PPC demonstrated enhanced functional coupling for lower level context updates relative to higher level context updates. Notably, in all cases, measures were merely correlational and cannot indicate causality, direction, or rule out relationships via intermediary nodes.

Separate condition-specific connectivity maps with the left aPFC as a seed were created for higher level context updates, lower level context updates, higher level context non-updates, and lower level context non-updates (see Materials and Methods). A 2×2 ANOVA with factors of level (higher, lower) and update (update, non-update) assessed changes in functional connectivity as a function of level and updating. This analysis revealed a significant level \times update interaction in the bilateral BG (Fig. 5A, Supplementary Table 1). This interaction was driven by greater updating-related increases in functional connectivity for higher level context updates relative to lower level context updates. When considering the simple effect of higher level context updating alone, sub-threshold clusters were found in the bilateral BG that did not meet our cluster extent criterion (64 voxel cluster in the left hemisphere, 54 voxel cluster in the right hemisphere). No voxels in the BG demonstrated increased connectivity with the left aPFC during lower level context updating. Exploratory whole-brain connectivity analyses revealed additional clusters in the anterior medial superior PFC and right anterior ventrolateral PFC which demonstrated greater updating-related connectivity increases with the left aPFC for higher level context updates than lower level context updates (Supplementary Table 1).

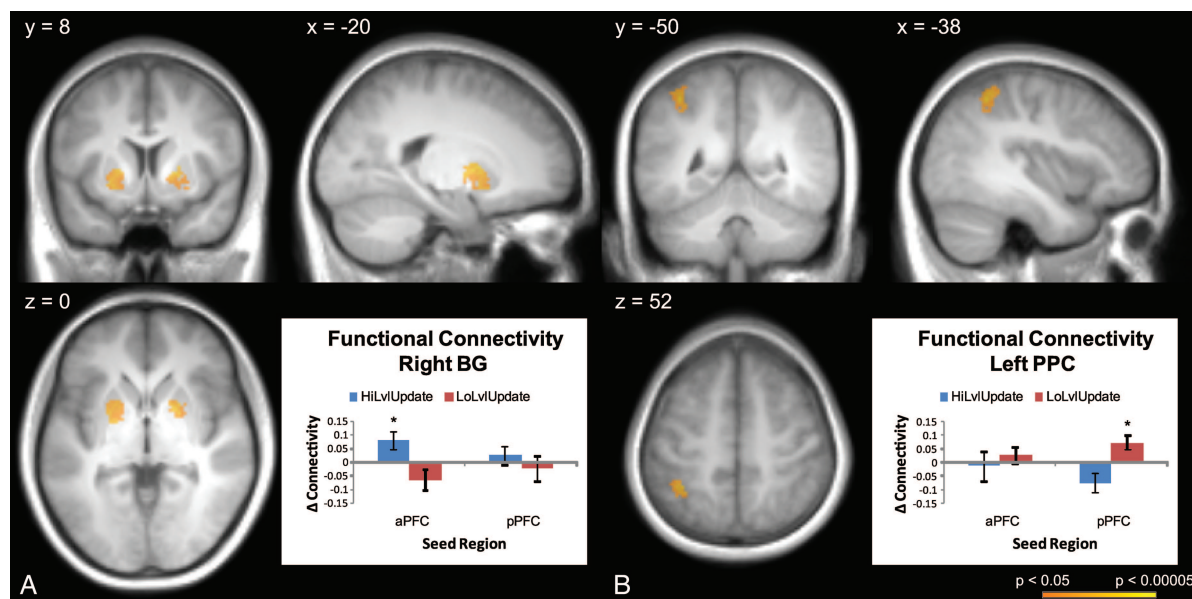


Figure 5. Functional connectivity analyses. (A) The aPFC demonstrated increased correlations with the BG during higher level context updates (update–non-update; HiLvlUpdate) relative to lower level context updates (LoLvlUpdate). Increases in connectivity were not found between the aPFC and BG during lower level context updates and there were no changes in connectivity between the BG and pPFC for either contrast. (B) By contrast, the pPFC demonstrated connectivity increases with the PPC during lower level context updates. These data indicate that dissociable aPFC–BG and pPFC–PPC networks are involved in higher and lower level context updating, respectively.

The above analysis was repeated with the left pPFC as a seed. This analysis revealed a significant level \times update interaction in the left PPC, which showed greater updating-related functional connectivity increases for lower level context updates relative to higher level context updates (Fig. 5B; Supplementary Table 1). The simple effect of lower level context updating revealed a similar left PPC cluster, whereas no significant simple effect was found for higher level context updating. Exploratory whole-brain connectivity analyses revealed no additional significant changes in connectivity with the left pPFC.

To summarize, both univariate and functional connectivity analyses confirmed a role of the aPFC and BG during updates of higher level context information in WM. By contrast, pPFC and PPC regions interacted during updates of lower level context information in WM. These results contradict models that assume that PFC–BG interactions underlie parallel updating mechanisms of hierarchical information. Instead, these results suggest that independent mechanisms underlie the updating of different forms of information in WM. We discuss what these potential mechanisms might be below. Before doing so, we first examine some potential alternative accounts of these data.

Alternative Accounts

One distinction between higher and lower level contexts in the present task is that higher level contexts must be robustly maintained across several intervening events, whereas lower level contexts must rarely span intervening events. It may be the case that PFC–BG interactions are only critical for forming distractor-resistant representations (McNab and Klingberg 2008) and the absence of BG activation for lower level contexts may be due to the reduced need to represent lower level contexts in a robust form. If so, lower level contexts should be easily disrupted by distracting stimuli. To test this hypothesis, we compared performance on trials that contained a distracting stimulus in-between the lower level update and response (i.e. those that contained a “Z”) with those that did not. Distraction had no effect on accuracy (95.04% without distraction; 95.37% with distraction; $t(20) = -0.26$, $P > 0.75$). RT was actually improved on trials with distraction (466.04 ms) than without (501.19 ms; $t(20) = -3.88$, $P < 0.001$). This improvement likely reflects the release of response inhibition that participants used to prepare for the possibility of a “Z” (“no-go”) stimulus. Hence, lower level contexts did not appear to be easily disrupted suggesting that distractor-resistance does not dissociate the networks found here. However, that more salient distraction might differentially affect higher and lower level contexts cannot be ruled out by these analyses and exploration of such conditions would be a target for future research.

It should also be noted that the models of WM updating explored here (Frank et al. 2001; O’Reilly and Frank 2006) suggest that the BG contain both cells involved in updating (“go” cells), as well as those that prevent items from entering WM (“no-go” cells). Hence, the BOLD signal may contain a mixture of the activity of both such cells. The way each type of cell influences the BOLD signal should have equal bearings upon all contrasts. So, the presence of a level \times update interaction in the least suggests dissociable BG mechanisms acting upon higher and lower level context updates. To further

elucidate whether a positive effect in such a contrast can be related to the action of “go” cells, we examined BG activation during the motor response (“X” or “Y”) compared with a comparable motor no-go (“Z”). Both the left ($t(20) = 4.49$, $P < 0.0005$) and right BG ($t(20) = 3.73$, $P < 0.005$) regions sensitive to higher level context updates were also more active for a motor response relative to withholding a response. As a result, it appears that our contrasts reflect the action of “go” cells, which do not appear to be sensitive to lower level context updates.

Discussion

Our data revealed that distinct neural correlates underlie the updating of hierarchical contexts in WM. Updating WM with higher level contextual information involved the aPFC and BG. By contrast, updating WM with lower level contextual information involved the pPFC and PPC. These regions were revealed by univariate investigations of activation patterns and their interactions were confirmed by examining changes in functional connectivity.

Gating and the BG

By the gating hypothesis, phasic activity in the BG acts as a trigger to allow new input into the PFC (Braver and Cohen 2000; Frank et al. 2001; Rougier et al. 2005; O’Reilly and Frank 2006; Reynolds and O’Reilly 2009). Consistent with these ideas, the BG and PFC are active during tasks that require updating WM (Callicott et al. 1999; Lewis et al. 2004). Moreover, patients with Parkinson’s Disease, which is marked by dopamine depletion in the dorsal striatum, are impaired in tasks that require updating of task-sets in WM (Cools et al. 2001, 2003; Cools 2006). These impairments are ameliorated by L-3,4-dihydroxyphenylalanine (L-DOPA) which increases phasic bursting in the dorsal striatum (Cools et al. 2003; Cools 2006). Here, the aPFC and BG demonstrated increased activation when subjects updated the higher level context. Moreover, functional connectivity analyses demonstrated that these regions showed increased functional coupling during higher level context updates consistent with the idea that these regions interact to update WM.

An issue with the gating hypothesis is elucidating how WM can be updated selectively. If phasic BG activity acts as a signal to allow new input into the PFC, then this signal might globally overwrite the contents of WM even if some parts of WM should persist across time. To solve this issue, Frank et al. (2001), O’Reilly and Frank (2006), Reynolds and O’Reilly (2009) suggested that information stored in WM is organized in distinct “stripes”, each of which can be selectively updated by corresponding BG circuits. Although data demonstrating abstraction gradients in the PFC may correspond to such “stripes” in the PFC (Koechlin et al. 2003; Koechlin and Jubault 2006; Badre and D’Esposito 2007; Badre 2008; Botvinick 2008; Badre and D’Esposito 2009; Kounieher et al. 2009; Badre et al. 2010), research has not clarified whether the BG are organized in a similar fashion. Our data were consistent with regional abstraction in the PFC revealing aPFC regions involved in updating higher level contexts and pPFC regions involved in updating lower level contexts. However, the BG were sensitive to higher but not to lower

level updates, suggesting distinct updating mechanisms across levels of abstraction.

Frank et al. (2001) and O'Reilly and Frank (2006) have also posited that BG gating mechanisms that act to update WM perform similar functions in the motor domain to allow a response to be executed. Indeed, the postulation of the BG as an output gate originated in the motor domain (Chevalier and Deniau 1990). Our data demonstrated activation of the BG when subjects made a motor response compared with when they withheld/inhibited a response (i.e. "no-go"). That the BG were active for both higher level context updates and responding is consistent with the idea that the BG perform a similar gating function in both WM and motor output (see also "Localization within the BG" below). However, our data indicate that the BG are not involved in all forms of updating in WM.

One assumption of the Frank and O'Reilly model is that selectively updating WM leaves non-updated (maintained) portions of WM undisturbed. Presently, the validity of this assumption is unclear. Consider the n -back task, which is regarded as a trademark task of WM updating (Cohen et al. 1997): subjects are required to keep the last n presented items in WM. When a new item is presented, this entails removing an item from WM (i.e. the $n + 1$ item) and adding the newly presented item into WM. Is this new set of n items a distinct representation (i.e. a new episode) or is it a modification of a prior representation? Although we do not know of data that can address this question, it seems plausible that each adjustment of WM leads to a distinctly new representation/episode. Indeed, such distinctive encoding would be optimal for protecting WM from proactive interference (Jonides and Nee 2006). So, selective updating in WM may not operate as presumed.

Attention Shifting as a Mechanism of Selective Updating

Supposing the above account, how can lower level context updates be characterized? One possibility is that when a higher level context is updated (e.g. "1"), WM represents the set of relevant action rules (e.g. "A-X → Target", "A-Y → Non-Target", "B-X → Non-Target", "B-Y → Non-Target"). Then, the presentation of a lower level context (e.g. "A") shifts attention to the subset of now relevant rules (e.g. "X → Target", "Y → Non-Target"). Thus, lower level context updates would consist of attention-shifts within WM rather than changes to WM itself (Fig. 6). This account is consistent with recent formulations of the architecture of WM and the processes that act upon it (Oberauer 2002; Jonides et al. 2008; Nee and Jonides 2008, 2011). By these models, WM consists of approximately 4 ± 1 items maintained in an accessible state, with a single item/chunk residing in the focus of attention for immediate processing. The focus can be shifted between representations in WM, highlighting a particular item/chunk for further processing (e.g. to guide action) without changing the content of WM. Some recent data suggest that such attention-shifts in WM recruit similar neural substrates as those involved in shifting attention in the sensory environment (Nobre et al. 2004; Nee and Jonides 2009; Tamber-Rosenau et al. 2011). Notably, prominent amongst these regions are the caudal superior frontal sulcus and PPC (see also Bledowski et al. 2009; Bledowski et al. 2010), the regions found activated to lower level context updates in the present data. Hence,

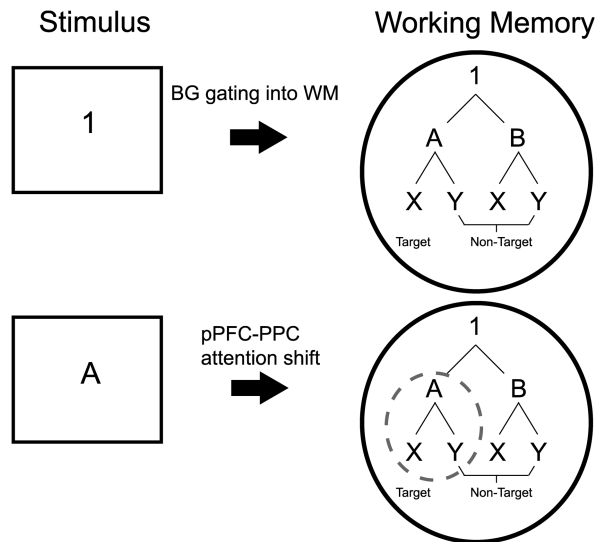


Figure 6. Schematic of Gating versus Shifting Processes. Visual stimuli are depicted on the left. Hypothesized contents of WM are presented on the right. Putative psychophysiological processes are represented by arrows. Top: When the higher level context cue is presented, it is gated into WM via the BG. Bottom: When the lower level context cue is presented, attention within WM shifts to the relevant set of rules (dashed gray circle) via pPFC and PPC processes.

responses to lower level context updates may be better modeled as attention-shifts than changes in WM content.

Taken together, our results suggest that how the brain organizes and updates hierarchical contexts in WM needs re-conceptualization. We suggest that changing the highest order context in WM requires updating WM with the set of currently relevant action rules. These updates involve PFC–BG interactions. Presentation of lower level contexts serves to limit the set of relevant rules and engages selective attention towards currently relevant rules. These attentional processes are subserved by pPFC and PPC interactions. Thus, hierarchical context updating in WM is accomplished by distinct mechanisms and neural substrates.

Sensitivity to Task-Demands

Whether updating or attention-shifting in WM occurs is likely to vary according to task-demands. We suggest that the BG may be involved in updating WM when the contents are globally overwritten. In such cases, loading even simple stimulus-response associations into WM should implicate the BG, consistent with other work (Cools 2006). Hierarchically structured content is complicated by the need to simultaneously maintain multiple levels of information. We hypothesize that in these cases, the BG will be involved for updating lower level content *only if* the higher level content can be simultaneously discarded. In the present design, this was not the case since the higher level context had to be maintained across trials. However, a design that included a higher level context cue on each trial, and thereby did not require persistence of the higher level context cue, may have afforded such processing.

A further consideration is the number of stimulus-response rules. Here, the number of rules was within putative limits of WM capacity. However, if the number of rules exceeds these limits, it is unclear how such information could be loaded

into WM to guide performance. It is possible that with enough practice, certain rule-sets could be combined into chunks. Then, attention may shift between rule-chunks, as well as within rule-chunks. If chunking cannot or has not been performed, then it may not be possible to simultaneously hold all relevant rules in mind. In such cases, subjects may adopt a strategy of holding cues, rather than rules, in WM. This would be akin to a retroactive strategy where cues are integrated with corresponding rules at the time of responding (Braver et al. 2007). In such cases, cues would be presumably be held in WM in a concrete form (e.g. phonological, visual) and activate corresponding posterior PFC regions regardless of hierarchical level. Investigation of how WM structures navigate such content would be an interesting future avenue.

Relationship to Previous Research in Hierarchical Abstraction

Characterizing the rostral–caudal axis of the PFC is a subject of much interest (Badre 2008; Badre and D’Esposito 2009; O’Reilly 2010). In their seminal work, Koehlin et al. (2003) suggested that abstraction is defined temporally. By this idea, more anterior regions of the PFC are involved as processing is required over longer intervals. In this framework, posterior regions of the PFC are involved when processing can be based on immediately present stimulus information, whereas more anterior regions of the PFC are involved when processing requires information from a context defined over a longer interval. Badre and D’Esposito (2007) suggested an alternative account. In their framework, abstraction is based on the specificity of representation. For example, “opening a door” forms an abstract class of actions, which are accomplished by a variety of more concrete operations (e.g. “turn”, “push”, “pull”), which can in-turn describe other more concrete operations (e.g. “turn knob”, “turn handle”). Although our purposes were not to distinguish between these or other accounts, it is worth noting that the higher level context in the 1–2-A-X task was more abstract than the lower level context both in terms of both time and representation. Temporally, the higher level context spanned a longer episode than the lower level context. In terms of representation, the higher level context denoted a more abstracted set of action rules than the lower level context. Hence, the aPFC regions found here may have been active in order to establish a longer lasting temporal episode or to select/represent an abstract rule set. Interestingly, a recent study with a similar paradigm to that used here added manipulations to dissociate factors of temporal and representational abstraction (Reynolds et al. 2012). That study did not find evidence of rostral–caudal distinctions by either factor. Instead, regions along the lateral PFC were adaptively active in a sustained or transient manner depending upon task-demands. Additional research will be needed to clarify how changes in task-demands elicit either rostral–caudal or sustained–transient distinctions.

It is also important to consider how the activations found here compare with previous studies investigating the rostral–caudal axis of the PFC. The aPFC region found in the present study was somewhat dorsal to the previous findings that were situated in the lateral frontal pole (BA 10; Badre and D’Esposito 2007), although it was similar in lateral and anterior extent (present peak: $-32\ 48\ 28$, Badre and D’Esposito peak: $-36\ 50$

6). O’Reilly (2010) has suggested that previous studies combined a mixture of rule complexity, thought to activate dorsal PFC regions, and category abstraction, thought to activate ventral PFC regions. As a result, activations crossed the dorsal–ventral axis along their rostral–caudal progression. Levels in our task varied only in rule complexity, which may explain why activations here were restricted to dorsal PFC regions. Interestingly, whereas the dorsal–ventral axis has been characterized in more posterior PFC regions, there is little evidence regarding whether it also extends to aPFC regions (Nee et al. 2012). That our aPFC activations differed from previous reports suggests that it may.

Another important consideration is the activation profile of different PFC regions. Badre and D’Esposito (2007) noted that activations in the lateral frontal pole demonstrated sustained, but not transient, activation profiles. Braver et al. (2003) noted a similar sustained, but not transient, effect in the lateral frontal pole during task-switching. Koehlin and Hyafil (2007) have suggested that the lateral frontal pole acts as a buffer to maintain pending task information, a process that they refer to as cognitive branching. Given our focus of transient updating processes, our design was not well-suited to examine sustained branching processes. This may explain the lack of lateral frontal pole activation in the present study.

Although anterior activations in the present study differed somewhat from previous reports, activations in the pPFC were in close proximity to previous studies (Koehlin et al. 2003; Koehlin and Jubault 2006; Badre and D’Esposito 2007). These previous studies have associated activation in this region with motor control/representation, which differs somewhat from our interpretation of attention control/representation. Previous research has demonstrated that regions related to motor preparation and attention are in close spatial proximity. This research has indicated that regions involved in attention are located just anterior to those involved in motor preparation (Boussaoud 2001; Picard and Strick 2001). The pPFC activation cluster found here extended caudally into putative motor preparation regions (dorsal preMotor cortex (PMd): $y = -14$), as well as rostrally into putative attention regions (pre-PMd; $y = 10$). As a result, these activations could potentially reflect a mixture of motor preparatory processes caudally and attention processes rostrally. To examine this possibility, we extracted the pPFC activation cluster sensitive to lower level context updates and re-analyzed it on a coronal slice-by-slice basis (see Supplementary material, Results: Contrasting Attention, Response Preparation, and Responding in the posterior PFC; Supplementary Figure 2). Separate indices of putative attentional control (response non-prep > lower level context non-update) and response preparation (response prep > response non-prep) were calculated and compared along with effects of responding and higher level context updating. Consistent with prior data, response preparation effects were significant in the caudal-most portions of the activation cluster ($y = -14$: $t(20) = 2.39$, $P < 0.05$; $y = -12$: $t(20) = 2.16$, $P < 0.05$). However, as activations proceeded more rostrally, these effects disappeared. By contrast, effects of higher level context updating were only significant in mid-caudal portions of the cluster ($-10 \leq y \leq 0$) just anterior to areas demonstrating effects of response preparation. Critically, effects of attentional control were strong throughout the entire cluster (all $t(20) > 3.79$, all $P < 0.005$). Hence, while there was evidence of a rostral–cognitive/caudal–motor gradient

consistent with previous data, effects of attention were present through the pPFC over-and-above all other effects. Hence, the present posterior PFC results are most easily accommodated by attentional explanations.

These results are consistent with an extensive literature linking the caudal superior frontal sulcus with selective attention (Kastner and Ungerleider 2000; Reynolds and Chelazzi 2004; Moore 2006). This process is thought to be accomplished through the representation of attentional priority (Serences et al. 2005). Such attentional prioritization is necessary for a variety of executive processes and a recent meta-analysis demonstrated that the caudal superior frontal sulcus is the most reliable site of activation across studies of executive function in WM (Nee et al. 2012). Although these data are consistent with the idea that pPFC activations found here are involved in shifting attention, it is less clear how they can be reconciled with similar activations found in previous studies of hierarchical control (Koechlin et al. 2003; Koechlin and Jubault 2006; Badre and D'Esposito 2007). Of note is that while pPFC activations in prior studies were sensitive to low level sensory or motor manipulations, they were also sensitive to manipulations of higher levels of abstraction. Such data have been interpreted as evidence of higher level PFC regions feeding backwards into posterior PFC regions and thereby increasing their activation level (Koechlin et al. 2003; Koechlin and Jubault 2006; Badre and D'Esposito 2007). While this is possible, it is also notable that manipulations of higher levels of abstraction in previous designs were commensurate with greater attentional demands and reaction times. As a result, it is possible that activations in the pPFC in previous studies may have been the result of increasing attentional demands.

Localization Within the BG

A number of theories posit distinct frontal–BG loops involved in cognitive and sensorimotor functions (Alexander et al. 1986; Parent and Hazrati 1995; Haber 2003). Although the BG areas found in the present study were identified through cognitive demands (i.e. WM updating), we found that these areas were also sensitive to sensorimotor operations (i.e. responding). This common sensitivity to cognitive and sensorimotor demands may seem at odds with distinct loop proposals. However, a full treatment of these proposals would require more detailed analyses that contrast cognitive and motor demands. Such analyses are outside the scope of the present study. Following anatomical data, it is likely that distinct portions of the BG are sensitive to cognitive and motor operations, but spatial normalization and smoothing preprocessing operations common in fMRI analysis may blur such distinctions. Future studies may definitively examine whether there is a gradient within the BG for different gating demands similar to the gradients present in the PFC. Ideally, such studies would incorporate high-resolution acquisition of fMRI data to uncover potentially subtle gradients (Lehericy et al. 2006).

Distinct areas of the frontal cortex are known to interact with distinct regions of the BG (Alexander et al. 1986; Parent and Hazrati 1995; Haber 2003). Much of this knowledge is based on invasive recordings in non-human animals, although recent non-invasive studies in humans have provided additional insights (Lehericy et al. 2004; Postuma and Dagher 2006; Di Martino et al. 2008). BG activations in the present

study were located in the ventral putamen approximately at the level of the anterior commissure and extended into the pallidum. Interestingly, these activations did not extend into the caudate, which is thought to interact with the dorsolateral–PFC to perform cognitive operations including WM (Alexander et al. 1986; Parent and Hazrati 1995; Haber 2003). However, areas of the putamen are also known to support cognitive operations. A recent study examining resting state connectivity demonstrated that ventral–rostral areas of the putamen similar in location to the present activations showed strong functional connectivity with the aPFC (Di Martino et al. 2008). The aPFC region found in that study was similar to the aPFC activations found here (−34 42 22 vs. −32 48 28). The co-activation of these regions is corroborated by a meta-analysis that revealed that similar aPFC regions in area 46 are frequently active in studies reporting activation in the putamen (Postuma and Dagher 2006). The convergence of these data suggest that the ventral–rostral putamen and area 46 consistently interact. However, our data suggest that whether anatomical connections between the PFC and BG are functionally engaged depends upon task-demands.

Conclusion

Hierarchically organized working memories afford a wide-range of flexible behaviors based upon contextual information. Influential models have been constructed to explain these important functions via parallel PFC–BG loops that operate across levels of abstraction (Frank et al. 2001; O'Reilly and Frank 2006). Our data indicate that hierarchically structured behavior is not as parallel as it might seem. Instead, independent neural and psychological mechanisms operate upon different levels of abstraction with PFC–BG interactions supporting updating the highest level of hierarchical contexts. By contrast, attention-shifts involving pPFC–PPC interactions serve to highlight different subsets of rules in order to accommodate lower level context information.

Supplementary Material

Supplementary material can be found at: <http://www.cercor.oxfordjournals.org/>.

Funding

This research was supported in part by AFOSR FA9550-07-1-0454 (J.B.), R03 DA023462 (J.B.), R01 DA026457 (J.B.), and the Indiana METACyt Initiative of Indiana University, funded in part through a major grant from the Lilly Endowment, Inc. Supported in part by the Intelligence Advanced Research Projects Activity (IARPA) via Department of the Interior (DOI) contract number D10PC20023.

Notes

The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of IARPA, DOI or the U.S. Government. The authors thank C. Chung, B. Pruce, and R. Fukunaga for help with data collection.

References

- Alexander GE, DeLong MR, Strick PL. 1986. Parallel organization of functionally segregated circuits linking basal ganglia and cortex. *Annu Rev Neurosci.* 9:357–381.
- Ashburner J, Friston K. 1997. Multimodal image coregistration and partitioning—a unified framework. *Neuroimage.* 6:209–217.
- Badre D. 2008. Cognitive control, hierarchy, and the rostro-caudal organization of the frontal lobes. *Trends Cogn Sci.* 12:193–200.
- Badre D, D'Esposito M. 2007. Functional magnetic resonance imaging evidence for a hierarchical organization of the prefrontal cortex. *J Cogn Neurosci.* 19:2082–2099.
- Badre D, D'Esposito M. 2009. Is the rostro-caudal axis of the frontal lobe hierarchical? *Nat Rev Neurosci.* 10:659–669.
- Badre D, Kayser AS, D'Esposito M. 2010. Frontal cortex and the discovery of abstract action rules. *Neuron.* 66:315–326.
- Barch DM, Berman MG, Engle R, Jones JH, Jonides J, Macdonald A, 3rd, Nee DE, Redick TS, Sponheim SR. 2009. CNTRICS final task selection: working memory. *Schizophr Bull.* 35:136–152.
- Bledowski C, Kaiser J, Rahm B. 2010. Basic operations in working memory: contributions from functional imaging studies. *Behav Brain Res.* 214:172–179.
- Bledowski C, Rahm B, Rowe JB. 2009. What “works” in working memory? Separate systems for selection and updating of critical information. *J Neurosci.* 29:13735–13741.
- Botvinick MM. 2008. Hierarchical models of behavior and prefrontal function. *Trends Cogn Sci.* 12:201–208.
- Boussaoud D. 2001. Attention versus intention in the primate premotor cortex. *Neuroimage.* 14:S40–45.
- Braver TS, Cohen JD. 2000. On the control of control: The role of dopamine in regulating prefrontal function and working memory. In: Monsell S, Driver J, editors. *Attention and performance XVIII: control of cognitive processes.* Cambridge, MA: MIT Press. p. 713–737.
- Braver TS, Gray JR, Burgess GC. 2007. Explaining the many varieties of working memory variation: dual mechanisms of cognitive control. In: Conway A, Jarrold C, Kane M, Miyake A, Towse J, editors. *Variation in Working Memory.* New York: Oxford University Press. p. 76–106.
- Braver TS, Reynolds JR, Donaldson DI. 2003. Neural mechanisms of transient and sustained cognitive control during task switching. *Neuron.* 39:713–726.
- Callicott JH, Mattay VS, Bertolino A, Finn K, Coppola R, Frank JA, Goldberg TE, Weinberger DR. 1999. Physiological characteristics of capacity constraints in working memory as revealed by functional MRI. *Cereb Cortex.* 9:20–26.
- Chevalier G, Deniau JM. 1990. Disinhibition as a basic process in the expression of striatal functions. *Trends Neurosci.* 13:277–280.
- Cohen JD, Perlstein WM, Braver TS, Nystrom LE, Noll DC, Jonides J, Smith EE. 1997. Temporal dynamics of brain activation during a working memory task. *Nature.* 386:604–608.
- Cools R. 2006. Dopaminergic modulation of cognitive function—implications for L-DOPA treatment in Parkinson's disease. *Neurosci Biobehav Rev.* 30:1–23.
- Cools R, Barker RA, Sahakian BJ, Robbins TW. 2003. L-DOPA medication remediates cognitive inflexibility, but increases impulsivity in patients with Parkinson's disease. *Neuropsychologia.* 41:1431–1441.
- Cools R, Barker RA, Sahakian BJ, Robbins TW. 2001. Mechanisms of cognitive set flexibility in Parkinson's disease. *Brain.* 124:2503–2512.
- Di Martino A, Scheres A, Margulies DS, Kelly AM, Uddin LQ, Shehzad Z, Biswal B, Walters JR, Castellanos FX, Milham MP. 2008. Functional connectivity of human striatum: a resting state fMRI study. *Cereb Cortex.* 18:2735–2747.
- Frank MJ, Loughry B, O'Reilly RC. 2001. Interactions between frontal cortex and basal ganglia in working memory: a computational model. *Cogn Affect Behav Neurosci.* 1:137–160.
- Fuster JM. 1990. Prefrontal cortex and the bridging of temporal gaps in the perception-action cycle. *Ann NY Acad Sci.* 608:318–329; discussion 330–316.
- Haber SN. 2003. The primate basal ganglia: parallel and integrative networks. *J Chem Neuroanat.* 26:317–330.
- Jonides J, Lewis RL, Nee DE, Lustig CA, Berman MG, Moore KS. 2008. The mind and brain of short-term memory. *Annu Rev Psychol.* 59:193–224.
- Jonides J, Nee DE. 2006. Brain mechanisms of proactive interference in working memory. *Neuroscience.* 139:181–193.
- Kastner S, Ungerleider LG. 2000. Mechanisms of visual attention in the human cortex. *Annu Rev Neurosci.* 23:315–341.
- Koechlin E, Hyafil A. 2007. Anterior prefrontal function and the limits of human decision-making. *Science.* 318:594–598.
- Koechlin E, Jubault T. 2006. Broca's area and the hierarchical organization of human behavior. *Neuron.* 50:963–974.
- Koechlin E, Ody C, Kouneiher F. 2003. The architecture of cognitive control in the human prefrontal cortex. *Science.* 302:1181–1185.
- Kouneiher F, Charron S, Koechlin E. 2009. Motivation and cognitive control in the human prefrontal cortex. *Nat Neurosci.* 12:939–945.
- Kriegeskorte N, Simmons WK, Bellgowan PS, Baker CI. 2009. Circular analysis in systems neuroscience: the dangers of double dipping. *Nat Neurosci.* 12:535–540.
- Lehericy S, Bardinet E, Tremblay L, Van de Moortele PF, Pochon JB, Dormont D, Kim DS, Yelnik J, Ugurbil K. 2006. Motor control in basal ganglia circuits using fMRI and brain atlas approaches. *Cereb Cortex.* 16:149–161.
- Lehericy S, Ducros M, Van de Moortele PF, Francois C, Thivard L, Poupon C, Swindale N, Ugurbil K, Kim DS. 2004. Diffusion tensor fiber tracking shows distinct corticostriatal circuits in humans. *Ann Neurol.* 55:522–529.
- Lewis SJ, Dove A, Robbins TW, Barker RA, Owen AM. 2004. Striatal contributions to working memory: a functional magnetic resonance imaging study in humans. *Eur J Neurosci.* 19:755–760.
- Lund TE, Norgaard MD, Rostrup E, Rowe JB, Paulson OB. 2005. Motion or activity: their role in intra- and inter-subject variation in fMRI. *Neuroimage.* 26:960–964.
- MacDonald AW, 3rd. 2008. Building a clinically relevant cognitive task: case study of the AX paradigm. *Schizophr Bull.* 34:619–628.
- McNab F, Klingberg T. 2008. Prefrontal cortex and basal ganglia control access to working memory. *Nat Neurosci.* 11:103–107.
- Moore T. 2006. The neurobiology of visual attention: finding sources. *Curr Opin Neurobiol.* 16:159–165.
- Nee DE, Brown JW, Askren MK, Berman MG, Demiralp E, Krawitz A, Jonides J. 2012. A meta-analysis of executive components of working memory. *Cereb Cortex.* Feb 7, [Epub ahead of print].
- Nee DE, Jonides J. 2009. Common and distinct neural correlates of perceptual and memorial selection. *Neuroimage.* 45:963–975.
- Nee DE, Jonides J. 2011. Dissociable contributions of prefrontal cortex and the hippocampus to short-term memory: evidence for a 3-state model of memory. *Neuroimage.* 54:1540–1548.
- Nee DE, Jonides J. 2008. Neural correlates of access to short-term memory. *Proc Natl Acad Sci USA.* 105:14228–14233.
- Nee DE, Kastner S, Brown JW. 2011. Functional heterogeneity of conflict, error, task-switching, and unexpectedness effects within medial prefrontal cortex. *Neuroimage.* 54:528–540.
- Nichols T, Brett M, Andersson J, Wager T, Poline JB. 2005. Valid conjunction inference with the minimum statistic. *Neuroimage.* 25:653–660.
- Nobre AC, Coull JT, Maquet P, Frith CD, Vandenberghe R, Mesulam MM. 2004. Orienting attention to locations in perceptual versus mental representations. *J Cogn Neurosci.* 16:363–373.
- Oberauer K. 2002. Access to information in working memory: exploring the focus of attention. *J Exp Psychol Learn Mem Cogn.* 28:411–421.
- Oppenheim AV, Schafer RW, Buck JR. 1999. *Discrete-time signal processing*, 2nd ed. Upper Saddle River, NJ: Prentice Hall.
- O'Reilly RC. 2010. The what and how of prefrontal cortical organization. *Trends Neurosci.* 33:355–361.
- O'Reilly RC, Frank MJ. 2006. Making working memory work: a computational model of learning in the prefrontal cortex and basal ganglia. *Neural Comput.* 18:283–328.

- Parent A, Hazrati LN. 1995. Functional anatomy of the basal ganglia. I. The cortico-basal ganglia-thalamo-cortical loop. *Brain Res Brain Res Rev.* 20:91–127.
- Picard N, Strick PL. 2001. Imaging the premotor areas. *Curr Opin Neurobiol.* 11:663–672.
- Postuma RB, Dagher A. 2006. Basal ganglia functional connectivity based on a meta-analysis of 126 positron emission tomography and functional magnetic resonance imaging publications. *Cereb Cortex.* 16:1508–1521.
- Reynolds JH, Chelazzi L. 2004. Attentional modulation of visual processing. *Annu Rev Neurosci.* 27:611–647.
- Reynolds JR, O'Reilly RC. 2009. Developing PFC representations using reinforcement learning. *Cognition.* 113:281–292.
- Reynolds JR, O'Reilly RC, Cohen JD, Braver TS. 2012. The function and organization of lateral prefrontal cortex: a test of competing hypotheses. *PLoS ONE.* 7:e30284.
- Rissman J, Gazzaley A, D'Esposito M. 2004. Measuring functional connectivity during distinct stages of a cognitive task. *Neuroimage.* 23:752–763.
- Rougier NP, Noelle DC, Braver TS, Cohen JD, O'Reilly RC. 2005. Prefrontal cortex and flexible cognitive control: rules without symbols. *Proc Natl Acad Sci USA.* 102:7338–7343.
- Serences JT, Shomstein S, Leber AB, Golay X, Egeth HE, Yantis S. 2005. Coordination of voluntary and stimulus-driven attentional control in human cortex. *Psychol Sci.* 16:114–122.
- Servan-Schreiber D, Cohen JD, Steingard S. 1996. Schizophrenic deficits in the processing of context. A test of a theoretical model. *Arch Gen Psychiatry.* 53:1105–1112.
- Tamber-Rosenau BJ, Esterman M, Chiu YC, Yantis S. 2011. Cortical mechanisms of cognitive control for shifting attention in vision and working memory. *J Cogn Neurosci.* 23:2905–2919.