

Prefrontal Cortex Organization: Dissociating Effects of Temporal Abstraction, Relational Abstraction, and Integration with fMRI

Derek Evan Nee¹, Andrew Jahn² and Joshua W. Brown²

¹Helen Wills Neuroscience Institute, University of California, Berkeley, CA 94720-3190, USA and ²Department of Psychological and Brain Sciences, Indiana University, Bloomington, IN 47405, USA

Address correspondence to Dr Derek Evan Nee, University of California, 132 Barker Hall, Berkeley, CA 94720-3190, USA. Email: denee@berkeley.edu

The functions of the prefrontal cortex (PFC) underlie higher-level cognition. Varying proposals suggest that the PFC is organized along a rostral-caudal gradient of abstraction with more abstract representations/processes associated with more rostral areas. However, the operational definition of abstraction is unclear. Here, we contrasted 2 prominent theories of abstraction—temporal and relational—using fMRI. We further examined whether integrating abstract rules—a function common to each theory—recruited the PFC independently of other abstraction effects. While robust effects of relational abstraction were present in the PFC, temporal abstraction effects were absent. Instead, we found activations specific to the integration of relational rules in areas previously shown to be associated with temporal abstraction. We suggest that previous effects of temporal abstraction were due to confounds with integration demands. We propose an integration framework to understand the functions of the PFC that resolves discrepancies in prior data.

Keywords: cognitive control, executive function, hierarchical, PFC, rule representation

Introduction

Humans have the remarkable ability to rapidly adjust behavior in accordance with contextual demands. This flexibility of action is thought to depend upon the functions of the prefrontal cortex (PFC). The centrality of the PFC in flexible action is evidenced by the disorganized, perseverative, and impulsive behaviors demonstrated by patients with damage to the PFC (Damasio and Anderson 1993). Hence, describing how the PFC enables dynamic and flexible behavior is paramount towards understanding intelligent cognition.

A challenge for understanding PFC function is how to explain its involvement in a diverse range of tasks that bear little surface resemblance to each other (Cabeza and Nyberg 2000; Duncan and Owen 2000). To meet this challenge, numerous theories have proposed organizing principles to describe the functions of the PFC (Christoff and Gabrieli 2000; Courtney 2004; Bunge and Zelazo 2006; Christoff and Keramatian 2007; Badre and D'Esposito 2009; Banich 2009; O'Reilly 2010; Nee et al. 2013). While PFC theories differ in a number of ways (for reviews, see Wood and Grafman 2003; Badre 2008), an emerging theme is that different areas of the PFC are responsible for processing and/or representing information at different levels of abstraction. Abstraction refers to the degree to which processing/representation is tied to or divorced from particular instances (e.g., “buy a car” is an abstract goal that can be realized through a number of concrete means). However, theories diverge in how abstraction is operationally defined. Here, we consider 2 potential forms of

abstraction that have received strong support from neuroimaging data: temporal abstraction and relational abstraction.

One prominent theory of PFC organization defines abstraction temporally (Koechlin et al. 2003; Koechlin and Summerfield 2007). Temporal abstraction follows from the notion that more abstract goals (e.g., “buy a car”) control responses for longer periods of time than more concrete subgoals (e.g., “go to dealership,” “test drive,” etc.). Under this view, responses that can be performed solely on the basis of existing environmental cues rely on caudal areas of frontal cortex, whereas responses requiring integration from temporally remote events recruit increasingly rostral areas of the PFC. To examine this proposal, Koechlin et al. (2003) varied whether responses could be determined on the basis of colored stimuli (sensory control) or whether color-to-response mappings depended on a temporally remote cue (episodic control). Demands on sensory and episodic control were manipulated by varying the information value of the stimuli and cues, respectively, where information value was defined as the degree to which stimuli/cues reduced the uncertainty of appropriate action. The authors found that increased demands on sensory control elicited activations in caudal frontal areas (dorsal premotor cortex, area 6), while increased demands on episodic control elicited activations in rostral lateral PFC (area 46; see also Kounieher et al. (2009) for a replication). The same effect of episodic control in rostral lateral PFC was observed when stimuli denoted a task (case or vowel-consonant judgments on letters) rather than mapping directly to a response, suggesting that the effect of episodic control is independent of task complexity. Hence, the authors argued that the rostral PFC is critically sensitive to maintaining a temporal signal to guide performance rather than merely responding to complexity (Koechlin and Summerfield 2007).

Another prominent account of PFC organization defines abstraction relationally (Christoff and Keramatian 2007). Under this view, responses based on simple stimulus features (e.g., “is this car big?”) involve caudal areas of the PFC. However, responses that require integrating relationships between features (e.g., “is this car as big as that one?”) recruit more rostral areas of the PFC. Early evidence for this account came from studies that employed the Raven's Progressive Matrices task (Christoff et al. 2001; Kroger et al. 2002). In this task, subjects were presented with a matrix of stimuli with 1 missing cell. The stimuli formed a pattern with one or more dimensions (e.g., shape, orientation, size) changing in an orderly manner across rows and down columns of the matrix. Based on the observed pattern, subjects were required to fill the missing cell. By manipulating the number of dimensions that needed to be integrated to solve the problems, these studies found that the PFC was sensitive to increased relational integration demands. This sensitivity to relational integration remained

after factoring out contributions of difficulty (Kroger et al. 2002) and time (Christoff et al. 2001), especially in the rostral-most areas of the PFC (frontopolar cortex, area 10). Similar patterns of activation have been noted using simpler tasks that have compared match/nonmatch decisions to nonrelational stimulus-based responses while also controlling for the duration over which such information was held in mind (Bunge et al. 2003). Hence, these studies suggest that the PFC is sensitive to demands to integrate relationships between information rather than simply holding information in mind over longer intervals.

It is important to note that temporal and relational abstractions are not necessarily mutually exclusive principles. Some of the strongest evidence to date of abstraction gradients in the PFC comes from a study that mixed both relational and temporal abstraction across different tasks (Badre and D'Esposito 2007). In that study, conditions requiring relational integration elicited more rostral activations in the PFC than conditions that did not, supporting relational abstraction accounts. However, even more rostral areas of the PFC (frontopolar cortex, area 10) were activated when episodic control demands were manipulated thereby demonstrating an effect of temporal abstraction. Therefore, it remains plausible that the PFC is sensitive to both temporal and relational abstraction.

A common factor across both abstraction accounts is the need for integration. In tasks that vary temporal abstraction, a temporally remote signal must be integrated within the context of the current sensory environment in order to determine the appropriate response. In tasks that vary relational abstraction, multiple stimulus features or relationships must be integrated in order to determine the appropriate response. So, it is possible that the PFC is sensitive to integration regardless of the type of information that needs to be integrated. Indeed, there is some evidence that the rostral-most areas of the PFC (frontopolar cortex, area 10) are specifically sensitive to integration demands (Ramnani and Owen 2004; Bunge et al. 2005; De Pisapia et al. 2007, 2012; Wendelken and Bunge 2010). However, to our knowledge, no study has separated demands of temporal abstraction, relational abstraction, and integration. As a result, the relative sensitivity of subregions of the PFC to each factor is unknown. Resolving this issue would enhance our understanding of the organization of the PFC.

The present study was designed to directly compare effects of temporal abstraction, relational abstraction, and integration. On each trial, subjects received 2 cues (see Fig. 1): one denoting a concrete feature (feature cue; e.g., "green") and another denoting an abstract relationship (relation cue; e.g., "and"). The order of cues was counterbalanced. Thereafter, subjects were shown 2 objects (probe) and made a decision on the basis of the feature and relationship (e.g., "are both objects green?"). Critically, cues and responses were separated in time so that the hemodynamic response for each event could be independently estimated. In previous research, a similar 2-cue design demonstrated dissociable rostral-caudal regions sensitive to each cue, but did not disentangle different forms of abstraction (Nee and Brown 2012a). By contrast, the present design enabled the examination of whether areas of the PFC are preferentially sensitive to temporal abstraction, relational abstraction, and/or integration (see Fig. 1). Furthermore, the information value of relational cues was manipulated (i.e., the degree to which a cue predicted the forthcoming response) since temporal abstraction is also described in terms of

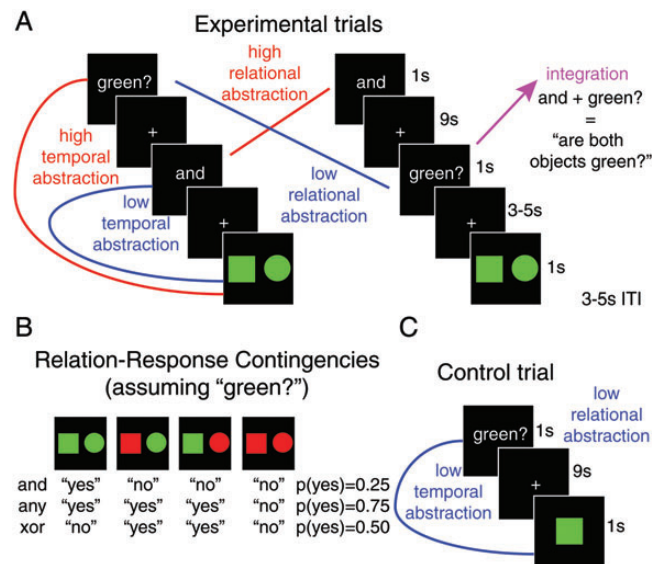


Figure 1. Depiction of the task. (A) On experimental trials, subjects were presented with 2 colored shape stimuli (probe) and made responses based upon the integration of 2 cues. The cues were separated in time so that the first cue was more temporally remote to the probe (i.e., occurred earlier) than the second, and therefore was more temporally abstract (i.e., first cue: high temporal abstraction, second cue: low temporal abstraction). One cue denoted a critical concrete feature (feature cue: low relational abstraction) and the other cue denoted a relational rule to perform (relation cue: high relational abstraction). In the depicted example, the cue "and" along with the cue "green?" indicates that subjects should make a "yes" response if both objects are green and a "no" response otherwise. (B) The 3 types of relation cues are depicted along with appropriate responses to example stimuli. "And" cues required a "yes" response if both objects contained the relevant feature and a "no" response otherwise. "Any" cues required a "yes" response if either or both objects contained the relevant feature and a "no" response otherwise. "Xor" cues required a "yes" response if only one object contained the relevant feature and a "no" response otherwise. "And" and "any" cues were informative since they indicated that a particular forthcoming response was more likely (i.e., reduced entropy) while "xor" cues were uninformative. The prior probability of a "yes" response varies systematically depending on which cue is presented. (C) On control trials, subjects responded to a single stimulus based upon a single feature cue. Compared with experimental trials, cues in control trials had low temporal abstraction, low relational abstraction, and required no integration.

information theory (Koechlin et al. 2003; Koechlin and Summerfield 2007). Together, our design permitted the assessment of several factors hypothesized to involve the PFC in order to disentangle their relative contributions.

Materials and Methods

Participants

We report results from 22 (12 females) right-handed native English speakers with normal or corrected-to-normal vision (mean age 24.3 years, range 20–31 years). Four additional participants performed a behavioral practice session but were excluded from the fMRI session for inability to follow task instructions. Informed consent was obtained in accordance with the Institutional Review Board at Indiana University. Subjects were compensated at a rate of \$25/h for participation plus a performance based bonus (mean \$4.46; range \$2.89–\$5.40).

Procedure

The task was designed to contrast effects of temporal abstraction, relational abstraction, and integration (Fig. 1). We also included a manipulation of information value to examine whether any effects were driven or modulated by the information value of cues. Subjects

responded to colored shapes based on cues. On experimental trials, subjects received 2 sequentially ordered cues. One cue denoted a particular feature (feature cue). A second cue denoted a relation (relation cue). Following presentation of the 2 cues, 2 probe objects appeared side-by-side and subjects made a yes/no decision regarding whether the objects held the appropriate relation with respect to the cued feature. Feature cues could indicate a target color (red or green) or shape (square or circle). Three types of relations were probed: “and” relations required that both probe objects contained the cued feature; “any” relations required that any one or both of the probe objects contained the cued feature; “xor” relations required that one and only one of the probe objects contained the cued feature (i.e., exclusive-or). The separation of cues by time in experimental trials provided a manipulation of “temporal abstraction” with the first cue providing a more temporally abstracted episode relative to the second cue. Concrete feature cues denoted a zero-order relation while abstract relation cues denoted a first-order relation thereby providing a manipulation of “relational abstraction.” The separation of cue events afforded the examination of “integration” during the presentation of the second cue when the meanings of the 2 cues were integrated. Finally, relation cues provided a manipulation of “information” as each relation cue had different probabilities of eliciting a “yes” or “no” response in accordance with feature matching probabilities (yes/no probability: “and”—0.25/0.75, “any”—0.75/0.25, “xor”—0.5/0.5). We also included control trials to act as an additional potential baseline. On control trials, subjects received only a feature cue. After a delay, a single probe object appeared and subjects made a yes/no decision regarding whether the object contained the cued feature. Control trials therefore required reduced temporal abstraction, relational abstraction, and integration demands relative to experimental trials.

All feature cues and all relation cues were presented with equal frequencies and randomized throughout the experiment. On experimental trials, the order of feature and relation cues was randomly counterbalanced. Subjects responded with the index finger of either hand with hand-to-response mapping counter-balanced between subjects. All cues and probes were presented for 1 s. On experimental trials, the first and second cues were separated by 9 s, and the probe was separated from the second cue by a jittered 3–5 s interval. On control trials, the cue and probe were separated by 9 s. All trials were separated by a jittered 3–5 s interval. Long constant intervals were used in order to assess both transient cue-related responses as well as sustained responses during delay intervals. Owing to time constraints, a shorter jittered interval between the second cue and the probe was used in order to statistically separate hemodynamic responses to the second cue and probe.

Control and experimental trials were separated into miniblocks of 4 trials each. Each mini-block began with an instruction (“Single Cue” or “Double Cue”) to alert subjects of the forthcoming trial type. After every third mini-block, subjects received the instruction “Fixate” followed by a 9-s fixation interval. There were twice as many experimental miniblocks as control miniblocks. Subjects performed 5 runs of 36 trials each divided into 9 miniblocks per run. The entire experiment consisted of 180 trials with 120 experimental trials and 60 control trials. The feature cue was presented first on 60 experimental trials, and the relation cue was presented first on the other 60 experimental trials. There were 40 trials each of each type of relation cue.

Within a week prior to scanning, subjects performed a behavioral practice session in order to familiarize themselves with the task. The behavioral practice session was identical to the scanning session except that subjects were seated at a computer instead of within the scanner. The day of scanning, the subjects performed a single warm-up run prior to scanning. The extensive practice was aimed to ensure that the task was well learned prior to scanning and to minimize learning effects during the scanning session.

Imaging Acquisition and Preprocessing

Images were acquired on a 3T Siemens Tim Trio equipped with a 32-channel head coil. Stimuli were presented to the subject via a projector at the rear of the scanner, reflected off a mirror mounted to the

head coil. Experimental tasks were presented using E-Prime software version 1.2 (Psychology Software Tools, Inc., Pittsburgh, PA, USA).

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 (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. High-resolution T_1 -weighted MPRAGE images were collected for spatial normalization ($384 \times 384 \times 256$ matrix of $0.667 \times 0.667 \times 0.7 \text{ mm}^3$ voxels; TR = 1800 ms; echo time = 2.82 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 and head movement using a least-squares approach and a 6-parameter rigid body spatial transformation. Images were corrected for distortion and movement-by-susceptibility artifacts using the FieldMap toolbox implemented in SPM5 (Andersson et al. 2001). 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/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

General Linear Model

fMRI data were analyzed using the general linear model implemented in SPM5. The model included transient, impulse regressors for cue and probe events. For the cue–cue interval in experimental trials and cue–probe interval in control trials, epoch regressors were included to capture delay period activation. These delay period regressors were segregated by surrounding cue and probe events by 4 s to permit separate assessment of cue, probe, and delay events (Zarahn et al. 1997). For experimental trials, separate cue regressors were included for each combination of relational abstraction (feature cue, relation cue) and time (first cue, second cue). Relation cues were further broken down by type with separate regressors for “and,” “any,” and “xor” cues. These combinations resulted in 8 cue regressors for experimental trials (feature first cue, “and” first cue, “any” first cue, “xor” first cue, feature second cue, “and” second cue, “any” second cue, “xor” second cue). Four regressors modeled the delay intervals between cues, one each for cue type (feature delay, “and” delay, “any” delay, “xor” delay). Six regressors modeled the probe, one for each combination of “yes” and “no” responses and relations. For control trials, the model included 1 regressor for the cue event (control cue), 1 for the delay interval (control delay), and 2 for the probe, separately estimating “yes” and “no” trials. Events from trials in which an error occurred were modeled separately and excluded from further analysis. One additional regressor was included to model the baseline fixation epochs. Nuisance regressors were included to model the instructions at the beginning of each miniblock, as well as block regressors to capture global run effects. All regressors (excluding the run regressors) were convolved with the canonical hemodynamic response function included in SPM5. Parameters estimated from this first-level model were then entered into second-level ANOVAs and *t*-tests for random effects analysis as described below.

PFC Search Space

In all analyses, we restricted our investigations to task-sensitive voxels. Task sensitivity was defined through a contrast of the average of all task events (cues, delays, and probes) minus the fixation baseline thresholded at $P < 0.05$ at the voxel level (uncorrected). This task-sensitive mask is depicted in Supplementary Figure 1A. As we were primarily interested in the lateral PFC, a PFC mask was created to

hone in on critical regions of interest (ROIs). The mask was created by generating a 10-mm sphere around lateral PFC peak activations from previous studies that have investigated effects of abstraction (Koechlin et al. 1999, 2003; Christoff et al. 2001, 2003, 2009; Bunge et al. 2003, 2005; Koechlin and Jubault 2006; Badre and D'Esposito 2007; Nee and Brown 2012a; Reynolds et al. 2012). As reported activations were overwhelmingly left lateralized, we extracted only the left hemisphere peaks and then reflected them around the x -axis to make the mask symmetrically bilateral. The PFC mask is depicted in Supplementary Figure 1B. In the analyses described below, we first looked for activations within the a priori mask thresholded at $P < 0.05$ corrected for multiple comparisons using false discovery rate (Genovese et al. 2002) with a 20-voxel extent criterion. Exploratory whole-brain analyses were then performed to uncover any activation missed by the masking procedure. These whole-brain analyses were thresholded at $P < 0.001$ at the voxel level with a 50-voxel extent providing cluster corrected threshold of $P < 0.05$ according to simulations using AlphaSim.

Temporal and Relational Abstraction and Integration

The primary analysis consisted of a 2×2 ANOVA on cue activations with factors of temporal and relational abstraction. This analysis was performed on cue events from experimental trials. For simplicity, different relation types were collapsed together resulting in the following cue types: feature first cue, relation first cue, feature second cue, relation second cue. Within the ANOVA framework, effects of temporal abstraction were explored by contrasting the first cue with the second cue, collapsing across relational abstraction (first cue > second cue). A follow-up analysis contrasted the sum of both types of first cues with the second cue (feature first cue + relation first cue > second cue) to better match working memory demands. Effects of relational abstraction were explored by contrasting relation cues with feature cues, collapsing across temporal abstraction (relation cue > feature cue). We also examined the interaction of temporal and relational abstraction factors. Next, we examined the effect of integrating cue information. This contrast consisted of looking for an overadditive effect (second cue > feature first cue + relation first cue). The logic of this contrast was that a region specialized for integration should show activation to the second cue that is greater than the sum of the individual cues. Finally, to assess whether any region demonstrated effects of both relational abstraction and integration, we performed valid conjunction analysis (Nichols et al. 2005).

As the primary analysis failed to identify areas sensitive to temporal abstraction, we conducted a second analysis to further examine potential effects of temporal abstraction. We contrasted feature cues drawn from the first cue phase of experimental trials with control cues. These cues were well matched visually, semantically, and in working memory load, but cues from the experimental trials elicited greater temporal abstraction relative to control cues. This analysis consisted of a 1-sample t -test. As this analysis also failed to find significant activations related to temporal abstraction, we also performed a similar analysis of the delay period activation comparing sustained activation following feature cues in experimental trials with sustained activation following control cues in control trials.

Independent ROIs

We created unbiased ROIs for visualization and to test sustained activations during the cue-cue interval. To do so, we examined the effect of relational abstraction, as well as the effect of integration using a leave-one-subject out procedure. For each contrast, a 1-sample t -test was performed on 21 of the 22 subjects. Spherical ROIs (6 mm) were then drawn around activation peaks for the respective contrast. Data from the held-out subject were extracted from the resultant independent group-defined ROIs, and the procedure was rotated to examine activations from each subject. ROIs from the relational abstraction contrast were drawn from the left insula (mean peak: -48.1 7.9 -0.5), left lateral premotor cortex (mean peak: -47.8 0.7 42.1), left inferior frontal gyrus, pars orbitalis (mean peak: -37.3 27.9 -3), left inferior frontal junction (mean peak: -46.0 0.1 33.8), and left rostral lateral PFC (mean peak: -45.7 36.5 16.5). ROIs from the integration contrast were drawn from the left caudal superior frontal

sulcus (mean peak: -31.4 -10.1 46.2) and left frontopolar cortex (mean peak: -35.6 51.8 16.1).

Time-Course Data

To examine the time course of activation in more detail, we re-estimated the data using a hybrid finite impulse response (FIR) model. In this model, 14 TR-by-TR delta regressors were included that spanned cue and delay events. For simplicity, all relation cues ("and," "any," "xor") were collapsed together. This resulted in 2 main conditions: feature->relation and relation->feature. Probe events were modeled as in the main text with separate regressors for "yes" and "no" trials convolved with a canonical hemodynamic response function. In this way, response-related information was regressed out of the FIR time courses. Time-course data were examined in the independent ROIs described above.

Temporal Abstraction ROIs

ROI analyses were performed to further explore effects of temporal abstraction, relational abstraction, and integration. As the analyses described above revealed effects of relational abstraction and integration, but not temporal abstraction, we examined ROIs drawn from studies that have provided evidence for temporal abstraction (Koechlin et al. 1999, 2003). Spherical ROIs (6 mm) were placed around peaks reported in frontopolar cortex (area 10; -36 57 9 and 36 66 21 ; Koechlin et al., 1999) and rostral lateral PFC (area 46; -40 32 20 and 32 32 20 ; Koechlin et al., 2003). Owing to signal dropout, activation in the right frontopolar cortex region could not be measured in 4 subjects, so analyses of the right frontopolar cortex were based on the remaining 18 subjects.

Information Theoretic Analysis

According to Koechlin et al. (2003), greater episodic control is associated with cues that provide greater information for action selection. Information is provided when the uncertainty (entropy) for selecting an action is reduced. With no additional information, the probability of a "yes" response is 0.5 and the probability of a "no" response is 0.5. Entropy can be computed as:

$$H(A) = - \sum_{a \in A} p(a) \log_2(p(a))$$

where A is the set of actions, a denotes a particular action, $P(a)$ is the probability of selecting action a , and $H(A)$ is the uncertainty (i.e., entropy) associated with selecting actions. In the present task, "yes" and "no" responses were equally likely ($P(A = \text{"yes"}) = P(A = \text{"no"}) = 0.5$). Hence, when no other information is known, entropy is 1 (bit). When a given cue provides guidance about how to respond, entropy is reduced to:

$$H(A|C = c) = - \sum_{a \in A} p(a|c) \log_2 p(a|c)$$

where c denotes a particular cue. Then, the amount of information given by the cue c is:

$$I(A, C = c) = H(A) - H(A|c)$$

In the present task, we are interested in the information provided by the cue when it is presented. This differs somewhat from Koechlin et al. (2003) where cues provided no information upon presentation, but provided essential information during integration. Here, we are interested in the information value of the cues themselves prior to integration. With this in mind, we can calculate the information value of the 3 types of relation cues:

$$\begin{aligned} I(A, \text{and}) &= 0.19 \\ I(A, \text{any}) &= 0.19 \\ I(A, \text{xor}) &= 0 \end{aligned}$$

Thus, "and" and "any" cues are equally informative while "xor" cues are uninformative. As a result, we investigated the effect of

information value by contrasting the average of “and” and “any” cues with “xor” cues ($(\text{“and”} + \text{“any”})/2 - \text{“xor”}$). This contrast was performed in the context of a 3×2 ANOVA with factors of relation type and time. Analyses of information collapsed across the time factor.

Results

Behavioral Results

Subjects performed the task well overall (mean error rate = 6.5%). Reaction time (RT) data were analyzed on correct trials only. As expected, subjects were less error-prone and faster on control trials (5.5%, 545.0 ms) than experimental trials (7.6%, 579.5 ms; error rate: $t_{(21)} = 2.73$, $P < 0.05$; RT: $t_{(21)} = 5.07$, $P < 0.0001$). Subjects also performed better on “yes” trials (5.4%, 543.0 ms) than “no” trials (7.6%, 582.1 ms; error rate: $t_{(21)} = 2.46$, $P < 0.05$; RT: $t_{(21)} = 5.33$, $P < 0.0001$). The performance difference between “yes” and “no” trials did not differ between control and experimental trials (error rate: $t_{(21)} = 0.68$, $P > 0.5$; RT: $t_{(21)} = 1.42$, $P > 0.15$).

The following analyses were restricted to experimental trials. The effect of cue order was not significant, indicating that whichever cue was presented first (feature cue or relation cue) had no effect on performance (error rate: $t_{(21)} = 0.22$, $P > 0.8$; RT: $t_{(21)} = 1.83$, $P > 0.05$). Next, we examined the effect of different relation cues. Relation cues differed in the extent to which the forthcoming response could be anticipated. “And” and “any” cues were informative, while “xor” cues were not (see Fig. 1; Materials and Methods section). A 1-way ANOVA on the 3 different types of relation cues revealed a significant effect in error rate ($F_{2,42} = 7.25$, $P < 0.005$), but not RT ($F_{2,42} = 2.18$, $P > 0.1$). Follow-up tests revealed that the effect in error rate was driven by more error-prone performance on “xor” trials relative to “and” trials (9.8% vs. 6.1%; $t_{(21)} = 3.02$, $P < 0.01$) and “any” trials (9.8% vs. 5.4%; $t_{(21)} = 2.97$, $P < 0.01$). However, error rates on “and” trials and “any” trials did not differ ($t_{(21)} = 0.76$, $P > 0.45$). Similar trends were present in the RT data that did not reach significance (“xor” RT > “and” RT: $t_{(21)} = 1.84$, $P = 0.08$; “xor” RT > “any” RT: $t_{(21)} = 1.86$, $P = 0.08$). This pattern is consistent with information theoretic analyses that indicate that “and” and “any” cues provide equivalent advanced information about forthcoming response probabilities (i.e., both reduce entropy) while “xor” cues provide no information (see Materials and Methods section for information theoretic calculations).

Imaging Results

We examined effects of temporal abstraction, relational abstraction, and integration as cues were presented to subjects prior to responding. Cue-related activations were submitted to a 2×2 ANOVA with factors of temporal abstraction and relational abstraction. Greater temporal abstraction was associated with the first cue relative to the second cue while greater relational abstraction was associated with the relation cue relative to the feature cue.

Temporal Abstraction

Within a priori-defined PFC mask (see Supplementary Fig. 1), only a single region demonstrated effects of temporal abstraction. Close inspection revealed that this region was located in the right putamen rather than the PFC (the a priori mask covered part of the putamen so as to cover nearby insular

peaks). Exploratory whole-brain analyses revealed a more substantial right putamen cluster as well as one in the left hemisphere (Supplementary Table 1). However, no additional clusters outside of the striatum were found.

One issue with the temporal abstraction contrast is that at the presentation of the second cue, the first cue was actively maintained. Hence, comparing the first cue directly to the second cue confounds working memory load, which is greater at the time of the second cue. As working memory maintenance is known to implicate the PFC (Curtis and D’Esposito 2003; Courtney 2004), the presence of a potential temporal abstraction effect may have been masked by working memory load. To overcome this issue, we “summed” parameter estimates from the feature first cue and relation first cue and compared them with the second cue, which was an “average” of the feature second cue and relation second cue. Assuming linearity, this contrast should control for working memory load, at the risk of being slightly liberal since working memory-related PFC responses are typically elevated during encoding events and decay slightly during delay intervals (Curtis and D’Esposito 2003). However, even this potentially liberal contrast did not reveal any voxels sensitive to temporal abstraction within the a priori PFC mask. Exploratory whole-brain analysis revealed clusters in the bilateral fusiform gyrus, bilateral parahippocampal gyrus, bilateral occipital cortex, and right amygdala.

To explore the lack of PFC involvement in more detail, we conducted a second analysis that contrasted the first feature cue from experimental trials with cues drawn from control trials. These cues were semantically and visually identical and matched for working memory load with the lone difference being that cues from experimental trials were held in mind over longer periods of time and thus had a higher degree of temporal abstraction. This analysis did not reveal any regions sensitive to temporal abstraction within the a priori PFC mask or elsewhere.

Finally, as temporal abstraction effects may show sustained rather than transient cue-related effects, we examined delay period activation comparing sustained activation following feature cues in experimental trials with sustained activation following control cues in control trials. This analysis also did not reveal any voxels sensitive to temporal abstraction in the PFC. A whole-brain search revealed areas in medial occipito-temporal cortex, left fusiform, and left postcentral gyrus demonstrating sustained effects of temporal abstraction.

Taken together, 4 analyses hypothesized to measure effects of temporal abstraction failed to find PFC regions sensitive to temporal abstraction. While some of these analyses revealed more posterior structures sensitive to temporal abstraction, these areas were not consistent between analyses. Given this inconsistency and lack of theoretical framework in which to interpret them, we will not discuss these areas further.

Relational Abstraction

In contrast to the lack of temporal abstraction effects, several left-lateralized PFC regions within the a priori mask demonstrated effects of relational abstraction (Fig. 2; Table 1). Starting caudally, effects of relational abstraction were observed in the lateral premotor cortex (area 6; $-50 -4 42$) and the inferior frontal junction (areas 6 and 44; $-44 0 30$) and extended ventrally into the insula (area 13; $-48 8 0$) and the left inferior frontal gyrus, pars orbitalis (area 47; $-38 26 -2$). Insular effects were also present in the right hemisphere. Effects of relational abstraction were observed in the rostral

lateral PFC (area 46; $-44\ 38\ 20$), as well. Exploratory whole-brain analyses revealed additional clusters in the presupplementary motor area (pre-SMA), right cerebellum, right fusiform gyrus, left temporoparietal junction, left thalamus, right anterior cingulate cortex, and left posterior cingulate cortex (Supplementary Table 1).

Activation profiles of different PFC subregions are depicted in Figure 2A–D segregated by cue order (feature->relation or relation->feature). Data are depicted from unbiased ROIs (see Materials and Methods section). As can be seen in Figure 2B, the lateral premotor cortex (area 6) demonstrated increased activation to the initial presentation of the relation first cue relative to the feature first cue. This activation difference was sustained through the delay interval, but reversed at the time of the second cue at which time there was greater activation for the relation second cue than feature second cue. A similar reversal was observed in the rostral lateral PFC (area 46) as depicted in Figure 2A. However, in this area, differences were not sustained during the delay interval. Sustained activity was also not observed in other areas of the PFC. To investigate this pattern in more detail, we contrasted parameter estimates of the cue–cue delay interval following relation versus feature cues. An ANOVA on contrast estimates (i.e., relation delay > feature delay) drawn from the lateral premotor cortex, rostral lateral PFC, inferior frontal junction, insula, and inferior frontal gyrus, pars orbitalis revealed a significant effect of region ($F_{4,84} = 4.03$, $P < 0.005$). While the lateral premotor cortex demonstrated increased sustained activation following relation cues relative to feature cues ($t_{(21)} = 1.86$, $P < 0.05$, 1-tailed), no other PFC region demonstrated a similar effect (all $t_{(21)} < 1.36$, $P > 0.05$, 1-tailed). Whole-brain examination revealed additional activations in bilateral occipital cortex demonstrating sustained maintenance of relation cues relative to feature cues (left: $-20\ -90\ -10$, peak $z = 4.82$, 123 voxels; right: $-16\ -92\ -10$, peak $z = 4.36$, 93 voxels). This indicates that the rostral lateral PFC and other PFC subregions were involved in the initial representation/selection of relationally abstract information, but that sustained maintenance of this information was restricted to the lateral premotor cortex and regions outside the PFC search space.

Abstraction Interaction

No region demonstrated an interaction between temporal and relational abstraction either in the a priori PFC mask or elsewhere.

Integration

In order to explore effects of integration, we examined activations to the second cue when both feature and relation cues were integrated together. These second cue activations were compared with the summed activation of feature and relation cues from the first cue phase. The logic of this contrast was that activations to the feature first cue and relation first cue provided an assay of the cue-related activation independently of integration. By summing these parameter estimates and comparing them to the second cue (averaged across feature and relation second cues), this analysis examined whether any region was sensitive to integration over and above to the activations to the individual cues themselves (Ramnani and Owen 2004). This analysis revealed activations in the bilateral caudal superior frontal sulcus/dorsal premotor cortex (area 6; $-28\ -8\ 50$; $36\ -6\ 54$), bilateral insula (area 13; $-30\ 20\ -6$; 32

$24\ -8$), and left rostral lateral PFC/frontopolar cortex (area 46, 10 ; $-34\ 50\ 22$). Interestingly, the areas sensitive to integration straddled those sensitive to relational abstraction (Fig. 2; Table 1) with integration effects present both caudal to and rostral to relational abstraction effects. Exploratory whole-brain analyses revealed additional clusters in the bilateral thalamus, bilateral occipital cortex, dorsal anterior cingulate cortex extending into the pre-SMA, and left rostral intraparietal sulcus (Supplementary Table 1).

Figure 2C,D depicts activation profiles of the caudal superior frontal sulcus (area 6) and frontopolar cortex (area 46, 10), respectively, drawn from unbiased ROIs (see Materials and Methods section). In contrast to the regions sensitive to relational abstraction, these areas did not demonstrate a cross-over pattern. Instead, activation in these areas ramped up starting from the offset of the first cue and peaking after the presentation of the second cue when integration occurred.

Relational Abstraction and Integration

To assess whether any regions were jointly sensitive to the effects of relational abstraction and integration, we performed valid conjunction analysis (Nichols et al. 2005). Within the a priori PFC mask, the left anterior insula/inferior frontal gyrus, pars orbitalis ($-30\ 24\ 0$; area 13, 47) demonstrated sensitivity to both relational abstraction and integration. This region showed both a cross-over cue pattern emblematic of relational abstraction-related regions and a ramp up and postsecond cue peak seen in integration-related regions. An exploratory whole-brain search revealed that the left thalamus was also jointly sensitive to relational abstraction and integration.

Information

Effects of temporal abstraction are also described in terms of information theory (Koechlin et al. 2003; Koechlin and Summerfield 2007). By this account, the cognitive control provided by a temporally abstracted episodic signal is proportional to the information value of the temporally remote cue. In the present task, the information value of cues varied: “and” and “any” cues provided information that reduced uncertainty (entropy), while “xor” cues did not (see Materials and Methods section for full details on information theoretic calculations). Hence, the observed effects of relational abstraction may have been carried by the information value of “and” and “any” cues. To investigate this possibility, we contrasted the average effect of “and” and “any” cues with “xor” cues. This contrast did not reveal any significant differences within the a priori PFC mask or elsewhere in the brain. As a result, it is unlikely that the relational abstraction results are reducible to the information value of cues.

Regions of Interest

When observed together, the regions sensitive to relational abstraction and integration correspond quite closely to the areas described by Koechlin et al. (2003) (see Supplementary Fig. 2). Nevertheless, it may be possible that our analyses missed regions involved in temporal abstraction. Hence, to further investigate effects of temporal abstraction, relational abstraction, and integration, we explored activation within ROIs previously reported to demonstrate effects consistent with temporal abstraction. ROIs were placed in the bilateral frontopolar cortex (area 10) around peaks reported by Koechlin et al. (1999), and in the bilateral rostral lateral PFC (area

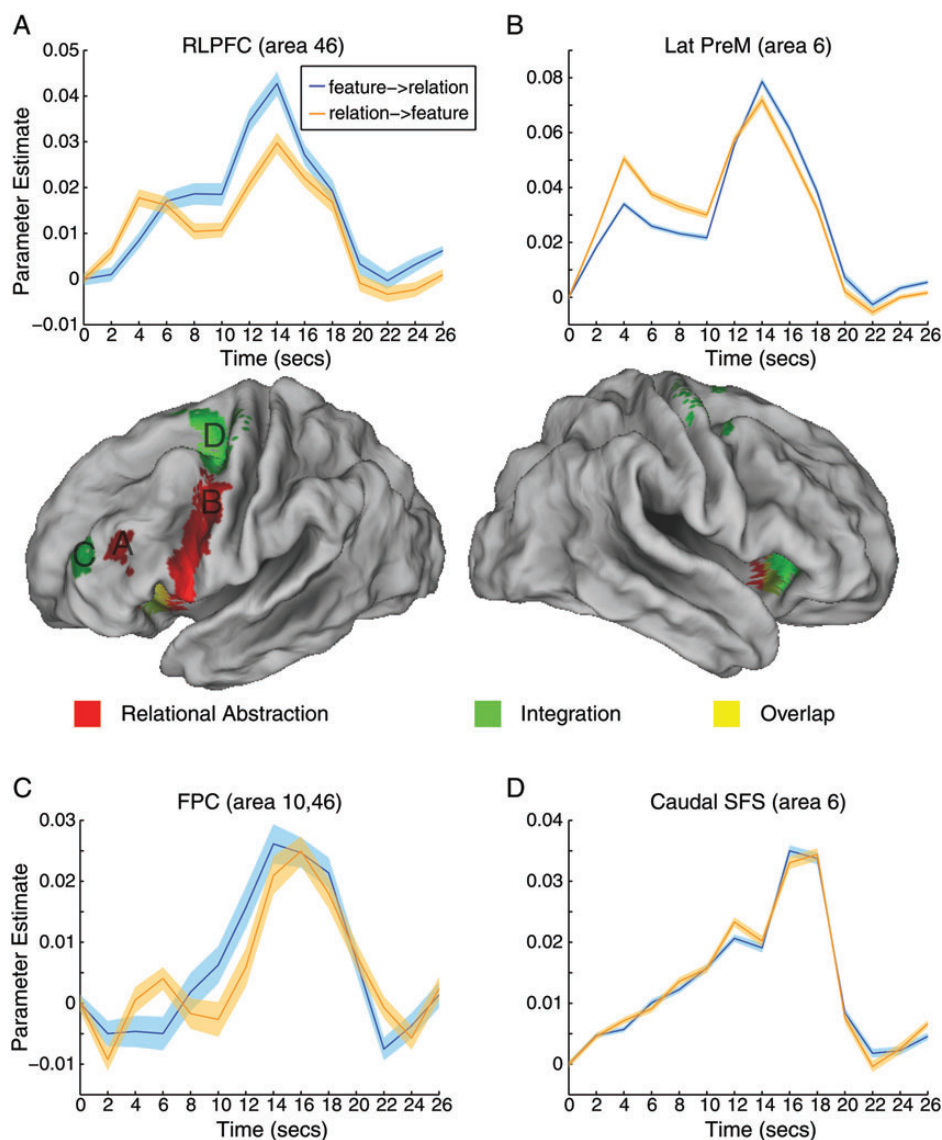


Figure 2. Neural correlates of relational abstraction and integration. Areas responsive to relational abstraction (red) were found along the lateral premotor cortex (A; Lat Premotor, area 6), extending ventrally into the inferior frontal gyrus and insula. A separate activation cluster was found in the rostral lateral prefrontal cortex (B; RLPFC, area 46). Areas responsive to integration (green) were found in the frontopolar cortex (C; FPC, area 10, 46) and caudal superior frontal sulcus (D; caudal SFS, area 6). Graphs (A–D) depict activation time courses drawn from unbiased regions of interest (see Supplementary Experimental Procedures). The first cue was presented at time 0 s and the second cue was presented at time 10 s. Probes, which were presented at time 14 to 16 s, were regressed out of the depicted graphs. Blue lines correspond to trials in which the feature cue was presented first and the relation cue second. Orange lines correspond to trials in which the relation cue was presented first and the feature cue second. Colored bands represent one standard error of the mean.

46) around peaks reported by Koechlin et al. (2003) (Fig. 3). None of the 4 ROIs demonstrated effects of temporal abstraction with any of the 4 different analysis methods described above (max $t = 1.51$, $P > 0.1$). By contrast, both the left and the right frontopolar cortex demonstrated effects of integration (left: $t_{(21)} = 2.20$, $P < 0.05$; right: $t_{(17)} = 2.40$, $P < 0.05$). Additionally, the left rostral lateral PFC demonstrated an effect of relational abstraction ($t_{(21)} = 2.60$, $P < 0.05$). Finally, no regions were sensitive to information value (all $t < 1$). Hence, these analyses replicated the effects described above.

Discussion

The present study compared effects of temporal abstraction, relational abstraction, and integration within the PFC. In

contrast to suggestions that the PFC is sensitive to temporal abstraction, our analyses failed to find effects of temporal abstraction in the PFC. We also failed to find evidence to support related suggestions that the PFC is sensitive to the information value of episodic signals. On the other hand, robust effects of relational abstraction were found in several lateral PFC areas that closely resembled previous reports. Finally, effects specific to the integration of contextual cues were found both rostral and caudal to regions sensitive to relational abstraction. We now reconsider putative models of PFC organization in light of these data.

Temporal Abstraction Revisited

In an influential study, Koechlin et al. (2003) demonstrated a cascade of rostral-caudal lateral frontal regions activated in

Region	Area	x	y	z	Vox	z-Score
Relational abstraction						
Left insula	13	-48	8	0	1656	6.27
Left lat premotor	6	-50	-4	42		4.94
Left IFG, orb	47	-38	26	-2		4.04
Left IFJ	6, 44	-44	0	30		3.54
Left insula	13	-32	16	-6		3.37
Left rostral MFG	46	-44	38	20	86	3.6
Right insula	13	32	16	-8	125	3.47
Left pre-SMA	6	-12	8	58		
Right insula	13	32	14	-10		
Right insula	13	46	10	0	22	2.47
Integration						
Left caudal SFS	6	-28	-8	50	788	6.27
Left pre-SMA	6	-14	6	64		4.22
Left pre-CG	6, 4	-36	-16	56		3.75
Left insula/IFG, orb	13, 47	-30	20	-6	283	4.99
Right insula/IFG, orb	13, 47	32	24	-8	309	4.87
Right dorsal premotor	6	36	-6	54	293	4.77
Right caudal SFS	6	32	0	50		3.77
Right pre-SMA	6	16	0	68		3.76
Right pre-CG	6	22	-18	68		3.07
Left rostral MFG/FPC	46, 10	-34	50	22	65	3.23

Note: FPC, frontopolar cortex; IFG, orb, inferior frontal gyrus, pars orbitalis; IFJ, inferior frontal junction; lat, lateral; MFG, middle frontal gyrus; pre-CG, precentral gyrus; pre-SMA, presupplementary motor area; SFS, superior frontal sulcus; vox, number of voxels.

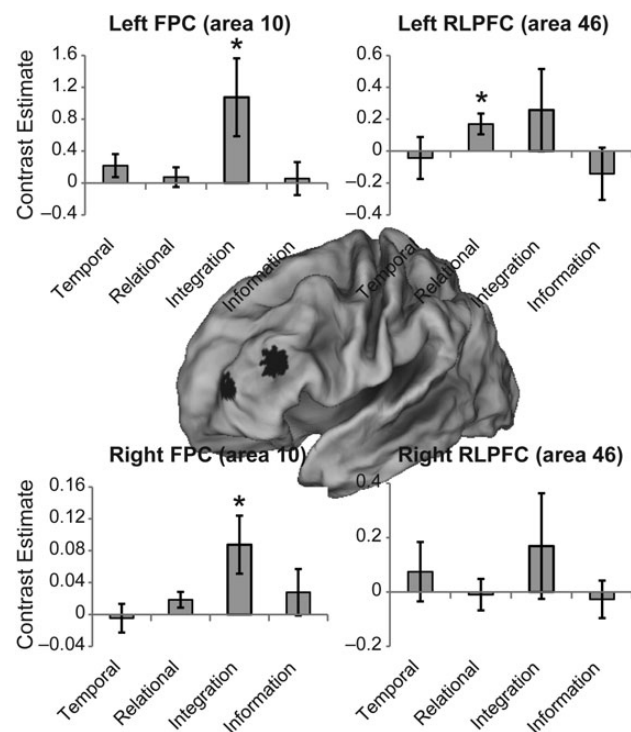


Figure 3. Effects within regions of interest. Regions of interest were placed in the bilateral rostral prefrontal cortex (RLPFC, BA 46) and bilateral frontopolar cortex (FPC, area 10) centered around peaks reported in previous studies that have documented effects of temporal abstraction. Bars represent contrasts described in the text. For temporal abstraction, the graphs depict the contrast of the first cue with the control cue, which tended to produce the strongest (nonsignificant) effects of temporal abstraction. The bilateral FPC demonstrated significant effects of integration while the left RLPFC demonstrated a significant effect of relational abstraction. No region demonstrated effects of temporal abstraction or information. * $P < 0.05$.

response to different levels of cognitive control. While these authors suggested that activations in caudal (area 6) and mid-caudal (area 44) regions were related to cognitive control demands manipulated by the sensory/contextual environment, more rostral activations (area 46) were hypothesized to be driven by a temporally abstracted episodic signal. The idea of a temporally driven cascade of control synergized well with previous data from the same group that demonstrated that information held in mind from a pending episode activated the rostral-most areas of the PFC (area 10; Koechlin et al. 1999). Hence, by their account, caudal areas of the frontal lobes processed current sensory and contextual signals, more rostral areas processed episodic signals, and the rostral-most areas processed pending episodic signals (Koechlin and Summerfield 2007).

More recent data have demonstrated inconsistent effects of temporal abstraction. In an elegant study, Badre and D'Esposito (2007) manipulated 4 levels of abstraction within a single experimental session. As in Koechlin et al. (2003), at the lowest level of abstraction, responses were determined by colored frames. At the second level, responses were based on the presence or absence of a particular feature with the feature of interest cued by the colored frame. At the third level, responses were based on whether 2 stimuli had matching features with the dimension of interest cued by the colored frame. At the fourth level, color-to-dimension mappings were manipulated by a cue presented at the beginning of a block. According to the framework of Koechlin and co-workers, the fourth level should have produced activations in the rostral lateral PFC (area 46) as the cue that determined color-to-dimension mappings provided episodic control. Instead, the fourth level produced activations in the frontopolar cortex (BA 10) corresponding closely to activations putatively involved in pending, but not current episodic signals (Koechlin et al. 1999; Koechlin and Hyafil 2007). Even more puzzling, rostral lateral PFC activations (area 46) were elicited by the third level, which provided a manipulation of relational abstraction, but not temporal abstraction. Hence, while the activations found in Badre and D'Esposito (2007) approximated those of Koechlin and coworkers, the roles of each region were inconsistent with respect to temporal abstraction. Further complicating this matter is a recent study by Reynolds et al. (2012) which compared effects of temporal abstraction and rule complexity. These authors found that midlateral (area 9/46) and caudal (area 6) regions of the frontal lobes were sensitive to both temporal abstraction and rule complexity. These data are in stark contrast with the present findings that no area of the PFC was sensitive to temporal abstraction. The inconsistency with how similar PFC regions respond to temporal abstraction across studies suggests that the PFC is sensitive to some factor that is only partly captured by the concept of temporal abstraction.

One difference between the present study and previous studies that have found effects of temporal abstraction is the explicit separation of integration demands in the present work. Koechlin et al. (2003) used a blocked design that did not permit separate assessment of episodic control signals from the integration of cue information with probe stimuli. Effects of temporal abstraction reported in Badre and D'Esposito (2007) and Reynolds et al. (2012) were uncovered through similar block analyses that confounded temporal abstraction and integration. Hence, given the present data, it

seems likely that effects of temporal abstraction in previous studies were due to integration rather than temporal abstraction per se.

Relational Abstraction Revisited

Could effects of relational abstraction also be explained by integration demands? In previous research, [Bunge et al. \(2003\)](#) observed activations in the lateral PFC to cues denoting a relational rule relative to cues denoting a nonrelational rule. As these effects were separated from the probe, the activations could not be explained by feature integration itself. Likewise, the present study also found lateral PFC activations independent of both probe-related integration and cue integration. Hence, relational abstraction effects are distinct from feature or cue integration.

These data are consistent with the hypothesis that the lateral PFC is involved in the representation of rules ([Wallis et al. 2001](#); [Bunge et al. 2003](#); [Wallis and Miller 2003](#); [Bunge and Zelazo 2006](#)) independent of the implementation of the rules themselves. Interestingly, in the present data, effects of relational abstraction were sustained in the lateral premotor cortex (area 6), but transient in more rostral areas of the PFC. Consistent with these data, single-unit recordings in monkeys have demonstrated rule representation in both premotor cortex and lateral PFC, but with greater sustained effects in premotor cortex ([Wallis and Miller 2003](#)). Hence, the lateral PFC may be important for the initial encoding and translation of a cue into a rule, while the premotor cortex maintains rule information in a more concrete form in preparation for future responding ([Nee and Brown 2012b](#)). Notably, rules also denote a kind of integration: they represent the integration of a state and action. As a result, regions involved in relational abstraction may also perform an integrative function by way of linking states and actions into rules.

Cue Integration

We found effects of cue integration in the rostral-most and caudal-most areas of our search space. The rostral area was found in a region bordering the rostral-end of area 46 and the caudal-end of area 10. To avoid confusion, we have referred to this region as frontopolar cortex to distinguish it from the rostral lateral PFC region sensitive to relational abstraction. ROI analyses centered more definitively in frontopolar cortex also showed effects of cue integration indicating that integration effects in frontopolar cortex are robust to differences in the exact location examined. Activations similar to those observed here were reported in a study that examined the integration of a number held in working memory into ongoing mathematical operations ([De Pisapia et al. 2007](#)). Compared with a condition requiring holding a number in mind for future recall, the need to integrate a number into ongoing operations produced a ramp up of activation in frontopolar cortex beginning shortly after encoding the to-be-integrated number and peaking at the time of integration. In addition, a study using the same task demonstrated that stimulating the frontopolar cortex affected the integration condition, but not closely matched conditions that did not require integration, indicating a causal role of frontopolar cortex in integration ([De Pisapia et al. 2012](#)).

Under what conditions are the integrative operations of the frontopolar cortex required? [Ramnani and Owen \(2004\)](#) have

suggested that the frontopolar cortex integrates the results of 2 or more separate cognitive operations when a single rule is insufficient to guide behavior. A critical feature of their account is that inputs into the frontopolar cortex are abstracted representations arising from supramodal cortex including the PFC and temporal pole. This assumption is supported by anatomical connectivity patterns ([Petrides and Pandya 2007](#)) and distinguishes the frontopolar cortex from other PFC regions that are connected with more downstream portions of cortex. Furthermore, the integrative role of the frontopolar cortex is supported by its increased dendritic spine density relative to other regions of cortex ([Jacobs et al. 2001](#)) providing a morphological basis to synthesize diverse inputs. This framework accounts well for the various roles of the frontopolar cortex during prospective memory ([Burgess et al. 2003](#)), branching ([Koechlin et al. 1999](#); [Charron and Koechlin 2010](#)), subgoal processing ([Braver and Bongiolatti 2002](#)), representing the value of alternative courses of action ([Boorman et al. 2009](#)), and exploring alternatives ([Daw et al. 2006](#)) in addition to relational abstraction ([Bunge et al. 2005](#); [Wendelken et al. 2008](#); [Wendelken and Bunge 2010](#)). In all of these cases, abstracted representations must be integrated into ongoing cognition in order to guide performance. Such representations differ from more concrete representations held in mind in typical delayed-match-to-sample working memory tasks that commensurately engage caudal frontal areas more consistently than the frontopolar cortex ([Curtis and D'Esposito 2003](#); [Courtney 2004](#)). A notable feature of this account is that the frontopolar cortex does not represent abstracted representations themselves, but instead provides a means to coordinate and link information represented elsewhere. This is consistent with data demonstrating that intentions cannot be decoded from the frontopolar cortex, but that the frontopolar cortex shows increased coupling with intention-representing regions during storage ([Gilbert 2011](#)). Furthermore, this account provides a necessary endpoint for abstraction gradients. Without such an endpoint, it would be difficult to determine the apex of abstraction.

Activations in the caudal superior frontal sulcus also demonstrated sensitivity to integration. We suggest that this pattern reflects attention for integrated feature combinations. The caudal superior frontal sulcus has a well-established role in visuospatial attention ([Kastner and Ungerleider 2000](#); [Reynolds and Chelazzi 2004](#); [Moore 2006](#)) thought to reflect top-down representation of attentional priority ([Serences et al. 2005](#)). In the present study, the combination of cues determined the set of stimuli associated with a “yes” response (e.g., 2 green stimuli). Hence, after presentation of the second cue, subjects knew what stimuli to look for in order to produce a “yes” response. Given its role in top-down attentional priority, it is likely that the caudal superior frontal sulcus represented this search set. Further, we suggest that the ramp up of activation in the caudal superior frontal sulcus just prior to the onset of the second cue reflects the preparation to encode the second cue. Thus, activations in the caudal superior frontal sulcus likely reflect attentional preparation. Note, unlike the frontopolar cortex, we do not suggest that the caudal superior frontal sulcus is necessarily involved in integration. Instead, we propose that integration enabled the concrete representation of the search set, the attention to which elicited activations in the caudal superior frontal sulcus.

An Integrative Framework

While we have argued that the PFC is involved in integration generally, it is clear that there are regional differences in integration sensitivity. What accounts for these differences? Recent theories propose that progressively rostral areas of the PFC represent increasingly abstracted content (Badre 2008; Badre and D'Esposito 2009; O'Reilly 2010). In the parlance of integration, more rostral areas integrate more concrete instances—represented caudally—into an abstracted representation. In other words, representations in more rostral areas form generalizations over caudal areas. We believe that such frameworks provide a useful means to understand the present data (see Supplementary Fig. 3). Starting caudally, we have suggested that the caudal superior frontal sulcus represents attention for specific visual instances, which are composed of the integration of particular features (e.g., 2 green items). That is, the caudal superior frontal sulcus maintains attention for a particular environmental state. Rostral to the caudal superior frontal sulcus, the lateral premotor cortex represents rules, which correspond to state-response associations. Given that the lateral premotor cortex was sensitive to relational abstraction even when particular features were unspecified, it seems likely that rules are encoded more abstractly than the representations of the caudal superior frontal sulcus. For example, the lateral premotor cortex may associate [2 features->“yes”] for an “and” cue. By contrast, the rostral lateral PFC may represent the relations more abstractly. This area may be important in encoding the abstract cues and translating them into rules that can be encoded by the lateral premotor cortex (e.g., [“and”->[2 features->“yes”]]). Finally, the frontopolar cortex may coordinate the integration of abstract representations enabling the ability to perform complex, multistep rules. Thus, different rostral-caudal areas of the PFC can be distinguished by the abstractness of the representations that they integrate.

Further Considerations

While we have concluded that the PFC is not sensitive to temporal abstraction over-and-above integration demands, additional experimentation is needed to fully address this possibility. In the present study, control trials required the maintenance of cue information across a 9-s delay interval. Although experimental trials theoretically should have elicited greater degrees of temporal abstraction than control trials, these differences may not have been substantial enough to detect effects of temporal abstraction with fMRI. Systematic manipulation of the cue-to-probe interval may reveal activations elicited purely by temporal abstraction. However, such manipulations must be careful to disentangle temporal abstraction and integration, which are confounded if probe-related activation cannot be statistically distinguished from cue-related activation. The relevant intervals that minimize temporal abstraction demands may prove infeasible with fMRI and may require insights from more temporally resolved methodologies.

Conclusion

We have suggested that integration is a useful framework to describe the organization of the PFC. In this framework, more rostral areas of the frontal lobes integrate representations subserved by more caudal areas. We propose that viewing the

PFC in this way resolves a growing corpus of data that attribute various forms of abstraction to the PFC. We have demonstrated that in the absence of integration, temporal abstraction effects are abolished. By contrast, areas of the PFC are sensitive to relational abstraction since relations describe the integration of a state and action (i.e., rule). Frontopolar areas previously associated with temporal abstraction effects demonstrate activation only with respect to integration operations. These data highlight the importance of disentangling the numerous complex cognitive events that give rise to frontal activations in order to precisely identify the operations of the frontal lobes.

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.), F32 NS082069 (D.N.), 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. The US 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 US Government.

Notes

The authors thank Mark D'Esposito for helpful comments on a previous version of this manuscript. *Conflict of Interest:* None declared.

References

- Andersson JL, Hutton C, Ashburner J, Turner R, Friston K. 2001. Modeling geometric deformations in EPI time series. *Neuroimage*. 13:903–919.
- 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.
- Banich MT. 2009. Executive function: the search for an integrated account. *Curr Dir Psychol Sci*. 18:89–94.
- Boorman ED, Behrens TE, Woolrich MW, Rushworth MF. 2009. How green is the grass on the other side? Frontopolar cortex and the evidence in favor of alternative courses of action. *Neuron*. 62:733–743.
- Braver TS, Bongiolatti SR. 2002. The role of frontopolar cortex in subgoal processing during working memory. *Neuroimage*. 15:523–536.

- Bunge SA, Kahn I, Wallis JD, Miller EK, Wagner AD. 2003. Neural circuits subserving the retrieval and maintenance of abstract rules. *J Neurophysiol.* 90:3419–3428.
- Bunge SA, Wendelken C, Badre D, Wagner AD. 2005. Analogical reasoning and prefrontal cortex: evidence for separable retrieval and integration mechanisms. *Cereb Cortex.* 15:239–249.
- Bunge SA, Zelazo PD. 2006. A brain-based account of the development of rule use in childhood. *Curr Dir Psychol Sci.* 15:118–121.
- Burgess PW, Scott SK, Frith CD. 2003. The role of the rostral frontal cortex (area 10) in prospective memory: a lateral versus medial dissociation. *Neuropsychologia.* 41:906–918.
- Cabeza R, Nyberg L. 2000. Imaging cognition II: an empirical review of 275 PET and fMRI studies. *J Cogn Neurosci.* 12:1–47.
- Charron S, Koechlin E. 2010. Divided representation of concurrent goals in the human frontal lobes. *Science.* 328:360–363.
- Christoff K, Gabrieli JD. 2000. The frontopolar cortex and human cognition: evidence for a rostrocaudal hierarchical organization within the human prefrontal cortex. *Psychobiology.* 28:168–186.
- Christoff K, Keramatian K. 2007. Abstraction of mental representations: theoretical considerations and neuroscientific evidence. In: Bunge SA, Wallis JD, editors. *Perspectives on rule-guided behavior.* New York: Oxford University Press. p. 107–126.
- Christoff K, Keramatian K, Gordon AM, Smith R, Madler B. 2009. Prefrontal organization of cognitive control according to levels of abstraction. *Brain Res.* 1286:94–105.
- Christoff K, Prabhakaran V, Dorfman J, Zhao Z, Kroger JK, Holyoak KJ, Gabrieli JD. 2001. Rostrolateral prefrontal cortex involvement in relational integration during reasoning. *Neuroimage.* 14:1136–1149.
- Christoff K, Ream JM, Geddes LP, Gabrieli JD. 2003. Evaluating self-generated information: anterior prefrontal contributions to human cognition. *Behav Neurosci.* 117:1161–1168.
- Courtney SM. 2004. Attention and cognitive control as emergent properties of information representation in working memory. *Cogn Affect Behav Neurosci.* 4:501–516.
- Curtis CE, D'Esposito M. 2003. Persistent activity in the prefrontal cortex during working memory. *Trends Cogn Sci.* 7:415–423.
- Damasio AR, Anderson SW. 1993. The frontal lobes. In: Heilman KM, Valenstein E, editors. *Clinical neuropsychology.* 3rd ed. New York: Oxford University Press. p. 409–460.
- Daw ND, O'Doherty JP, Dayan P, Seymour B, Dolan RJ. 2006. Cortical substrates for exploratory decisions in humans. *Nature.* 441:876–879.
- De Pisapia N, Sandrini M, Braver TS, Cattaneo L. 2012. Integration in working memory: a magnetic stimulation study on the role of left anterior prefrontal cortex. *PLoS one.* 7:e43731.
- De Pisapia N, Slomski JA, Braver TS. 2007. Functional specializations in lateral prefrontal cortex associated with the integration and segregation of information in working memory. *Cereb Cortex.* 17:993–1006.
- Duncan J, Owen AM. 2000. Common regions of the human frontal lobe recruited by diverse cognitive demands. *Trends Neurosci.* 23:475–483.
- Genovese CR, Lazar NA, Nichols T. 2002. Thresholding of statistical maps in functional neuroimaging using the false discovery rate. *Neuroimage.* 15:870–878.
- Gilbert SJ. 2011. Decoding the content of delayed intentions. *J Neurosci.* 31:2888–2894.
- Jacobs B, Schall M, Prather M, Kapler E, Driscoll L, Baca S, Jacobs J, Ford K, Wainwright M, Trembl M. 2001. Regional dendritic and spine variation in human cerebral cortex: a quantitative golgi study. *Cereb Cortex.* 11:558–571.
- Kastner S, Ungerleider LG. 2000. Mechanisms of visual attention in the human cortex. *Annu Rev Neurosci.* 23:315–341.
- Koechlin E, Basso G, Pietrini P, Panzer S, Grafman J. 1999. The role of the anterior prefrontal cortex in human cognition. *Nature.* 399:148–151.
- 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.
- Koechlin E, Summerfield C. 2007. An information theoretical approach to prefrontal executive function. *Trends Cogn Sci.* 11:229–235.
- Kouneiher F, Charron S, Koechlin E. 2009. Motivation and cognitive control in the human prefrontal cortex. *Nat Neurosci.* 12:939–945.
- Kroger JK, Sabb FW, Fales CL, Bookheimer SY, Cohen MS, Holyoak KJ. 2002. Recruitment of anterior dorsolateral prefrontal cortex in human reasoning: a parametric study of relational complexity. *Cereb Cortex.* 12:477–485.
- Moore T. 2006. The neurobiology of visual attention: finding sources. *Curr Opin Neurobiol.* 16:159–165.
- Nee DE, Brown JW. 2012a. Dissociable frontal-striatal and frontoparietal networks involved in updating hierarchical contexts in working memory. *Cereb Cortex.* [Epub ahead of print].
- Nee DE, Brown JW. 2012b. Rostral-caudal gradients of abstraction revealed by multi-variate pattern analysis of working memory. *Neuroimage.* 63:1285–1294.
- Nee DE, Brown JW, Askren MK, Berman MG, Demiralp E, Krawitz A, Jonides J. 2013. A meta-analysis of executive components of working memory. *Cereb Cortex.* 23:264–282.
- Nichols T, Brett M, Andersson J, Wager T, Poline JB. 2005. Valid conjunction inference with the minimum statistic. *Neuroimage.* 25:653–660.
- O'Reilly RC. 2010. The What and How of prefrontal cortical organization. *Trends Neurosci.* 33:355–361.
- Petrides M, Pandya DN. 2007. Efferent association pathways from the rostral prefrontal cortex in the macaque monkey. *J Neurosci.* 27:11573–11586.
- Ramnani N, Owen AM. 2004. Anterior prefrontal cortex: insights into function from anatomy and neuroimaging. *Nat Rev Neurosci.* 5:184–194.
- Reynolds JH, Chelazzi L. 2004. Attentional modulation of visual processing. *Annu Rev Neurosci.* 27:611–647.
- 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.
- 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.
- Wallis JD, Anderson KC, Miller EK. 2001. Single neurons in prefrontal cortex encode abstract rules. *Nature.* 411:953–956.
- Wallis JD, Miller EK. 2003. From rule to response: neuronal processes in the premotor and prefrontal cortex. *J Neurophysiol.* 90:1790–1806.
- Wendelken C, Bunge SA. 2010. Transitive inference: distinct contributions of rostral lateral prefrontal cortex and the hippocampus. *J Cogn Neurosci.* 22:837–847.
- Wendelken C, Nakhachenko D, Donohue SE, Carter CS, Bunge SA. 2008. "Brain is to thought as stomach is to ??": investigating the role of rostral lateral prefrontal cortex in relational reasoning. *J Cogn Neurosci.* 20:682–693.
- Wood JN, Grafman J. 2003. Human prefrontal cortex: processing and representational perspectives. *Nat Rev Neurosci.* 4:139–147.
- Zarahn E, Aguirre G, D'Esposito M. 1997. A trial-based experimental design for fMRI. *Neuroimage.* 6:122–138.